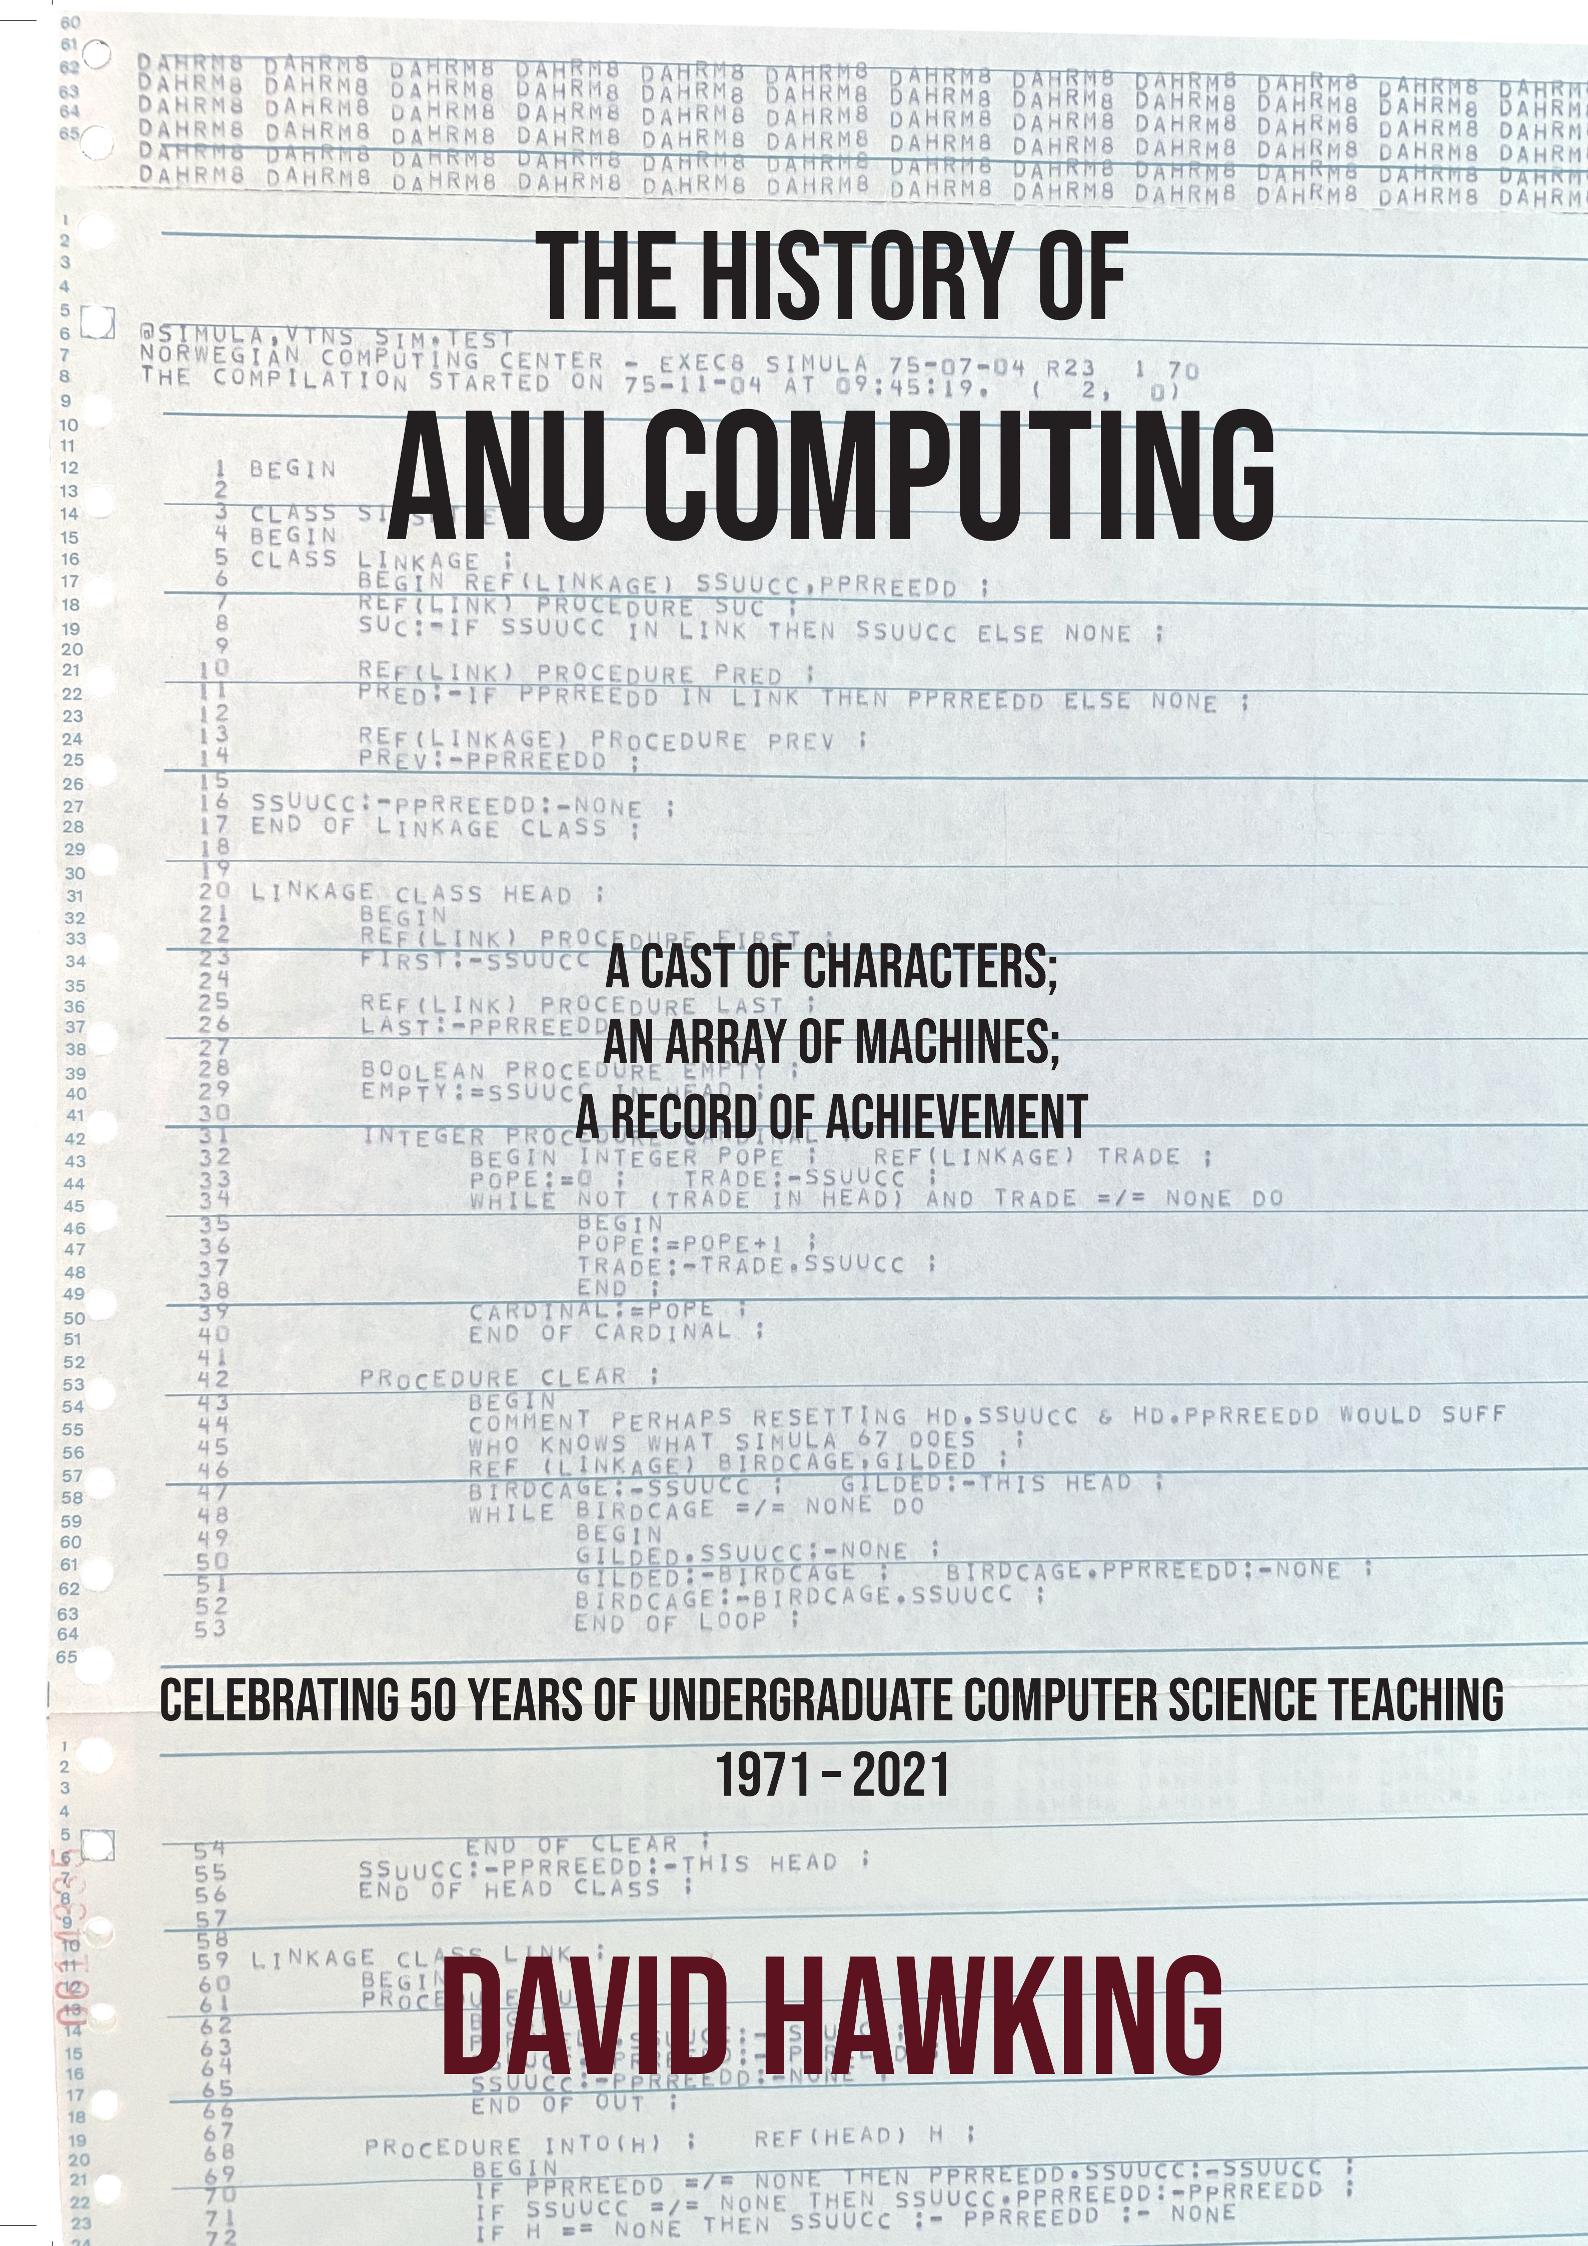

```
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM8 DAHRM
```

THE HISTORY OF ANU COMPUTING

```
@SIMULA,VTNS SIM.TEST
NORWEGIAN COMPUTING CENTER - EXEC8 SIMULA 75-07-04 R23  1 70
THE COMPILATION STARTED ON 75-11-04 AT 09:45:19.  (  2,  0)

BEGIN

CLASS SI
BEGIN
CLASS LINKAGE ;
      BEGIN REF(LINKAGE) SSUUCC,PPRREEDD ;
      REF(LINK) PROCEDURE SUC ;
      SUC:-IF SSUUCC IN LINK THEN SSUUCC ELSE NONE ;

      REF(LINK) PROCEDURE PRED ;
      PRED:-IF PPRREEDD IN LINK THEN PPRREEDD ELSE NONE ;

      REF(LINKAGE) PROCEDURE PREV ;
      PREV:-PPRREEDD ;

SSUUCC:-PPRREEDD:-NONE ;
END OF LINKAGE CLASS ;


LINKAGE CLASS HEAD ;
      BEGIN
      REF(LINK) PROCE
      FIRST:-SSUUCC

      REF(LINK) PROCEDURE LAST ;
      LAST:-PPRREEDD

      BOOLEAN PROCEDURE EMPTY ;
      EMPTY:=SSUUCC IN HEAD ;

      INTEGER PROC
            BEGIN INTEGER POPE ;   REF(LINKAGE) TRADE ;
            POPE:=0 ;    TRADE:-SSUUCC ;
            WHILE NOT (TRADE IN HEAD) AND TRADE =/= NONE DO
                  BEGIN
                  POPE:=POPE+1 ;
                  TRADE:-TRADE.SSUUCC ;
                  END ;
            CARDINAL:=POPE ;
            END OF CARDINAL ;

      PROCEDURE CLEAR ;
            BEGIN
            COMMENT PERHAPS RESETTING HD.SSUUCC & HD.PPRREEDD WOULD SUFF
            WHO KNOWS WHAT SIMULA 67 DOES  ;
            REF (LINKAGE) BIRDCAGE,GILDED ;
            BIRDCAGE:-SSUUCC ;    GILDED:-THIS HEAD ;
            WHILE BIRDCAGE =/= NONE DO
                  BEGIN
                  GILDED.SSUUCC:-NONE ;
                  GILDED:-BIRDCAGE ;    BIRDCAGE.PPRREEDD:-NONE ;
                  BIRDCAGE:-BIRDCAGE.SSUUCC ;
                  END OF LOOP ;
```

A CAST OF CHARACTERS;
AN ARRAY OF MACHINES;
A RECORD OF ACHIEVEMENT

CELEBRATING 50 YEARS OF UNDERGRADUATE COMPUTER SCIENCE TEACHING
1971 – 2021

```
            END OF CLEAR ;
      SSUUCC:-PPRREEDD:-THIS HEAD ;
      END OF HEAD CLASS ;


LINKAGE CLASS LINK ;
      BEGIN
      PROCE
            BEGIN
            SSUUCC:-PPRREEDD:-NONE ;
            END OF OUT ;

      PROCEDURE INTO(H) ;   REF(HEAD) H ;
            BEGIN
            IF PPRREEDD =/= NONE THEN PPRREEDD.SSUUCC:-SSUUCC ;
            IF SSUUCC =/= NONE THEN SSUUCC.PPRREEDD:-PPRREEDD ;
            IF H == NONE THEN SSUUCC :- PPRREEDD :- NONE
```

DAVID HAWKING

THE HISTORY OF
ANU
COMPUTING

A Cast of Characters
An Array of Machines
A Record of Achievement

Celebrating 50 Years of Undergraduate Computer Science Teaching
1971 – 2021

David Hawking
`david.hawking@anu.edu.au`

Text copyright © David Hawking 2021. Quotations and illustrations, as per individual acknowledgments.
The moral right of the author is asserted.

ISBN 978-0-6451743-0-4 (Print) – ISBN 978-0-6451743-1-1 (PDF) – ISBN 978-0-6451743-2-8 (EPub)

Typeset in 11pt Palatino using LaTeX. Last edited: 16 Mar 2022
Cover design: Jack Griffiths-Hawking
Back cover photo: 2007: The author (left) in a PhD supervisory session with Paul Thomas on Willows Oval. *Photo: Andrew Taylor / Age, Sydney Morning Herald, © Nine Publishing. Remarkably, Andrew threw the frisbee as well as taking the photo.*

This is quite a superb book! Congratulations! It's all that I expected, and more. I loved the scene-setting stuff in the 1970s! Every page brings a smile; it really is a disarming and affectionate epistle. You are upfront about it being a personal recount, and you have solicited input from people repeatedly.

Henry Gardner, former Head, DCS

What a tremendous work and so much detail I did not know about – I only worked there for 21 years!!

Robin Erskine, former Director, IT Services

Chapter 12 – You wrote in so powerful way about the contributions, successes, and struggles of our women that I think something in me has changed forever. Thank you, David, for capturing and connecting the female voices!

Hanna Suominen, School of Computing and Data61

Well, that was a trip down memory lane and was very interesting reading. I had forgotten some of the people you mentioned but it all came flooding back!

Barbara Bankovsky, former programmer, MISD

First of all – congratulations. You have captured an era before it disappeared! Another decade and it may not have been possible.

Brian Molinari, former Head, DCS

Well done – I really enjoyed reading this especially as I got to meet some of the players.

Ron Morrison, Professor, University of St Andrews, Scotland

Thanks for that – a fab read! It certainly helped me understand some of the history of the FCU prior to my arrival.

Melanie Rooney, former Acting Head / Team Lead, Faculties Computer Unit

Chapter 9 – that was a walk down memory lane! Great work!

Geoff Huston AM, Chief Scientist, APNIC

Chapter 10: Thanks for a great read. I'm looking forward to reading the book and thank you so much for taking the time and effort to write this important piece of ANU history.

Rick Van Haeften, former Director, Corporate Information Services

I have read the entire chapter and can't find anything that requires amendment. It is a very informative summary and I am most impressed by the details that you cover and more importantly, personally recount.

Bill Gibson, former CIO, Australian Tax Office

Great reading! I like the relaxed style very much. Brought back a lot of memories.

Bob McKay, former Professor, Seoul National University, Korea

I had SO much fun reading it, and look forward to reading the finished piece.

Les Landau, former programmer, ANUCC

What a wonderful trip down memory lane, thank you!! And an excellent read to boot.

Roger Clarke, former Reader, Information Systems

Chapter 4: Wow, I think we need to hand over this chapter to every new student and staff member as a mandatory induction item with an exam in the end. Wow, I thought I knew our people well, but I learned so much more as them as all rounded individuals.

Hanna Suominen, School of Computing and Data61

Foreword

Twenty-five years ago I was lucky enough to start as a student in the ANU Department of Computer Science (DCS). At the time, I had no idea that I'd end up spending the better part of a decade studying there, but I knew immediately it was a special place.

DCS made you feel like you were in the middle of it all. From the first steps through the front door you were reminded of where you stood in history, walking past a room filled with computers of the past (a 'Computer History Museum'), and passing a neon-lit glass cylinder artfully filled with electronic components of the past (now definitely serving the purpose of Art not Logic). This was a discipline and a department where much had happened, and where there was much more to come.

Reading through this book it struck me just how much the department was not about all the technology and formal lectures, but rather, the people. The lecturers in the department all had unique roles in the formation of Computer Science as a discipline. Stories were told around the department of a lecturer who wrote 'Australia's first laser printer driver', or around that time Microsoft tried to 'annex' Samba FS, or one professor's work on bible code debunking. Our lecturers not only had these experiences, they humbly shared them with us as well. I remember an inspiring advanced algorithms lecture based on a graph scribbled on the back of an envelope, and one time when a harried lecturer walked in a little disheveled and told the class "one buffer overflow, and a million people think you are an idiot." Another lecturer paused our lectures every 15 minutes to tell us a campfire-esque story about a disaster caused by poor software engineering. These kinds of stories continually reminded us that these weren't just teachers enabling a whole new generation of CS engineers, they were working with us to shape the future of Computer Science.

The teaching in DCS engendered enthusiasm and involvement from students, and occasionally kept us in the labs all night long. This infectious enthusiasm spanned everything from logic design and kernel hacking, to group software engineering projects. It also reflected what was going on in the wider world of computer science. In our final year our class embarked on a software engineering project to build a news aggregator ('My News Network'); which was clearly formative since I ran Google News 15 years later! The lab also had strong partnerships with Research and Industry through being nestled next to the CSIRO's Department of Mathematical and Information Sciences (CMIS) and the Research School of Information Sciences (RSISE). These partnerships were hugely impactful in my life, by providing access to work in research that had quantifiable impact. Working with Dave's Panoptic Search team in CMIS is where my love of Search and Information Retrieval started, and there's no doubt that without the unique opportunity I'd not be where I am today.

One of the passages I was thrilled to read is how the diversity within the Department has improved. In my day, out of a class of 300 people, I recall only 5 women; moreover in my Honours year every student was male. That was the year when the inaugural Women's Paul Thistlewaite Honours scholarship was awarded, however there were no female candidates and so it came to me. While I was honoured by and gratefully received it, it has been very heartening to hear that programs and awards like this are now greatly helping to encourage more student diversity.

This book anchors ANU's DCS in a larger world of immense excitement, disruption and opportunity. In this book Dave manages to weave together a unique collection of stories, experiences, and occasional hijinks that pulls together this one-of-a-kind department, located in a top university with world-ranked computing facilities.

Trystan Upstill
Vice President of Engineering, Google
07 May, 2021

Contents

Acknowledgements

I gratefully acknowledge the valuable contributions of Chris Johnson, Brian Molinari, Robin Erskine and Keith France, who made available photographs and many primary sources such as early Annual Reports and board minutes. Chris also recorded interviews with John Hurst and Vicki Peterson and contributed to memory-raising discussions with me and John Hurst. Thanks are due to John for conveying his keenness to have a history written, and for sharing memories.

I am in debt to the many people who have contributed material and/or engaged in useful discussions: Philip Allnutt, Brian Anderson, Ann Apthorp (née Lorrel Sherar), Peter Bailey, Julie Bakalor, David Baldwin, Barbara Bankovsky, Amanda Barnard, Nick Barnes, Tony Barry, Chris Barter, Genevieve Bell, Anthony Bennett, David Berriman, Jan Bitmead, Steve Blackburn, Lindsay Botten, Clive Boughton, Richard Brent, Tim Brook, Wray Buntine, Mike Carden, Peter Christen, Roger Clarke, Robert Cohen, Michael Cook, Brian Corrie, Nick Craswell, Barbara Davidson, John Dravnieks, Roger Edberg, Bob Edwards, Peter Elford, Robin Erskine, Peter Farmer, Hugh Fisher, Shayne Flint, Keith France, Henry Gardner, Bill Gibson, Bob Gingold, Stephen Gould, Gopal Gupta, Kathy Handel, Tony Hosking, Geoff Huston, Ruth Jarvis, Lynette Johns-Boast, Stuart Kendall, Yvonne Kulesza, Hanna Kurniawati, Rika Kobayashi, Thena Kyprianou, Les Landau, Hongdong Li, John Lloyd, Siew-Gim McGregor, Brendan McKay, Paul Mackerras, Rob Mahony, John McGee, George McLaughlin, Bruce Millar, Brian Molinari, Ron Morrison, Malcolm Newey, John O'Callaghan, Emelyn Pecar, Mark Reed, Kathy Reid, Mark Reid, Alistair Rendell, Jochen Renz, Brian Robson, Mike Robson, Melanie Rooney, Ramesh Sankaranarayana, Scott Sanner, Suthagar Seevaratnam, John Slaney, Robin Stanton, Peter Strazdins, Hanna Suominen, Kerry Taylor, Sylvie Thiébaux, Paul Thomas, Andrew Tridgell, Jochen Trumpf, Rick Van Haeften, Trevor Vickers, Robert Whaley, Allan Williams, Graham Williams, Bob Williamson, Peter Wishart, Lexing Xie and Alex Zelinsky.

There are of course many others who could have contributed. Perhaps I didn't know of your role in the history; perhaps I didn't know how to contact you; perhaps you weren't able to respond by my deadline.

Thanks to Richard Brent, Robin Erskine, Robin Stanton, and the two dozen "beta testers" for reviewing draft chapters. Thanks to Kathy Griffiths for reviewing the whole thing, and for contributing as a very skilled researcher. Thanks to everyone who took the time to meet me – those chats were a most enjoyable part of what turned out to be a substantial project. Thanks to Beth Lonergan and the ANU Archives for very useful assistance. The ANU Archives is a critical resource for anyone writing histories or trying to learn from the past. It's regrettable that ANU departments have not been more diligent in depositing records, photographs, and reports.

I've tried to accurately acknowledge the sources of all photographs and quoted material. Please contact me if there are any issues.

Chapter 1

Introduction

I believe that the 50th anniversary of undergraduate computer science (CS) teaching at the Australian National University (ANU) fell on 02 March, 2021. It seems appropriate that the tumultuous half century of ANU computing, as a technology and as a discipline, should now be recorded.

Discussions about writing a history started around the 45th anniversary. Eventually, and with the support and encouragement of early members of the department, particularly Chris Johnson and John Hurst, I volunteered to write a personal account. Because that inevitably contains large gaps, I have incorporated large amounts of (duly attributed) material from former colleagues and others.

I have of course made use of primary sources held by individuals, and by the ANU Archives. Unfortunately, there are huge gaps in formal records. Indeed, ANU departments seem no longer to produce annual reports. I have found very few photographs indeed of early ANU computing equipment, a notable exception being an album of photographs collected by Keith France, titled *Computer Services Centre: A Pictorial History* which he is donating to the ANU Archives.

I have chosen to consider all aspects of computer science (and computing in general) at ANU, rather than restricting myself to the Department of Computer Science (DCS) or to teaching. Many other parts of ANU have contributed to the story of computing at ANU, including the Computer Centre (ANUCC), Vice-Chancellor's Computing Research Group (VCCRG), Computer Services Centre (CSC), Computer Sciences Lab (CSL), Automated Reasoning Project (ARP), the Centre for Information Science Research (CISR), the ANU Supercomputing Facility (ANUSF), the Department of Systems Engineering, the Department of Engineering Physics, the Research School of Information Science and Engineering (RSISE), the ACSys Cooperative Research Centre, the Faculty of Engineering and Information Technology (FEIT), National ICT Australia (NICTA), the Australian Partnership for Advanced Computation (APAC), the National Computational Infrastructure (NCI), the Information Systems group within the Faculty of Economics & Commerce, the 3A Institute, and the College of Engineering and Computer Science (CECS).

In 2021, computing is a vital part of most disciplines. Even outside computer science, ANU students use computers to obtain course materials from Wattle,[1] to attend virtual lectures and tutorials, to gather information via the web, to prepare assignments and presentations, to analyse data, to model systems, to perform calculations, to collaborate with fellow students, to submit assignments, and for many other purposes. Use of computers has become ordinary, expected, and mundane. Including smartphones (which are more powerful than early supercomputers), there may be many tens of thousands of computers on campus.

In 1971, it wasn't at all like that. There were only a handful of computers on campus and undergraduates were rarely permitted to touch them or their peripherals. A lucky few were able to prepare programs on cards, and submit them for processing. The establishment of a Department of Computer Science was controversial. "Is computer science really a science? Is it really a discipline worthy of being taught at a university?"

This history attempts to tell the story of how computer science and computing at ANU evolved from 1971 to 2021. It covers the evolution of the computer science curriculum, and the evolution of

[1] The Web Access to Teaching and Learning Environments system

the ANU organisational framework supporting computer science, and computing more generally. It charts the explosion in power and prevalence of computers on campus, and the radical changes in their interconnectedness. Importantly, it tells stories of many of the people involved in the gentle computing revolution at ANU. There is a bias toward the history of DCS. This is partly to appropriately address the 50th anniversary occasion, and partly because of limits to my knowledge and perspectives.

I start by painting a picture of the social and general context prevailing at ANU and in broader society in 1971 when DCS began. I step back in time to write about ANU's first computers and the people who operated and used them. I also attempt to document the current situation, in the hope that this will be of interest to old-timers and useful to the person writing the centenary history!

The history of computing at ANU, like that of most organisations, is dotted with occasional fights, enmities, trickery, empire building, and scrambles for resources. All that would make a very interesting story, but in the interests of a quiet life, I have left most of it out. I do remember Robin Stanton returning from a meeting one day:[2]

Brian: "How did the meeting go Robin?"

Robin: "Ah well. There seemed to be an outbreak of facial elongation."

Dave: "Facial elongation???"

Robin: "Yes, I wasn't sure if it was the Pinnochio effect, or snouts reaching for the trough!"

I've tried to make the story entertaining, but I have assumed that people reading a computer science history will have at least some interest in and knowledge of computers and computing.

1.1 Why Me?

Computer science was first taught as an undergraduate subject at ANU in 1971. I arrived as an undergraduate that year but was not able to study computer science because Computer Science I was only available as a second-year subject. Over considerable protest from the Arts Faculty, ANU was in the process of changing from year-long subjects, taught in three terms, to semesters – when I started computer science in 1972, it was no longer CS I, but CS B01 and B02.

Since then I've had a continuous but highly varied connection with ANU computer science, detailed in Appendix J, including as a student, teacher, programmer, researcher and honorary professor. I remember a lot of its early history because coming to university was very significant for me and I had a feeling that I was immersed in a wonderfully exciting time in the development of computing and information technology. I came from the country (Beechworth, Victoria) and was the first in my family to attend university. I could see so many exciting possibilities for computing!

Kind people have provided me with a lot of documents and photographs, and I've also kept many mementos from the early days, exemplified by my Christmas Eve payslip from 1975.

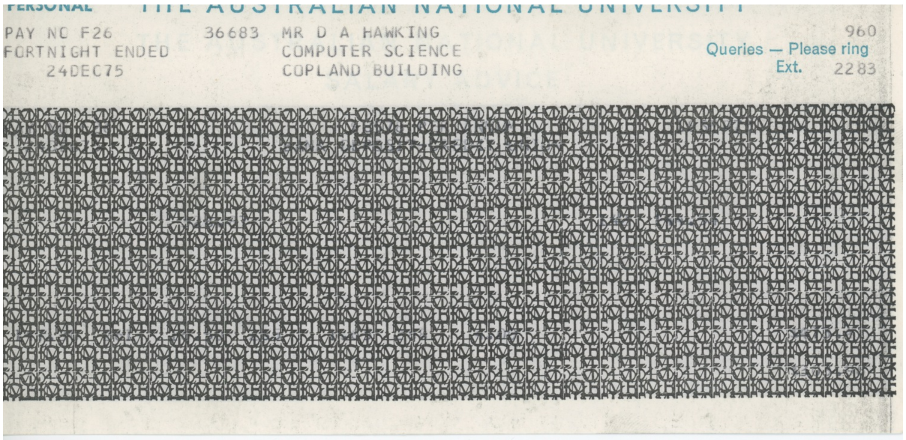

PERSONAL THE AUSTRALIAN NATIONAL UNIVERSITY

PAY NO F26 36683 MR D A HAWKING 960
FORTNIGHT ENDED COMPUTER SCIENCE
24DEC75 COPLAND BUILDING

Queries — Please ring
Ext. 2283

[2] Robin says he doesn't remember making this witty observation, and says that Brian doesn't remember it either. But it's such a good story that it has stuck in my memory.

1.2 Writing conventions

A problem with covering a period of more than half a century is that names and titles of people, names of buildings and names of organisational units change. I've tried to follow the following conventions:

- Except sometimes when quoting other people, I've avoided the use of honorifics;
- When people change their name, I generally use the name most commonly used while at ANU and, at least for the first reference, link to the other name using *formerly known as* or *later known as*;
- I apply the same approach to buildings;
- I usually refer to the organisational unit responsible for undergraduate computer science teaching as *DCS*, even when it was the Computer Science sub-department of the Department of Statistics, the School of Computer Science, the Research School of Computer Science or the School of Computing;
- I generally introduce acronyms at the first reference in each chapter.
- There are separate indexes to People, Hardware & Networks, and Software & Services. The content of the appendices is not indexed.

The text contains hundreds of links to web pages which cannot be guaranteed to remain active. On 30 March 2021 I checked all the links in the PDF by clicking on them using Preview on an Apple MacBook Pro, with the Firefox browser. All were active at the time of the check, except for a few referenced in the appendix on NCI, which was itself resurrected from the WayBack Machine.[3] In the event that a URL ceases to be active, I can only suggest searching in case the information is available at another web address, or looking for the original URL in the Wayback Machine.

1.3 Reporting Alleged Errors and Omissions

I fear that, despite my best intentions, errors and omissions remain. Inspired by Donald Knuth's practice regarding errors,[4] I propose the following mechanism for correcting errors. (Corrections may be applied in future editions.)

1. Carefully document the error or significant omission and your proposed remedy.
2. Write a cheque payable to David Hawking for the amount appropriate to the year in which you lodge your report. In 2021 the amount is $A327.68. Each calendar year after 2021, the amount halves. After 2036, you needn't send a cheque.
3. Post Items 1 and 2 to David Hawking, PO Box 6284, O'Connor ACT 2602. Don't forget a stamp!

The above is, of course, a joke. Feel free to email me at `david.hawking@acm.org`.

[3] `https://archive.org/web/`

[4] See Page 281 and `https://en.wikipedia.org/wiki/Knuth_reward_check` if you don't know what I'm talking about.

Chapter 2

Beginnings

The beginnings of computing, and computer science, at ANU marked a period of upheaval, even revolution, both in technology and in society.

In the 1960s, the lives of ordinary Australians were only slightly affected by computers. Some might have received machine-readable mark-sense cards or documents from the government, emblazoned with the instructions:

DO NOT FOLD, STAPLE, SPINDLE OR MUTILATE

Few would have realised how much computers and computer communication would come to dominate their lives.

I commence my history with a brief account of the societal and ANU context in which ANU's computing revolution began.

2.1 1970s: External Context

in 1971, Billy McMahon became prime minister, deposing John Gorton thanks to the latter's casting vote. In some states, women were not allowed in public bars,[1] and automatically lost government jobs when they married. Shamefully, right wing politicians stoked fear of "the Yellow Peril", and "the Red Menace". Indeed belief in the domino theory of communism had led to Australian involvement in the Vietnam War (1962 – 1972). Philip Baxter of the Australian Atomic Energy Agency[2] campaigned vigorously for Australia to build a full-scale nuclear reactor and an atomic bomb. In 1971 ANU's Director, RSPhysS was Ernest Titterton who had also worked on the Manhattan Project and supervised nuclear tests at Bikini Atoll (1946) and Maralinga (1956).

The Cold War was in full swing, and many citizens of Europe were fearful for their lives. Swiss houses were required to include fall-out shelters. At the same time, the environmental movement seemed to be gathering momentum around the world.

An Environment Society (EnSoc) was formed at ANU, with me as an office holder. The Club of Rome attempted to highlight the finiteness of the earth's resources and predicted that oil and various other critical resources would run out within a few decades – it actually turns out that Earth has too much coal, oil and gas to burn safely. The Zero Population Growth (ZPG) movement, founded by Paul Ehrlich and others, was influential at the time but seems to have failed.

Australian forces were fighting in Vietnam. Although unable to vote until 21, male Australians turning 20 were required to register for National Service. If their birthday was drawn out of a barrel,

[1] Roger Clarke recalls the barmaid in Sydney's Town Hall Hotel saying severely (in 1973), "We don't serve ladies in the public bar!"

[2] Philip Baxter was part of the Manhattan Project at Oak Ridge from 1944 until he was recruited by NSW University of Technology (later UNSW) in 1949 where he stayed, mostly as vice-chancellor, until 1969. He served as chairman of the Australian Atomic Energy Commission between 1957 and 1972, and as chairman of the International Atomic Energy Agency from 1969 to 1970. As well as founding Lucas Heights, he also founded NIDA and oversaw the completion of the Sydney Opera House.

they were required to serve two years in the military and could be sent to fight in Vietnam. ANU, like other universities, was a hot bed of draft resisters. A famous resister called Steve Padgham lived for a while in a kombi-van in the carpark next to the Chifley Library, and at other times with his girlfriend in females-only Ursula College.

A major headquarters for student activism was in a group house at 30 Canning St, Ainslie. Jack Waterford, who later became Editor of the Canberra Times, was Business Manager of Woroni,[3] and prominent among the activists.

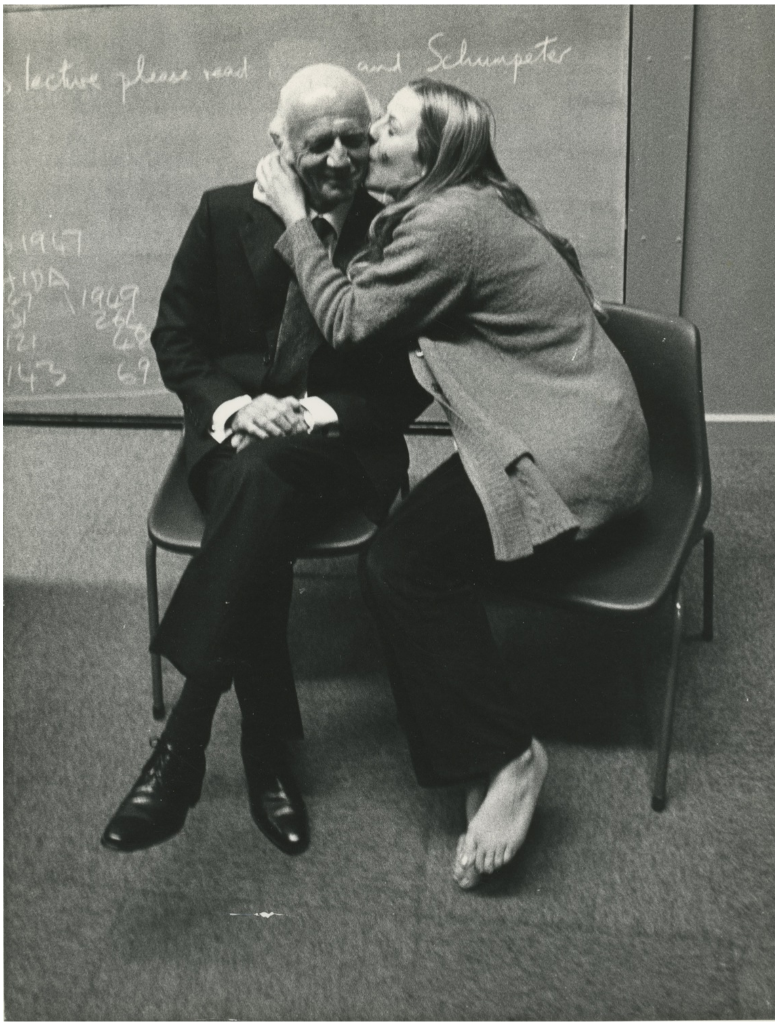

1972: Rigmor Berg welcomes Prime Minister Billy McMahon to the ANU campus.
Photo: Huw Price[4]

In 1971, the musical *Hair* was still playing, hippie culture was in full swing, and the counter-cultural Aquarius Festival of University Arts was held at ANU. (The next and last such festival was held in Nimbin.) Thousands of interstate students attended. Along with lots of music, highlighted by a massive dope riddled Daddy Cool concert in a marquee near the Chifley Library, there were large scale student protests, notably the "Day of Rage against Apartheid, Conscription and the Vietnam War." Immediately prior to this the McMahon government had passed the Public Order Act[5] which permitted a Sergeant of Police (or above) to order the dispersal of gatherings of twelve or more persons. Failure to disperse within 15 minutes, left you liable to up to 6 months jail.

[3] Woroni has been the ANU student newspaper since 1950. I co-edited an Environment issue with Rigmor Berg in 1972. `https://www.woroni.com.au/`

[4] Huw Price is an ANU Alumnus. He was Bertrand Russell Professor of Philosophy, Cambridge University from 2011 until his retirement in 2020.

[5] `https://www.legislation.gov.au/Details/C2011C00283`

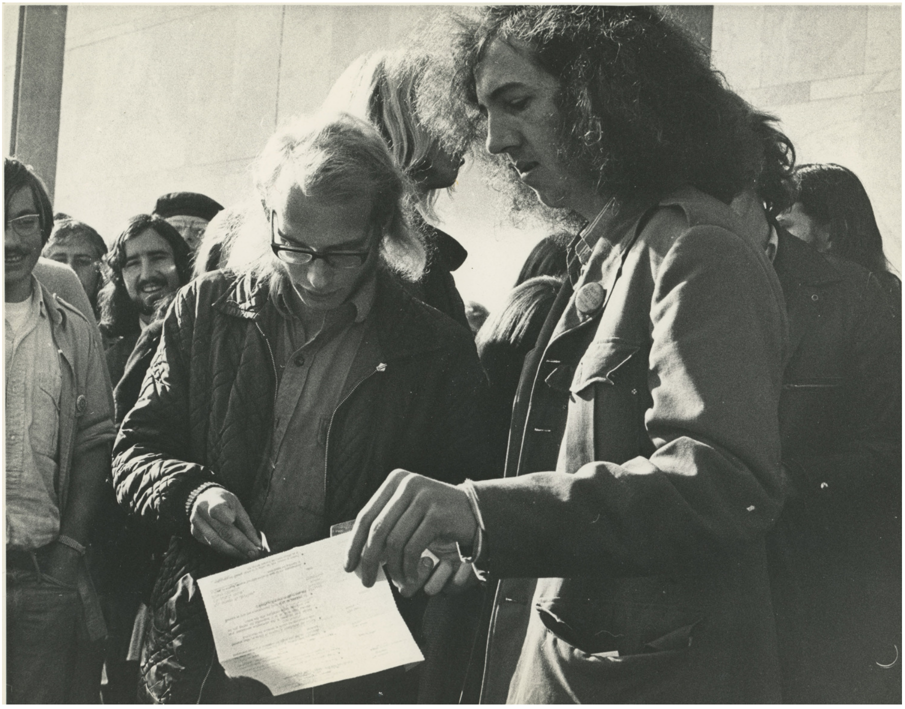

1971: Burning a call-up notice. *Photo: Huw Price*

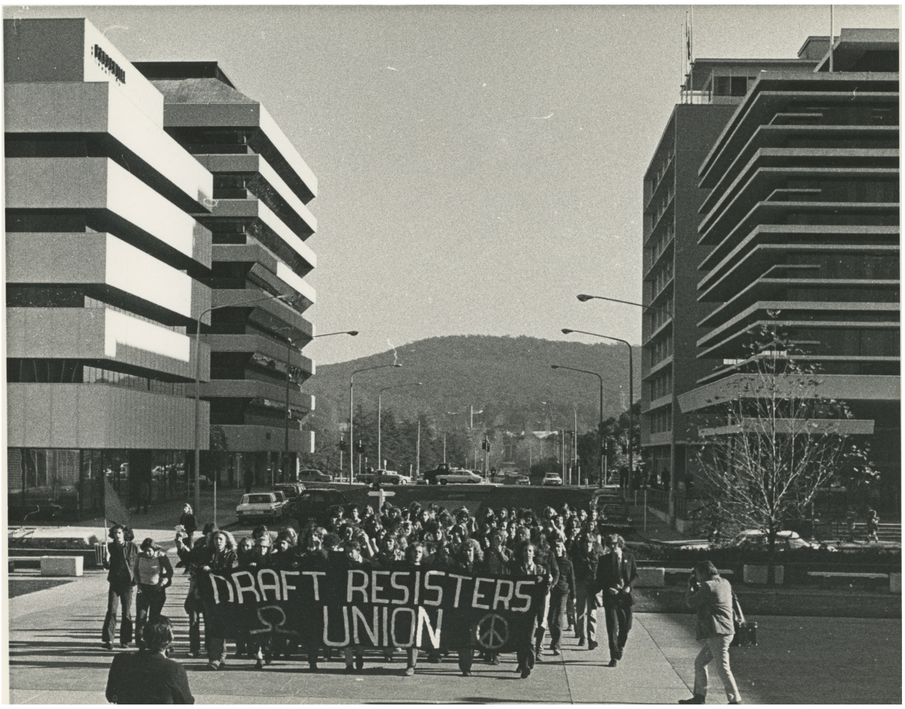

1971: Protest march in University Avenue. The front row includes: Jack Waterford, future Editor of the Canberra Times; Alan Gould, Canberra poet; Steve Padgham, draft resister and future Canberra teacher; Richard Refshauge, future ACT Director of Public Prosecutions and Justice of the ACT Supreme Court. *Photo: Huw Price*

On the Day of Rage a large crowd of students marched from ANU into Garema Place (not yet closed to traffic) and sat down in the road. As each student leader in turn attempted to address the crowd via megaphone they were arrested and taken off to a waiting police bus. Tiring of this, the crowd of demonstrators then marched and chanted their way through Civic, and across Commonwealth Avenue to protest at [Old] Parliament House, where arrests continued. Finally, the march returned to the city police station to demand the release of the arrestees.

Crowds such as this often gathered on the Library Lawn during the Aquarius Festival of University Arts held from 15 to 22 May. During the week a multitude of activities was conducted on campus in the biggest festival of its kind yet held.

Photo by courtesy of G. Carpay

Photo extracted from 1971 ANU Annual Report. **The report says that 6000 students came from out of town. That's 30% more than the then total number of ANU students. Note the newly constructed Melville Hall.**

The scene in the courtyard between the Police Station, the Supreme Court, and the Reserve Bank was frenetic and very tense. Police photographers and officers with eidetic imagery[6] were hanging out of the upper-floor windows of the police station, attempting to identify draft resisters in the crowd. A cordon of police hemmed in thousands of chanting, banner waving protestors. A senior police officer eventually read something through a megaphone (no doubt the direction to disperse) but no-one could hear it because of the noise. Finally, someone threw a missile at a police photographer, and the police charged. There ensued a stampede down University Avenue, all the way to Sullivans Creek, with police seemingly arbitrarily arresting students in the street.

My memory is that there were 186 arrests that day. Diana Riddell (after whom a building in Kambri[7] is now named) was Administrative Secretary of the Students Association and a great support for the arrested students. Some new students thought her name was, "Dira Dell".

Outside the Aquarius Festival, there was a great deal of student activism. Students and others mounted a constant 24-hour "Honk against Apartheid" vigil outside the pure-white South African Embassy on State Circle. (Capital Circle hadn't been completed at that stage so there was quite a lot of traffic and a lot of honking.) Campaigns to harrass the embassy staff were very active. Gravel was ordered from a supplier, to be dumped in the embassy driveway. Every public phone (remember them?) in the student union (then located in the architecturally-interesting but now-demolished building across the lawn from the Chifley Library) had a sign encouraging phone users to ring the embassy (numbers provided) and argue against apartheid.

The anti-Apartheid campaign reached its peak later that year, when the Springbok rugby team arrived to play a match at Manuka Oval. 750 NSW police were brought to Canberra and sworn in as special constables. Their train from Sydney received bomb threats and completed the journey at walking pace. The playing area at Manuka was surrounded with 2 metre high barbed wire fences, but despite this the ground was bombarded with flares, smoke devices and other missiles. Nothing

[6] Nowadays it's computers who have the eidetic imagery!

[7] Kambri is the new student facilities precinct along University Avenue just East of Sullivans Creek, which was opened in 2019.

could be heard from the field because of shouting and thousands of umpire whistles among the crowd. 47 demonstrators were arrested. It's still the only rugby match I have ever watched.

Sex was a hot topic in the early 1970s. The Democratic Labor Party led by Senator Vince Gair waged a long campaign against sex on university campuses. His opposition reached a climax when the ANU Orientation Week Handbook for 1971 featured a zodiac calendar of sexual positions.

At that time, contraception was taboo. It was illegal to advertise contraceptives or to display them publicly.[8] Despite this, an early issue of Woroni in 1971 featured a full page image, due to John Reid, of a doorstep with milk bottles full of condoms. The image was also published in *Nation Review*[9] In Orientation Week, 1971, buckets of condoms were thrown to an audience and a lecture on male homosexuality was given by a gay man. All this seemed incredibly daring and revolutionary.

In fact, male homosexuality was illegal throughout Australia until South Australia introduced the defence of "consenting adults in private", in 1972.[10] In Tasmania, homosexuality remained illegal until 1997, and a man was jailed there for homosexuality in 1984.[11]

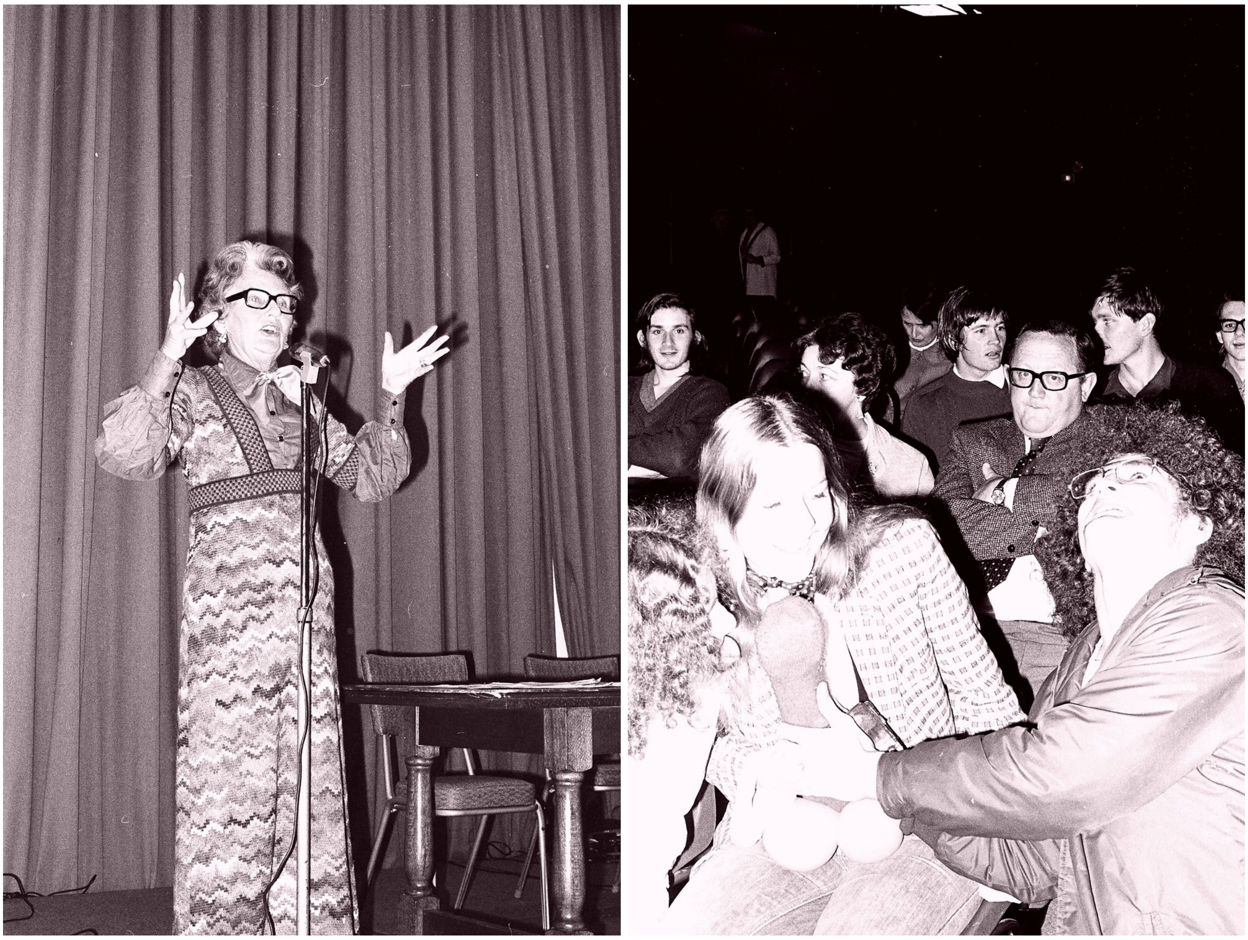

1973: L: Mary Whitehouse (of the Festival of Light) speaks in ANU's Melville Hall. R: In the audience, Rigmor Berg and Brutus O'Dowd convey a contrary message, to the apparent disapproval of a man in the row behind. ***Photos: Huw Price***

In 1971, many progressive people campaigned against the institution of marriage, and the vestiges of the system which treated a wife as her husband's chattel. In contrast, leading up to the referendum in 2017, many progressive people campaigned to *extend* the institution of marriage to include same-sex couples.

[8] Oddly, contraceptives, such as the pill, were subject to a 27.5% luxury tax. https://www.nma.gov.au/defining-moments/resources/the-pill#!

[9] For the young among you, in those days a "milko" used to deliver milk in glass bottles to your front door and take away the empties for refilling.

[10] https://en.wikipedia.org/wiki/LGBT_history_in_Australia

[11] https://en.wikipedia.org/wiki/Capital_punishment_for_homosexuality

When Mary Whitehouse visited the campus in 1973 to campaign for chastity and morality, the organiser, in blissful ignorance, told the audience how difficult it was to get her to come, and the front rows were occupied by students stroking large phalluses, or dressed in flasher raincoats.

The influence of Christian religion was still very strong in the 1970s, and Sunday sporting and commercial activities were very restricted. Until the Sydney Swans joined the AFL in the 1980s, all VFL/AFL matches were played on Saturday afternoons.[12] In many states, Sunday trading was prohibited for many types of business. Even in the 1980s, Frank Penhalluriack was sent to jail for opening his Victorian hardware shop on a Sunday.[13] In NSW, pubs which were "Licensed to sell fermented and spirituous liquors" could only serve alcohol on Sundays to "bonafide travellers", i.e., you were encouraged to drive in order to drink!

Driving on Australian roads in the 1970s was very dangerous. In 1971, 3590 people died in car crashes, a rate of 44 fatalities per billion kilometres driven. By 2018, the number of fatalities had fallen to 1135, and the rate had fallen by an order of magnitude.[14] Since 1971, wearing seatbelts became mandatory, automotive design rules tightened, airbags became standard, ABS brakes became commonplace, breath testing became common, speed limits were more strongly enforced, road rules and signs were harmonised across states, and vast sums were spent on improving roads. In 1971, there were two rail underpasses on the Hume Highway south of Sydney with a 5 mph (8 km/hr) speed limit. Almost none of the main routes between Canberra and Sydney and Sydney and Melbourne were dual carriageway. Now all of it is.

If an accident were reported on the TV news, it would be shown in black and white, because colour television wasn't introduced in Australia until 1974.

If you were a student and your car broke down on a country road you couldn't call your parents on your mobile – mobile what?? You had to hope some other motorist would give you a lift to a public phone box in the nearest town, which could be miles away. (Not kilometres because metrication of measurements in Australia only began in 1971 and wasn't complete until 1981.) Once at the phone box, you needed to know if the phone box was in the same area code as your parents'. If so, you could insert the necessary coins and dial the six digit local number. (For example, the main switchboard for ANU was 49 5111.) When the call was answered you pressed Button A, the machine swallowed your coins and the call was connected. For local calls you could talk for hours without further charge.

If, on the other hand, your parents' phone had a different area code, you dialled 0175 to connect to an operator. The operator would tell you how much it would cost for each period of 3 minutes. You would insert the money and, when directed, press Button A. If the call could not be connected, you pressed Button B to return the coins. As your allotted 3 minutes were about to expire, there were a series of audible "pips" and you could insert more money to extend the call. Such "trunk calls" were quite expensive, and motivated many clever schemes to avoid the charges.

In December, 1972 Gough Whitlam was elected as Australian Prime Minister after 23 years of Liberal-Country Party rule. He soon abolished conscription, abolished university tuition fees, ended Australia's involvement in the Vietnam war, and recognised the People's Republic of China. 1972 also saw the establishment of the Aboriginal Tent Embassy outside [old] Parliament House. It's still there.

I was rather grateful for Whitlam's election, since otherwise, two months after the election, I would have been required to decide whether to register for National Service. Unfortunately the Whitlam government came to a sad end. In November 1975, I was attending a training course on teaching methods, when a man obviously in a state of shock burst through the door, gasping, "They've sacked Whitlam!" Shock and outrage spread like wildfire across the campus.

1973 was most notable for an oil crisis, which hit Canberra badly. Trains stopped running, bus service was severely impacted, service stations ran out of fuel, and petrol was rationed. At the time oil heating was very common in Canberra, and oil tanks could be seen on the walls of many houses. The dramatic increase in oil price after 1973, caused a mass migration to other fuels.

[12] https://en.wikipedia.org/wiki/History_of_the_Australian_Football_League

[13] https://en.wikipedia.org/wiki/Frank_Penhalluriack

[14] https://en.m.wikipedia.org/wiki/List_of_motor_vehicle_deaths_in_Australia_by_year

In the early 1970s, ANU operated a large oil-fired boiler in what is now the Yencken Building. From there, underground steam pipes delivered heat to radiators in the older buildings on campus, such as Copland. After the oil shock, ANU decommissioned the boiler and converted the heat exchanger in the Copland building to use about 1000 electric elements, which were automatically controlled to minimise cost. All elements were switched on during the early morning in winter, and progressively switched off as ANU employees came to work, thus minimising ANU's peak demand.

2.2 1970s: ANU Context

Throughout the 1970s and beyond, ANU was an amalgamation of two distinct parts, the Institute of Advanced Studies (IAS) and the School of General Studies (SGS). IAS was the original ANU, established largely to attract Australian research talent back home from Britain and the USA, while SGS developed from the Canberra University College (CUC) which had amalgamated with ANU in 1960. CUC was an outpost of the University of Melbourne. It taught only Bachelor of Arts students and included Psychology, because Psychology was a science subject, and University of Melbourne BA degrees were required to include one. IAS was funded by a block grant from the Commonwealth, and had double the budget of SGS. The IAS and the SGS shared the costs of facilities such as libraries and the Computer Centre.

In 1971, the Chancellor was H.C. Coombs, the Vice-Chancellor was John Crawford, and David Dunbar was Deputy Vice-Chancellor (DVC). The 1971 Annual Report does not list other senior officers, but my memory is that key players in the Chancelry were Ross Hohnen, George Hamilton, David Hodgkin, George Dicker, Colin Plowman, Jane Flecknoe and Mollie Bouquet. Richard Refshauge was president of the ANU Students Association.

In the ANU Report of the Council for 1971[15] the Vice-Chancellor argues that the size of ANU would appropriately stabilise at a level of 7200 full-time-equivalent students, including 1300 postgraduates. He thought that the minimum viable size was 5000. In 1971, the number of students enrolled in the School of General Studies was 4454, 2478 full-time and 1976 part-time. This was considerably below the stated desirable minimum. Of the total, 450 were enrolled for postgraduate degrees. In addition, 362 students were enrolled in the Institute of Advanced Studies. More than 60% of the student load came from within 40km of the University.

The University Librarian was Jack Graneek, who is reported to have personally inspected every volume received by the library until shortly before his retirement in 1972.

ANU Creative Arts Fellows for 1971/72 included some famous names: Stanislav Ostoja-Kotkowski (lasers), Arthur Boyd (paintings, ceramics and tapestries), Don Banks (music) and Norma Redpath (sculpture).

ANU employed a Travel Clerk (whose name I temporarily forget). Yes, a single travel clerk to handle the university's travel arrangements. She was subject to inducements from airlines and other travel providers. ("Let us provide you with a free flight and accommodation so you can confirm for yourself that your staff will be flying with the required level of support and convenience.") An edict came from a senior officer that no ANU staff were to accept freebies from travel companies, but when Qantas launched its non-stop 747 service to Los Angeles and offered a free seat, the Travel Clerk was unable to resist. Lining up for the toilet, she found herself next to the man who'd forbidden travel freebies.

– "What are you doing here???"

– "Freeloading, just like you."

[15] https://openresearch-repository.anu.edu.au/bitstream/1885/14782/1/AnnualReport-1971.pdf

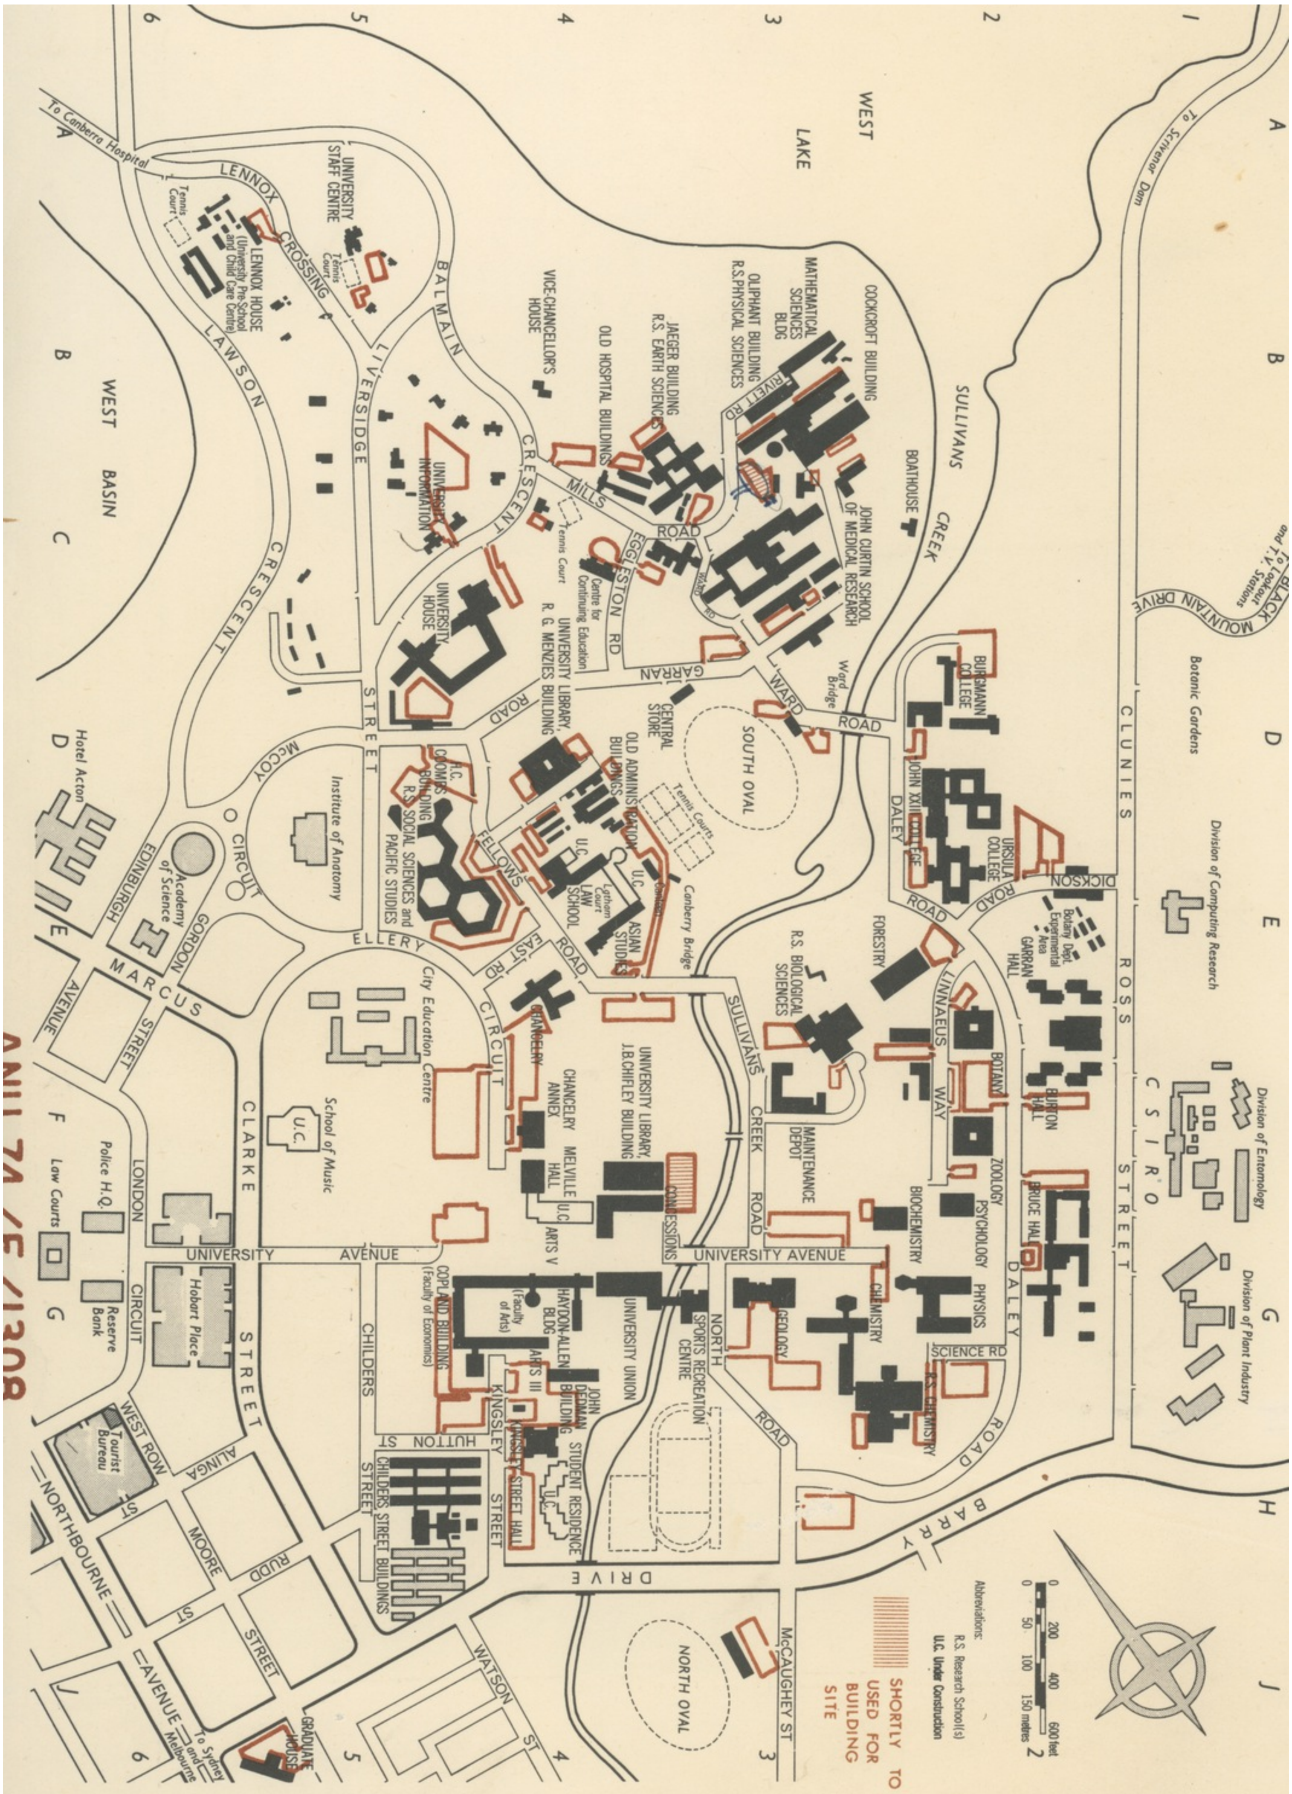

1974 Map of the ANU campus. The past 46 years have seen dramatic changes! This map shows a break in University Avenue. In 1971, there was an unimpeded road from Bruce Hall to the forecourt of the Supreme Court.

Computer companies in the 1970s and 1980s had reasonable budgets for wining and dining customers. If you had a position which might conceivably be regarded as having influence on university spending, you were likely to be invited to business lunches reasonably often. Lunches were sometimes an opportunity to talk about the university's requirements and to find out more about the company's thinking and directions. Sometimes they were purely social, an opportunity for junior vendor staff and university staff to get to know each other, and have an up-market free lunch. I was never subject to dishonest inducements or to overt pressure.[16] However, Peter Grimshaw, Business Manager of the Joint Schools told me that he had a blanket rule – never accept anything, however small, from anyone who might do business with the University.

By far the best freebie I received was when Control Data Corporation (CDC) organised a full-day launch of its products in the Academy of Science. Federal Parliament had unexpectedly decided to sit, causing dozens of government acceptances to cancel at the last minute. To avoid the embarrassment of empty seats, Bill Caelli, a 1971 ANU PhD graduate who worked for CDC, rang Robin Stanton and encouraged us all to come along. For members of DCS the timing was perfect – examination results for the semester had just been finalised! Well before the era of data projectors, the CDC presentations used a computer-controlled bank of Kodak Carousel 35mm slide projectors. Delicious Devonshire teas were served at morning and afternoon tea, and we retreated to a sumptuous, multi-course banquet lunch at University House, featuring Beef Wellington, cognac and cigars.

Business lunches became far less common when Australia introduced a Fringe Benefits Tax in 1986. Even then there were ways around it. In the late 1980s I travelled with Don Hardman, then head of ANU's administrative computing, to visit Imagineering for the purpose of finding out about network technologies which might be applicable in his area. Imagineering specialised in selling software for Macintoshes and IBM-compatibles. They had grown rapidly and occupied a building covering a whole city block near Sydney airport. The floor we met on had a large number of glass meeting rooms around the perimeter. Nearby restaurants delivered fancy, but FBT-free, food and drink to each of the meeting rooms, and attendees worked as they ate and drank. The network expert Don and I met had an ostentatious sports car with the number plate `DR LAN` and, in addition to his day job, played jazz clarinet in Sydney venues until 2am. We also met company owner Jodee Rich, who later co-founded One.Tel, with James Packer and Lachlan Murdoch on the board.[17]

The IAS consisted of seven Reseach Schools: Physical Sciences (RSPhysS), John Curtin School of Medical Research (JCSMR), Social Sciences (RSSS), Pacific Studies (RSPacS), Biological Sciences (RSBS), Chemistry (RSC) and Earth Sciences (RSES). Directors of research schools had full control of their budgets and wielded considerable power. As you will see from the 1974 Campus Map, in the 1970s, the IAS was almost all on the South side of the campus, geographically separated from SGS. The exception was late-coming RSC which had been deliberately located next to SGS Chemistry.

The SGS consisted of the Faculties of Arts, Economics, Science, Law and Asian Studies. The number of undergraduate degrees was very limited compared to today. You could enrol in BA, BEc, BSc, LLB, BA (Asian Studies), and BSc (Forestry) and various double-degree combinations. The Deputy Chair of the Board of the School of General Studies was Cecil Gibb (Psychology). The Business Manager was Val Brett.

Up until Gough Whitlam's ascendance, university students paid tuition fees, although many held Commonwealth Scholarships which covered them and provided a living allowance depending upon where their parents lived and how much they earned. I lived in Garran Hall, where an annual fee of a bit over $600 (yes, six hundred!) covered a room, three meals a day, and maid service (making the bed, vacuuming the room, and cleaning up the sink) five days a week!

Students paid a compulsory general services fee which funded the operation of the Students Association (SA), Student Representative Council (SRC), the Student Union, and the Sports Union. Because of his memories of compulsory fees being funnelled to leftist organisations such as the National Union of Students, Christopher Pyne made it a priority to abolish them after he became Minister for

[16] Perhaps I wasn't senior enough, or looked too honest.

[17] On p.4 of *Screw Business as Usual*, Richard Branson reports Desmond Tutu as saying, "People keep accusing me of name-dropping. Only last week I was at Buckingham Palace and the Queen said to me, 'Arch, you're name-dropping again.'"

Education in 2013. The SA sponsored a large range of clubs and societies, and gave grants to cover [flagon] wine and cheese at AGMs and other expenses. The Student Union operated a bar and cafeteria and organised concerts. I was elected to the SRC just before the SA abolished it, and before I'd had a chance to attend a meeting.

In 1971 both the SA and the Union were housed in an architecturally distinctive building, now demolished, directly across the lawn from the front entrance of the Chifley Library. The library lawn was a major focus of student activities, both organised and spontaneous. Students liked hanging out on the grass. "Concessions" such as bank branches, post office, and the University Co-operative Bookshop were located in huts close to the Law Faculty. There were a couple of beautiful flood-lit specimen *E. mannifera ssp. maculosa* trees near them. In 1973, all the services were moved to Union Court, along with the University Co-operative Credit Society, and the huts were demolished.

In 1971, there was a large collection of poor quality buildings on the North Eastern corner of the campus, now demolished. They were known as the Childers St Buildings and Childers St Hall. Childers St Hall was the scene of many a concert, play, and satirical review. Examinations were held in Childers St and another building was occupied by the department of Prehistory and Anthropology. In 1974 they moved to the top floor of the Crisp Building which was an extension to the northern wing of the Copland Building. My first office in 1975 was close to those of the distinguished prehistorian John Mulvaney and distinguished anthropologist Anthony Forge.

Later, in the 1980s, prior to construction of the ANUTech building on the corner of Barry Drive and North Road, ANUTech and its PCTech subsidiary which handled sales under the ANU-Apple Consortium Agreement were located in one of the Childers St buildings.

In 1973 a new Union building was created in a much less interesting architecural style and the old union housed ANU administrative staff. The new Union housed offices, the bar, and various eating venues including one whose furniture seemed to be styled on the Korova Milkbar from *A Clockwork Orange*. The new union was internally painted in rich colours, but the toilets still stank. The paved amphitheatre of Union Court reflected and amplified heat, and was never as much a focus of student activities as the library lawn had been. Across the bridge over Sullivan's Creek was the Sports Union, which was dominated by many (perhaps six) heavily booked squash courts.

In about 1975, a considerable sum of non-ANU money was raised to build the ANU Arts Centre between the Chifley Library and Sullivans Creek. Many editions of the *Old Time Music Hall* were held there. An in-house restaurant called *Vivaldi's* was decorated with theatrical playbills and photos, and *The Gods* for many years served the best coffee on campus.

In 2019, the old Union, the new Union, Union Court and the Arts Centre were all demolished to make way for Kambri.

In the 1970s, airconditioners were used only in computer rooms and laboratories. Requests to spend money on airconditioning for people and teaching rooms were automatically refused. That meant that working conditions were sometimes rather hot. It was just as well that undergraduate teaching was not undertaken in summer. Richard Brent recalls, "I managed to get an airconditioner installed in Bev Johnstone's office on the grounds that there was a computer there. Since it was adjacent to my office, I could leave the communicating door open to get some cool air."

2.2.1 1970s: ANU Teaching Methods

In the early 1970s the Faculty of Economics had a practice, if not an official policy, of not scheduling lectures or tutorials on either Mondays or Fridays, presumably to support the economy of the NSW South Coast.

In 1971, nearly all ANU teaching rooms were fitted with blackboards and boxes of chalk. Some lecturers used nothing but "talk and chalk." The photo on Page 24 shows that blackboard technology is not dead. Over the decade, whiteboards made an appearance but there were always difficulties finding working pens.

In 1971, most lecturers made use of overhead projectors, either writing and drawing onto a continuous transparency roll, or bringing pre-prepared hand-drawn transparencies. Sensible lecturers carried their own set of pens. Eventually, photocopier technology improved to allow material, including computer output, to be copied on to transparencies. Prior to this a number of expensive photocopier drum replacements were required due to melted foils.

Use of overhead projectors continued at ANU until the 1990s. I attended the Alan Turing Centenary Conference in Manchester in 2012 where a keynote speaker, Roger Penrose, showed that this technology is still not quite dead. He used an overhead projector (dragged out of storage in the Manchester Town Hall and deep-cleaned) to brilliant effect, even to the extent of showing an animation involving strips of transparency. His performance was the highlight of the conference, outshining the half dozen Turing Award winners who also spoke.[18]

In some ANU subjects 35mm slides were projected but I don't remember this happening in computer science. Some subjects also screened 16mm films – I worked as a casual projectionist for the Visual Aids Unit (based in I Block, I think it was) and showed films for Asian Studies and Anthropology among others. I used to go to Visual Aids half an hour before the lecture, borrow a Bell and Howell projector, set up in the gap between lectures, screen the reel, hoping there weren't any twists and that there was no sprocket damage, rewind after the show, and return the equipment to I Block.

These days it is the norm for lecture notes, comprehensive course materials, and assignments, to be created and distributed electronically. In the early 1970s, comprehensive course materials were not generally provided. Students were expected to attend lectures and take their own notes. Printed assignment sheets, and such other notes as there were, were duplicated using a Gestetner duplicator and handed out in class.

In 2020 ANU mathematicians still prefer blackboards. Bob Edwards demonstrates the blackboards in the new Hanna Neumann Building. I'm told that the chalk is a special variety, imported from Japan. ***Photo: David Hawking***

In the first edition of Computer Science C01, a part-time student from the Public Service took away a 40-page document belonging to the lecturer and came back with 80 copies, one for each C01 student, made by his department's staff. This was a very valuable public service!

Lecturers were expected to put key text books and other course resources on "Short Loan" at the Reserve Desk of the Chifley Library. You could reserve an item for a particular time and I think you could borrow it for two hours, or overnight if you collected it just before the library shut at 11pm. Sometimes there was intense contention for items on Short Loan.

For critical items not on Short Loan, some students would reserve a book for their own dedicated use, by shelving it in the wrong section. Library staff were always on the lookout for this, and also for

[18] Roger Penrose was a winner of the 2020 Nobel Prize for Physics.

CCAE students sneaking into the ANU library. A couple of the library staff were suspicious-minded and authoritarian. Toilet rolls in the library were labeled: *Arts Degrees. Tear one off and put to good use.*

To borrow anything from the library you needed a Library Card. They were embossed with your details. To borrow an item, you filled out a flimsy paper loan slip including a carbon layer, and put the slip in a small roller machine with your card, then handed the item and the loan slip at the Loans Desk, where the slip was retained and the required return date was stamped onto a form glued into the back of the book. There was an out-of-hours return chute, and one of our friends, now a professor at ANU, once drove us in his Land Cruiser up the steps of the library to drop off a book without having to get out of the vehicle.

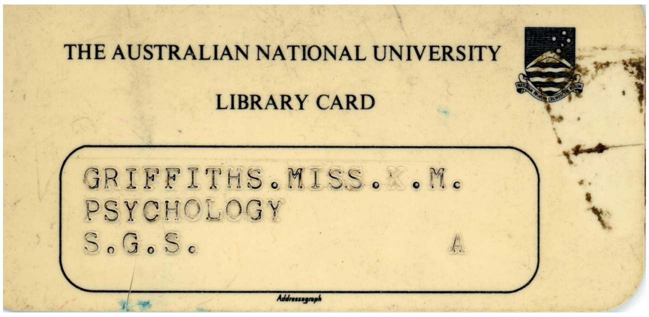

THE AUSTRALIAN NATIONAL UNIVERSITY

LIBRARY CARD

GRIFFITHS.MISS.K.M.
PSYCHOLOGY
S.G.S. A

Addressograph

The Chifley Library used the Library of Congress classification system rather than Dewey Decimal, and there was an author/title card index in wooden drawer units at the entry to the library buildings. There was also a subject card index but this had very little coverage.

The library card catalogues were the only "search engine" to which I had access as an undergraduate. However, I believe that researchers in JCSMR were able to submit queries by post to the MEDLARS (Medical Literature Analysis and Retrieval System) at the US National Library of Medicine. Ten days or so later they would receive a printout of matching articles.

ANU students who actually read books and journal articles, would likely also frequent the National Library, the RSPhysS Library in the Le Couteur building (now demolished), with its magnificent views over the lake, and perhaps even the library at the Royal Military College, Duntroon (which became ADFA in 1986). Having become a more serious student half-way through third year, the RSPhysS library was my favourite place to work, followed by perching in the bow windows of the Chifley.

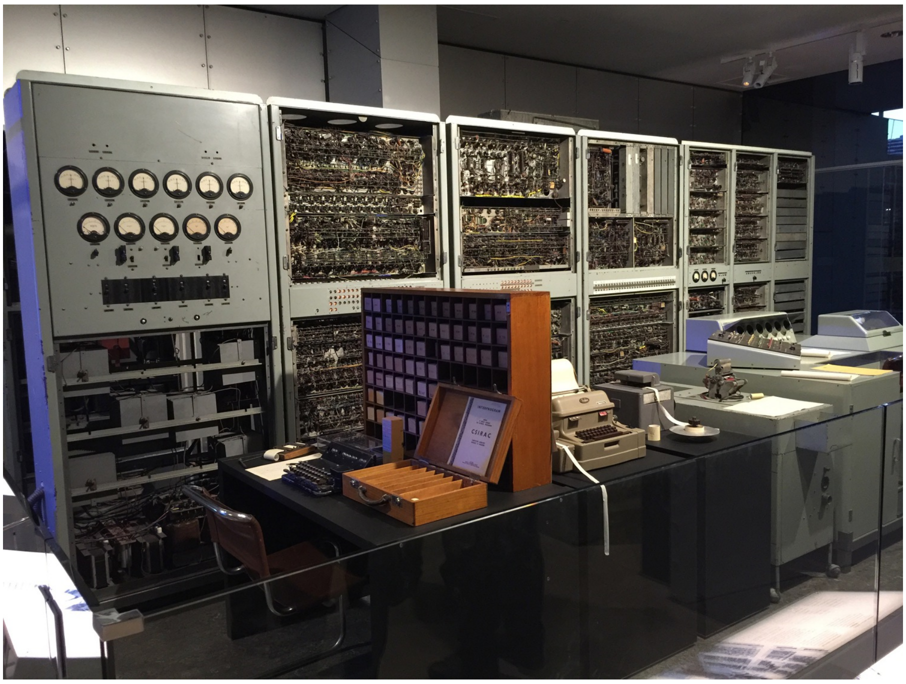

CSIRAC in Melbourne Museum. *Photo: David Hawking, June 2016.*

2.3 Australian Computing Before 1971

ANU was actually rather late getting into computing.

According to the Australian Computer Society heritage project[19] computing in Australia began in November 1949 with the powering up of the CSIR Mark 1 designed by Trevor Pearcey and built by him, Maston Beard and Geoff Hill. Later called CSIRAC, it was only the fourth stored program electronic computer ever built and is the oldest such machine still in existence.

In August 1951 the first computer conference in Australia was organised by CSIR and the Department of Electrical Engineering at the University of Sydney.[20] Key players were Trevor Pearcey, David Myers, and Ross Blunden. The second Australian computer conference, *Data Processing and Automatic Computing Machines*, did not occur until 1957, organised by John Ovenstone at the Weapons Research Establishment in Salisbury, South Australia.[21]

In the mid 1950s Harry Messel in the Department of Physics at Sydney University, ordered the construction of a copy of ILLIAC, which came into service in March 1956, and was known as SILLIAC. It was managed by John Bennett, who delivered Australia's first university course in computer programming in June of 1956.

According to a history by Robin Vowels[22], there is room for debate about whether the English Electric DEUCE computer installed at the NSW University of Technology, now UNSW, pre-dated or post-dated SILLIAC. Of course the DEUCE was built in England. It was closely derived from Alan Turing's ACE design.

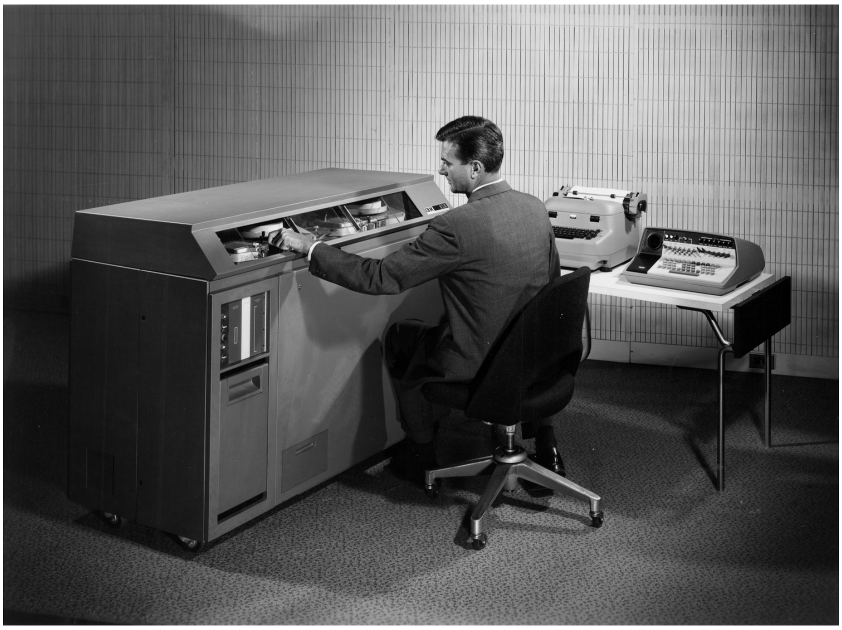

An IBM 610, of the same type as ANU's first computer. Programs were represented on paper tape rather than in memory (i.e. it was not a stored-program computer.) The door at the lower left of the cabinet conceals the fold-out plugboard. ***Photo courtesy of IBM, ©International Business Machines Corporation.***

[19] `https://ia.acs.org.au/article/2016/the-acs-heritage-project--chapter-1.html`

[20] `https://collections.museumsvictoria.com.au/articles/2718`

[21] See `https://50years.acs.org.au/heritage-projects/acs-heritage-project--chapter-7.html` for more information about Ovenstone and computing at WRE.

[22] *UTECOM – An English Electric DEUCE.* `http://members.dodo.com.au/~robin51/deuce.htm`

ANU's first computer

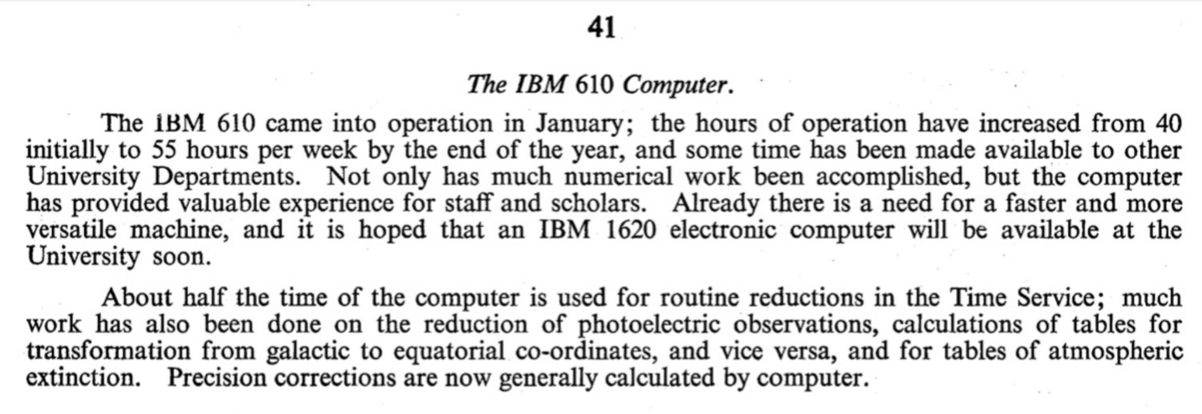

41

The IBM 610 *Computer.*

The IBM 610 came into operation in January; the hours of operation have increased from 40 initially to 55 hours per week by the end of the year, and some time has been made available to other University Departments. Not only has much numerical work been accomplished, but the computer has provided valuable experience for staff and scholars. Already there is a need for a faster and more versatile machine, and it is hoped that an IBM 1620 electronic computer will be available at the University soon.

About half the time of the computer is used for routine reductions in the Time Service; much work has also been done on the reduction of photoelectric observations, calculations of tables for transformation from galactic to equatorial co-ordinates, and vice versa, and for tables of atmospheric extinction. Precision corrections are now generally calculated by computer.

From Mount Stromlo section of 1960 ANU Annual Report (`https://openresearch-repository.anu.edu.au/bitstream/1885/14770/1/AnnualReport-1960.pdf`)

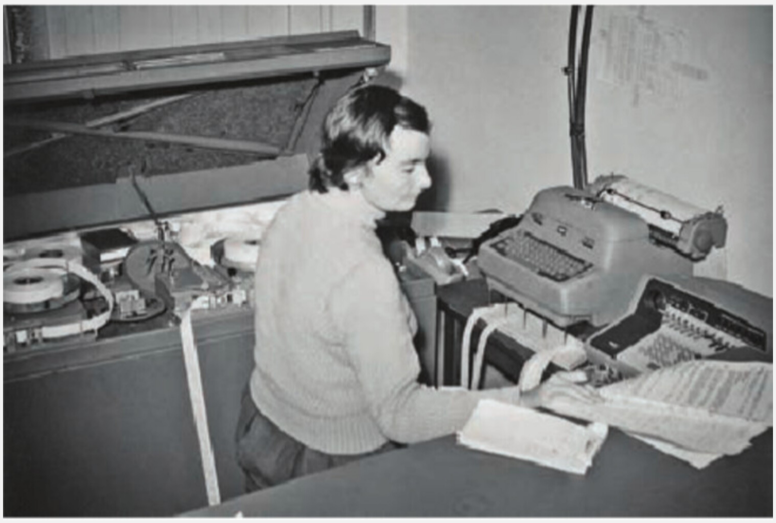

Claire Wehner enters data into the Mt Stromlo IBM 610. *Photo: Mt Stromlo archives*

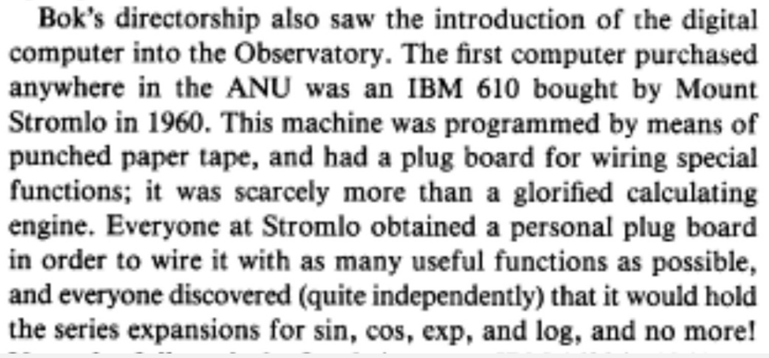

Bok's directorship also saw the introduction of the digital computer into the Observatory. The first computer purchased anywhere in the ANU was an IBM 610 bought by Mount Stromlo in 1960. This machine was programmed by means of punched paper tape, and had a plug board for wiring special functions; it was scarcely more than a glorified calculating engine. Everyone at Stromlo obtained a personal plug board in order to wire it with as many useful functions as possible, and everyone discovered (quite independently) that it would hold the series expansions for sin, cos, exp, and log, and no more!

From A.R. Hyland and D.J. Faulkner, *From the Sun to the Universe – The Woolley and Bok Directorships at Mount Stromlo*, Proc ASA 8(2), 1989, p. 224. (`http://articles.adsabs.harvard.edu/cgi-bin/nph-iarticle_query?bibcode=1989PASAu...8..216H&db_key=AST&page_ind=8&data_type=GIF&type=SCREEN_VIEW&classic=YES`)

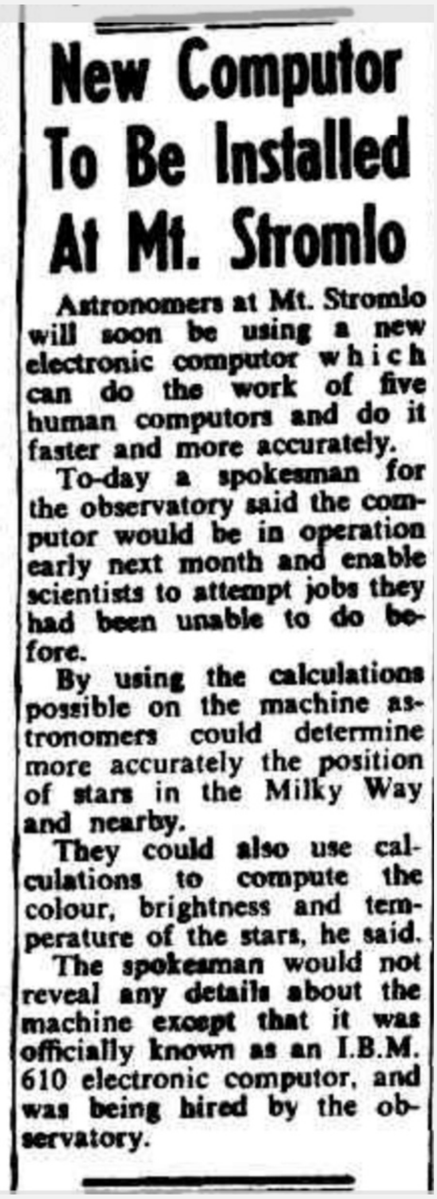

New Computor To Be Installed At Mt. Stromlo

Astronomers at Mt. Stromlo will soon be using a new electronic computor which can do the work of five human computors and do it faster and more accurately.

To-day a spokesman for the observatory said the computor would be in operation early next month and enable scientists to attempt jobs they had been unable to do before.

By using the calculations possible on the machine astronomers could determine more accurately the position of stars in the Milky Way and nearby.

They could also use calculations to compute the colour, brightness and temperature of the stars, he said.

The spokesman would not reveal any details about the machine except that it was officially known as an I.B.M. 610 electronic computor, and was being hired by the observatory.

R: Canberra Times, 15 Jan 1960. (`https://trove.nla.gov.au/newspaper/article/103075045`)

ANU's 1959 annual report[23] foreshadows the arrival of an IBM 610 on 01 Feb, 1960 for use by the National Time Service at Mt Stromlo. (The latter provided time signals for transmission by Belconnen Naval Wireless Station.) The report says that Mr Boots [ANU] and Mr Grenot [IBM] were preparing the programmes for it which would calculate the apparent declinations and times of transit of the PZT stars, and reduce the observations for clock corrections and latitude variation.

Weighing as little as 360kg, and described as a mobile deskside machine which could be operated by a single person, some regard the 610 as IBM's first personal computer – indeed it was originally called the Personal Automatic Computer or PAC. It was one of the first computers to accept input from a keyboard, and it sent printed output to a typewriter. Alternatively, input/ouput via paper tape was supported by the reader/punches at the top of the cabinet. Internally it used vacuum tubes and a magnetic drum, and had a memory of 2604 decimal digits, configured as 84 registers of 31 digits each. It was called "Auto-point" because it allowed automatic handling of a decimal point within the registers.

Mount Stromlo Observatory acquired the IBM 610[24] computer in January, 1960. The Canberra Times report at the time suggests that it was rented rather than purchased. The academic price was USD460 per month in 1957, approximately USD4185 p.m in 2019 dollars, or A$65k p.a. plus extra cost in Australian market. According to an article by Hyland and Faulkner (see extract on Page 27) it was the first computer anywhere in ANU. It could be configured using a plugboard seemingly much smaller and simpler than the one for the Univac 1004, which is pictured on Page 60. An interesting video of the 610 is at `https://youtu.be/5pkKbyTLIdI`.

Kenneth Le Couteur was ANU's foundation Professor of Theoretical Physics, appointed by Mark Oliphant in 1956. During World War II, Le Couteur had worked as a code breaker at Bletchley Park, where he used the Colossus computer and its predecessor, known as the Heath Robinson machine. In 1962, he arranged for the Department of Theoretical Physics to acquire an IBM 1620 computer.

The IBM 1620 Model I was the first stored program computer to be located at ANU. It seems from the ANU Annual Report of 1960 that it was rented.[25]

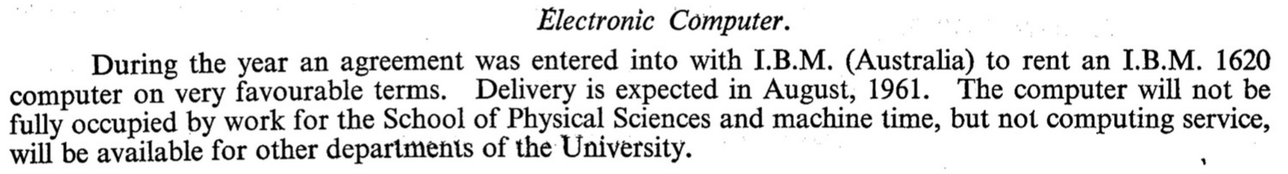

Electronic Computer.

During the year an agreement was entered into with I.B.M. (Australia) to rent an I.B.M. 1620 computer on very favourable terms. Delivery is expected in August, 1961. The computer will not be fully occupied by work for the School of Physical Sciences and machine time, but not computing service, will be available for other departments of the University.

1960: From the ANU Annual Report, p.52.[26]

Mark Oliphant's section of the ANU Annual Report for 1961 provides context for the acquisition of the 1620.

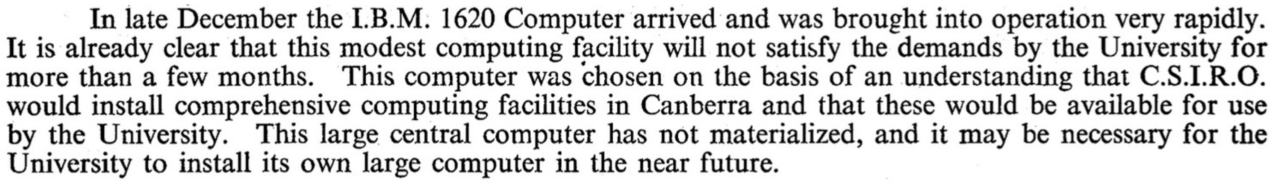

In late December the I.B.M. 1620 Computer arrived and was brought into operation very rapidly. It is already clear that this modest computing facility will not satisfy the demands by the University for more than a few months. This computer was chosen on the basis of an understanding that C.S.I.R.O. would install comprehensive computing facilities in Canberra and that these would be available for use by the University. This large central computer has not materialized, and it may be necessary for the University to install its own large computer in the near future.

1961: From the ANU Annual Report, p.43. [27] Elsewhere, it is reported that the 1620 arrived on 02 January, 1962.

[23] `https://openresearch-repository.anu.edu.au/handle/1885/14769`

[24] `https://www.ibm.com/ibm/history/exhibits/vintage/vintage_4506VV4001.html`
`http://www.columbia.edu/cu/computinghistory/610.html`

[25] I have read that in this era, IBM only rented out its computers. Outright purchase only became possible later.

[26] `https://openresearch-repository.anu.edu.au/bitstream/1885/14770/1/AnnualReport-1960.pdf`

[27] `https://openresearch-repository.anu.edu.au/bitstream/1885/14771/1/AnnualReport-1961.pdf`

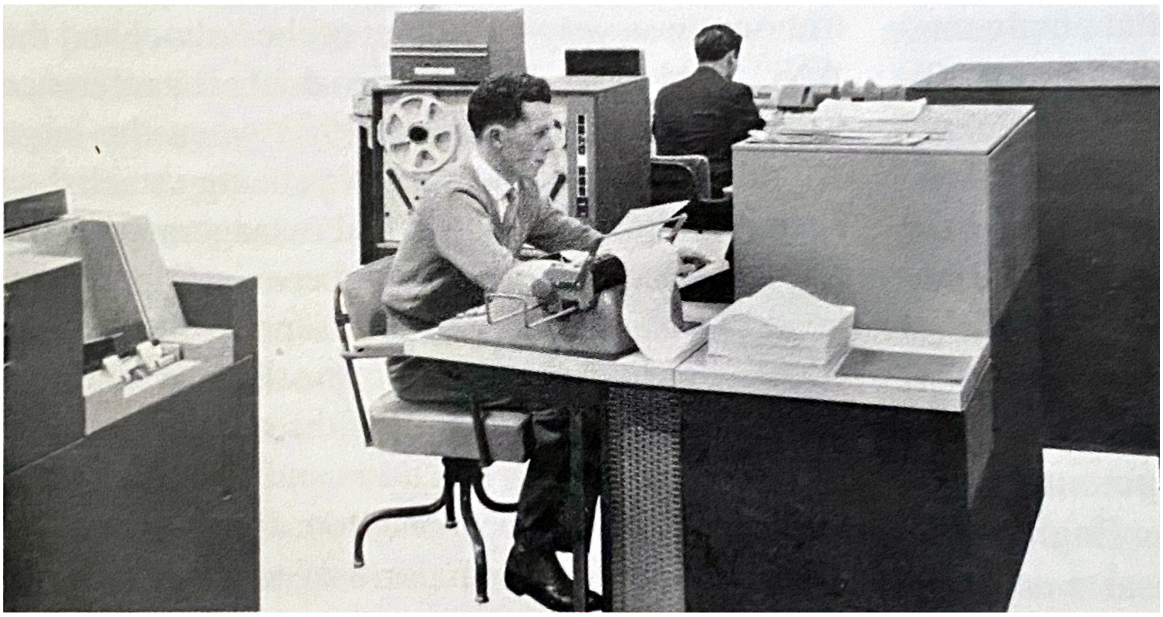

1962: Brian Robson at the console of the IBM 1620. *Photo: Brian Robson, via p.39 Ophel and Jenkin. Fire In The Belly:The history of the ANU physics research school from 1946 to 1996.*[28]

12 Feb 2021: Brian Robson, manager of ANU's IBM 1620s from January, 1962, at his desk in the Oliphant Building. *Photo: David Hawking*

[28] `https://physics.anu.edu.au/about/history/fire_in_the_belly/`

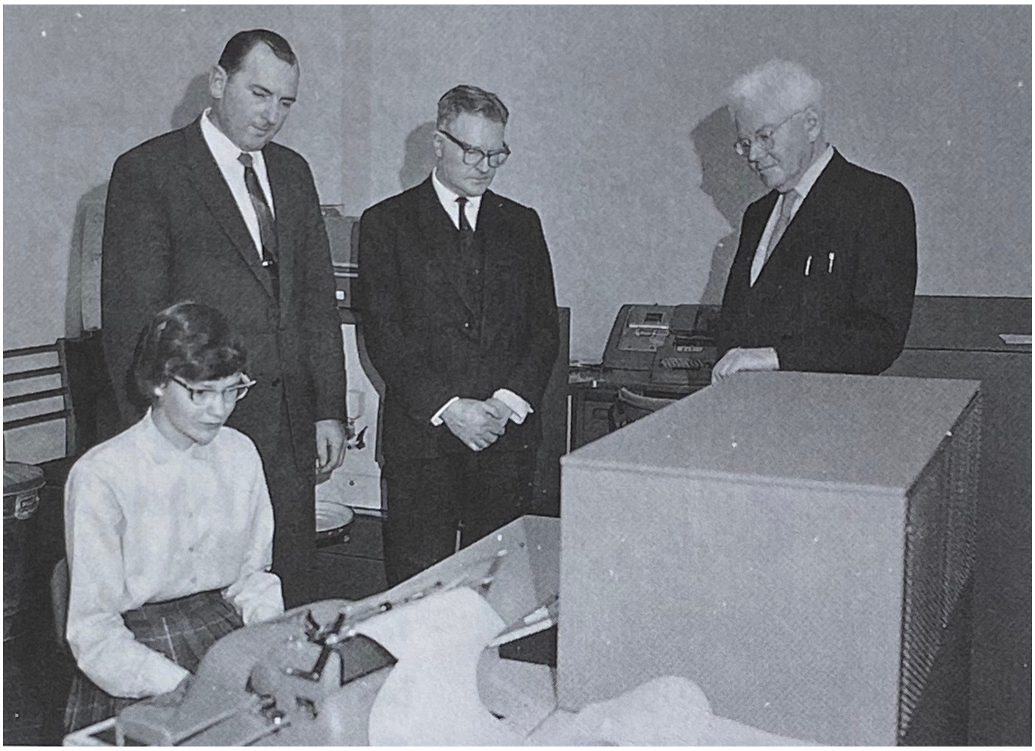

1962: Presumably at the launch of the IBM 1620. L to R: Margaret Campbell (Operator), unknown IBM employee, Leonard Huxley (Vice-Chancellor), Mark Oliphant (Director, RSPhysS). (This image appeared in Foster & Varghese, *The Making of the Australian National University* with a caption containing many apparent errors.) *Photo: Australian News and Information Bureau.*

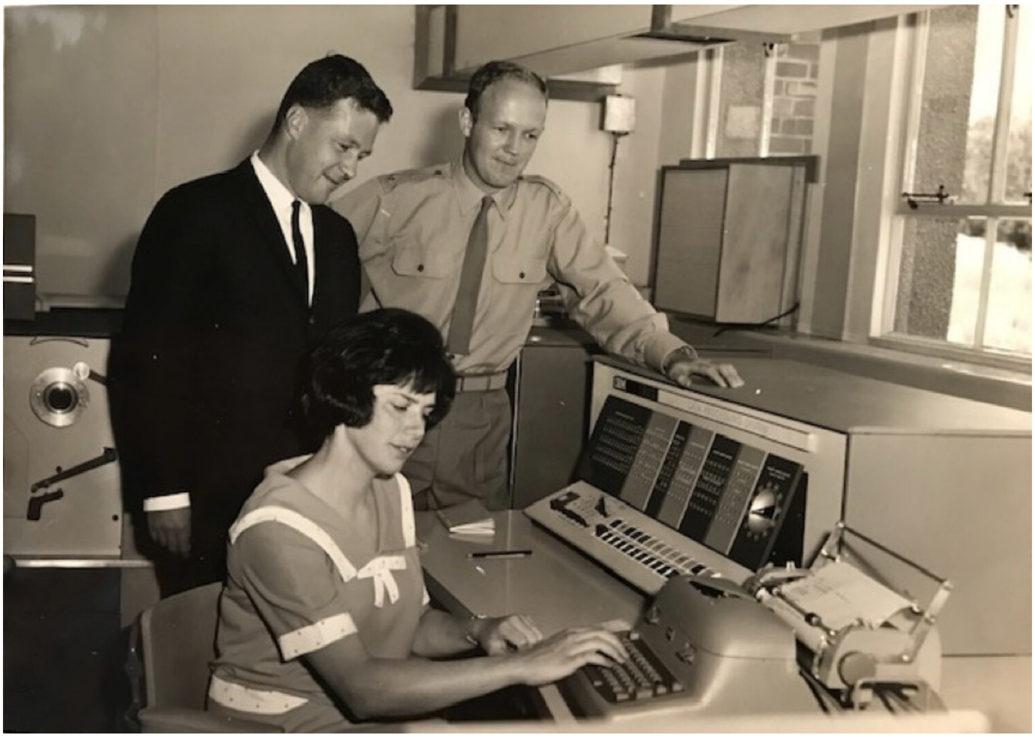

1964/5: Lorrel Sherar (later known as Ann Apthorp) at the console of ANU's IBM 1620, watched by Mike Goldrick (believed to be head of IBM Canberra), and an unidentified man. *Photo supplied by Ann Apthorp.*

The 1620 was managed by Brian Robson. Imagine my delight on discovering, in February 2021, that 59 years later, Brian Robson is still an active researcher with an office in the Oliphant Building, despite having compulsorily retired in 1999! He's currently completing the writing of a book, *Understanding Gravity: The Generation Model Approach.*

Brian began his postgraduate studies in Theoretical Physics at the University of Melbourne in 1955, the year in which CSIR Mark I, Australia's only operational computer, was shut down in Sydney and trucked to the University of Melbourne.[29] The relocated machine was launched as CSIRAC on 14 June, 1956 and Brian used it from soon after that until he submitted his PhD thesis on 01 April, 1960. He always ran computations on CSIRAC three times, and used the majority answer! Despite this distrust of CSIRAC's reliability, he recently recomputed scientific tables that he had prepared on CSIRAC and found them to be quite accurate.

He started work in the Department of Theoretical Physics at ANU in April 1960, and was put in charge of the IBM 1620 when it arrived on 02 January, 1962.

The IBM 1620 was located on the ground floor (northern end, facing where the lake is now) of the Oliphant Building (then known as the Chifley Building). Calculations for the National Time Service were transferred from the 610 to the 1620.

According to `https://genderinstitute.anu.edu.au/elizabeth-reid-transforming-adversity-impact` the first operator/programmer for the 1620 was Elizabeth Reid, who went on to a highly distinguished career after serving as Adviser on Women's Affairs to Prime Minister Gough Whitlam in 1972.

> In 1961 she taught herself Fortran and became the first operator/programmer for the spanking new IBM 1620. By 1965 she had received a Bachelor of Arts (Hons) First Class from ANU and was on her way to England on a Commonwealth Travelling Scholarship to study philosophy at Oxford.

Elizabeth Reid (first Operator/Programmer of ANU's 1620) photographed in 1973. ***Photo: ANU Archives. ANU Photographic Services.***

Initially there were two operators. Margaret Campbell and Peter Tindale, who later worked as a programmer in the Computer Centre. Several additional programmers were later recruited to assist

[29] Alistair Moffat, *Fifty Years of Computing at the University of Melbourne*.`https://people.eng.unimelb.edu.au/ammoffat/fifty-years/mof06history.pdf`

with the transition to the IBM 360/50. By the time it arrived in 1966, there were two operators and four programmers. Brian tells me,

> ... of the six operators/programmers of the IBM 1620, who transferred to the IBM 360 model 50 in 1965-66, only Peter Tindale was male.

One of the females was Barbara Davidson whose story appears on Page 273.

Brian Robson met a number of very interesting people in the course of his career. One of them was the grand-daughter of Marie and Pierre Curie. Another was Basil de Ferranti (a British MP associated with Ferranti Computers), who predicted the advent of personal computers decades before they appeared. Finally, C.S. Daley[30], who was vitally important to the building of Canberra, and after whom Daley Road is named, came to an Open Day, and played some of the primitive games available on the IBM 1620. One was a reaction time game in which the winner pressed `STOP` in the shortest time after `START` – due to much practice, Brian usually won.

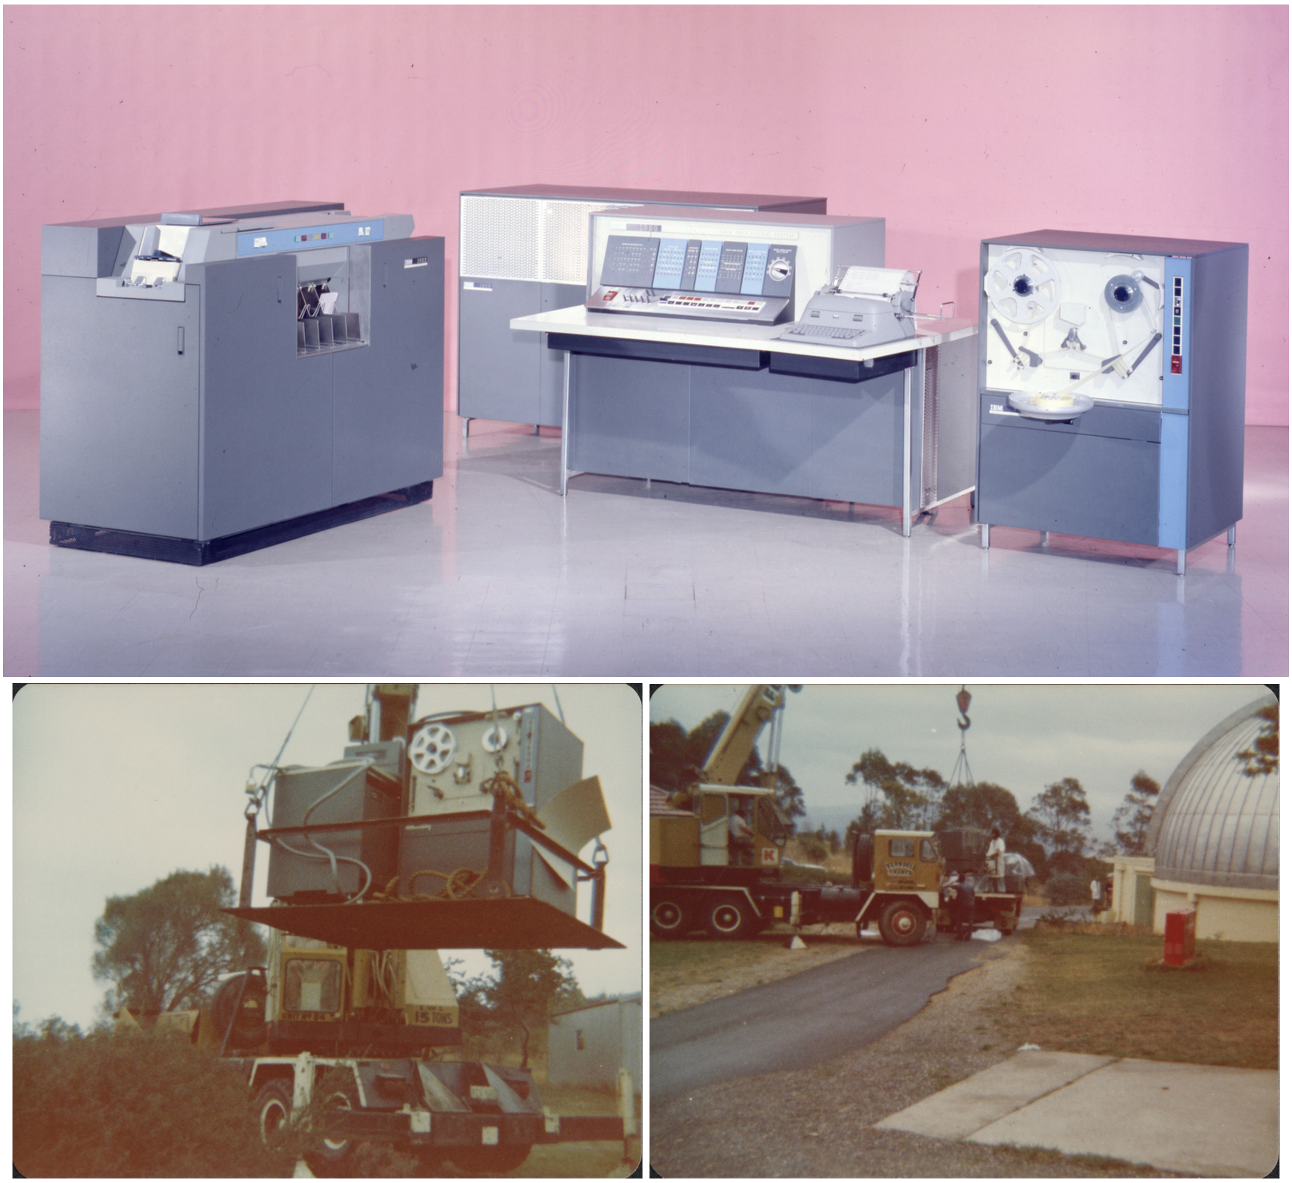

Top: A clearer picture of an IBM 1620 installation. Unfortunately, not ANU's. The magnetic core memory unit at the rear held 20K or 40K decimal digits, and took six people to carry. *Photo courtesy of IBM, ©International Business Machines Corporation.* **Bottom: Mid-1970s: Removing ANU's first 1620 from a basement at Mount Stromlo Observatory for transfer to DCS in the Copland Building.** *Photo ANU Archives: Records of the Commonwealth Solar Observatory (Mt Stromlo) and Kurt Gottlieb.*

[30] `https://adb.anu.edu.au/biography/daley-charles-studdy-5865`

The story of ANU's two IBM 1620s

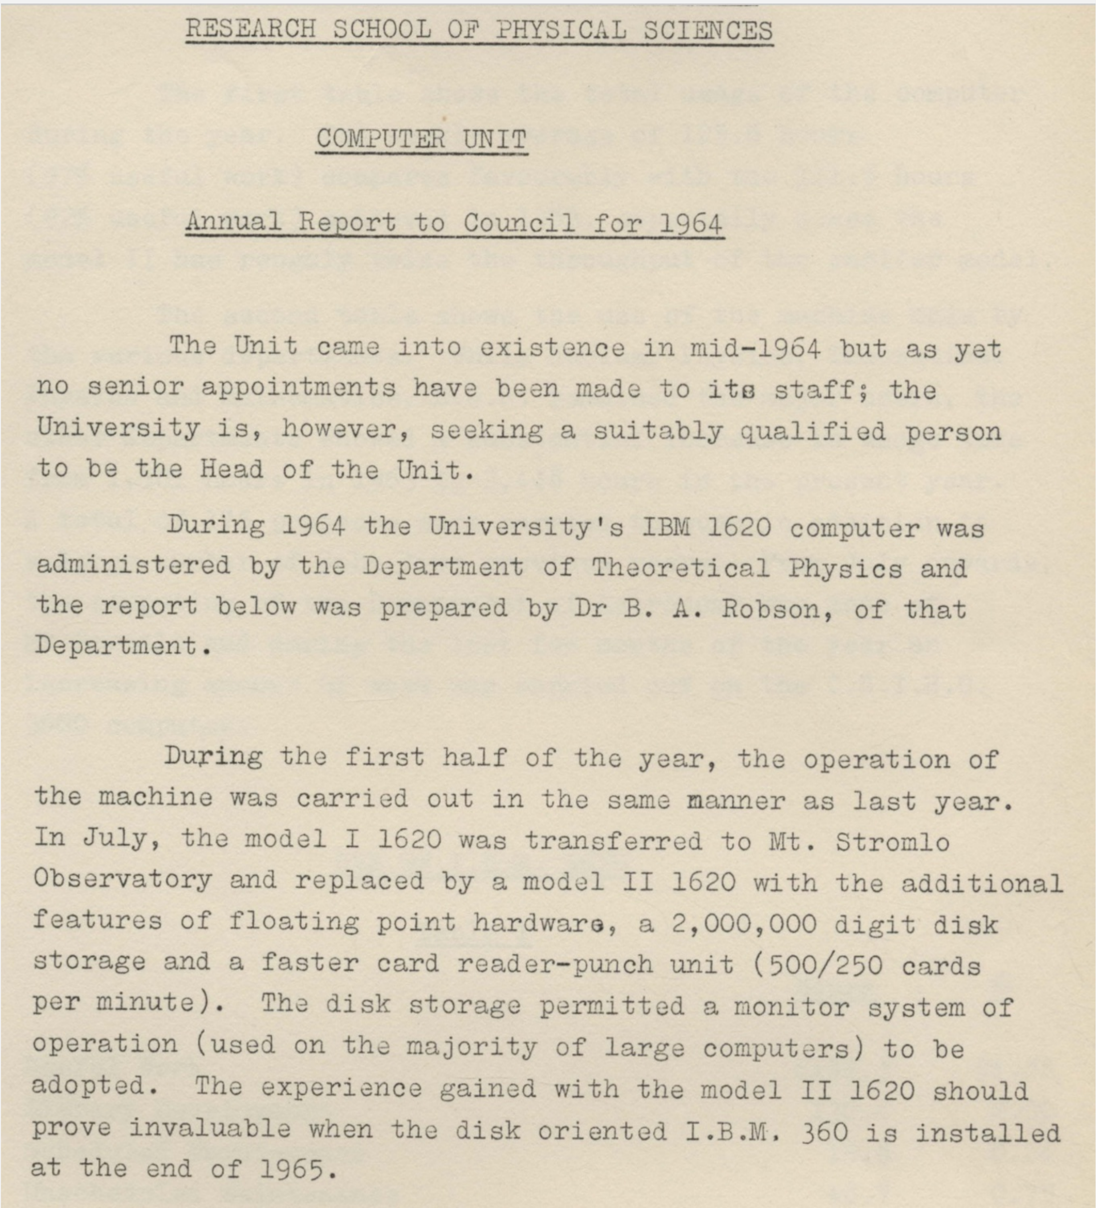

RESEARCH SCHOOL OF PHYSICAL SCIENCES

COMPUTER UNIT

Annual Report to Council for 1964

The Unit came into existence in mid-1964 but as yet no senior appointments have been made to its staff; the University is, however, seeking a suitably qualified person to be the Head of the Unit.

During 1964 the University's IBM 1620 computer was administered by the Department of Theoretical Physics and the report below was prepared by Dr B. A. Robson, of that Department.

During the first half of the year, the operation of the machine was carried out in the same manner as last year. In July, the model I 1620 was transferred to Mt. Stromlo Observatory and replaced by a model II 1620 with the additional features of floating point hardware, a 2,000,000 digit disk storage and a faster card reader-punch unit (500/250 cards per minute). The disk storage permitted a monitor system of operation (used on the majority of large computers) to be adopted. The experience gained with the model II 1620 should prove invaluable when the disk oriented I.B.M. 360 is installed at the end of 1965.

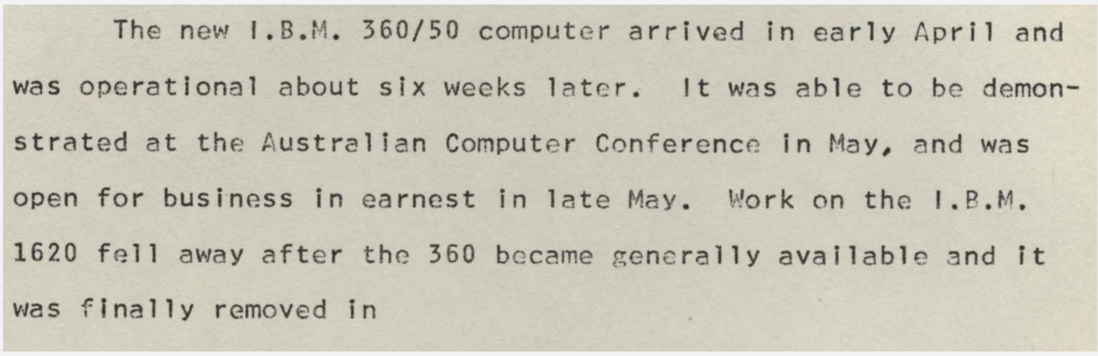

The new I.B.M. 360/50 computer arrived in early April and was operational about six weeks later. It was able to be demonstrated at the Australian Computer Conference in May, and was open for business in earnest in late May. Work on the I.B.M. 1620 fell away after the 360 became generally available and it was finally removed in

Top: Extract from Computer Unit Annual Report for 1964.[a] **Bottom: Extract from Computer Centre Annual Report for 1966. The last sentence is incomplete in the original.** [b]

[a] https://openresearch-repository.anu.edu.au/bitstream/1885/149339/2/2.1.8.3%283%29-17-1964.pdf

[b] https://openresearch-repository.anu.edu.au/bitstream/1885/149339/4/2.1.8.4(3)-13-1966.pdf

The Department of Statistics was located in the Mathematical Sciences Building (later named the Le Couteur Building, now demolished and in the process of being rebuilt). Brian Robson claims to have taught, in 1964, the first ANU undergraduate computing lectures, to students enrolled in Ted Hannan's Statistics III course.

The RSPhyS Computer Unit Annual Report for 1964[31] records that the university acquired an IBM 1620 model II in 1964 to replace the model I and transferred the older machine to Mt Stromlo, replacing the IBM 610.

Later, Brian was a member of the Procurement Committee which oversaw the purchase of an IBM 360/50. Other members were Trevor Swan (Chair, Professor of Economics), and Dale Hebbard (Nuclear Physics). The committee considered a Control Data (CDC) 3600, but it wasn't affordable for ANU. (CSIRO acquired one and ANU people sometimes used it.)

The 1966 Annual Report for the Computer Centre, reports that the model II IBM 1620 was removed sometime during 1966 after the arrival of the IBM 360/50. See the panel on Page 33.

An IBM 360 installation. Unfortunately, I was unable to locate a photo of ANU's 360/50. ***Photo courtesy of IBM, ©International Business Machines Corporation.***

Michael Cook joined the ANU Psychology Department in 1961. After attending a FORTRAN course at CSIRO, he started programming statistical analyses for his Psychology colleagues on CSIRO's CDC 3600. Michael later transferred his focus to an Olivetti Programma 101[32], acquired by Psychology in the mid 1960s. It was a sophisticated programmable calculator, with programs stored on magnetic strips. In the 1970s Michael began using a PDP-8, first for data analysis and then for controlling experiments. After using a PDP-11 while on sabbatical in Edinburgh, he persuaded the department to buy one. In addition to data analysis it was used for document processing by Michael and Ken Vine and ran John Trotter's Discourse Supervisor (DS),[33] coded in LISP.

[31] https://openresearch-repository.anu.edu.au/bitstream/1885/149339/2/2.1.8.3%283%29-17-1964.pdf

[32] https://en.wikipedia.org/wiki/Programma_101

[33] The idea was that an academic discussion would proceed by the participants typing statements into the computer, where the DS would reject any statement which was not an agreed premise or did not follow from the premises by a valid chain of logic.

ANU's first CS PhD completion?

John O'Callaghan completed in 1968, in Engineering Physics with Steve Kaneff as supervisor. His thesis was titled, *Pattern recognition using some principles of the organism-environment interaction.*[a] See Appendix C for some early CS PhD completions and discussion of other contenders.

[a] https://openresearch-repository.anu.edu.au/handle/1885/139353

Significant consequences of prime number computation, as told by Richard Brent

It wasn't Mersenne primes that burned out a magnetic core on the IBM 360/50, but it was related to primes. I was looking at the distribution of gaps between adjacent primes, in particular how large a gap can you get for primes $< x$ say? The answer seems to be about $(\log x)^2$ but no-one has any idea how to prove this. The average gap is about $\log x$ (which can be proved).

So my program was "sieving" for primes – the program was actually written for the IBM 360/91 while I was at Yorktown Heights, but it ran fine (if slower) on the 360/50. Because of the way sieving works, it is likely to access certain words of memory in a periodic fashion, with many different periods possible. It turned out that one of those periods was close to a resonant frequency in the wires that were threaded through the computer memory. (The memory was magnetic core, with an x/y grid of wires so a pair of wires passed through each magnetic core, which could store one bit.)

This resonance set the wires vibrating and eventually wore the insulation off them, causing a 1-bit memory fault. I noticed this and just remapped my program to avoid the offending bit, but the payroll program was not so lucky! IBM claimed this was impossible, but it later turned out that an IBM engineer had predicted that it might happen. I used to run my program overnight, bribing the operators to start it up when they finished running admin programs for the night. The program would write checkpoints to a file, so could be interrupted and restarted any time.

Later on, I ran various programs to do with prime numbers on the Univac 1108. These played a role in detecting the infamous Intel Pentium division bug. It was discovered by Thomas Nicely, who initially ran his program on Intel 80486 machines (32-bit) and verified his results by comparison with my results on the Univac 1108 (36-bit). When he tried to run on a Pentium, he occasionally got errors (symptoms of a table used by the division algorithm not being completely initialised), so he was able to confirm that the results were wrong (and the 80486 results correct) by comparing with my results. It took some effort to convince Intel, and eventually cost them hundreds of millions of US dollars to recall/replace faulty Pentium chips.

2.4 1970s: ANU Computing Facilities

In 1971, when DCS came into existence as a sub-department of the Department of Statistics, ANU's main computing resource was an IBM 360/50 located in the Cockroft Building. It met all of ANU administrative computing needs, and ran batch jobs for academics and students. It ran a *Multi-Programming with a Fixed Number of Tasks* (MPFT) operating system which meant it could run data processing jobs in one slice of memory, and student FORTRAN jobs in another. . . .All batch jobs of course. Richard Brent ran his prime number explorer whenever the machine was otherwise idle.

A staff of four Operators ran the machine for most of the day, loading magnetic tapes and switching 7MB or even 14MB removable disk packs. In 1971, the Computer Supervisor was Alan Harris and the operators were Margaret Brinkley, Keith France, Cher Lek, and Errol Rumpf.[34] Overnight, researchers and graduate students could book time slots for exclusive use of the machine. I was friends with a PhD student and once or twice piggybacked on his slots in the middle of the night.

Just about all programs and datasets were represented as decks of punched cards. All programs and data for administration, and some for research, were written on printed coding sheets and punched by the Data Processing Unit, housed in the ancient buildings (F-Block?) to the north of the Menzies Library. DPU was headed by Ms Skaidrite Darius who was awarded an Honorary

[34] According to 1971 Computer Centre Annual Report, accessible from https://openresearch-repository.anu.edu.au/handle/1885/149187

Doctorate by ANU in 2019.

Among the few other computers around the campus were the previously mentioned IBM 1620 at Mt Stromlo, a PDP-15 in Engineering Physics, two PDP-8s in Psychology, one in Zoology, and a couple of Hewlett Packard machines in RSES. The Hewlett Packards were run and extensively tailored for RSES use by Peter Arriens, a geochronologist with an Antarctic glacier named after him.

PDP-8s were a remarkable machine: 12-bit word length, maximum integer value +2047, and a memory of $2^{12} = 4096$ words. Despite underwhelming resources they were often used as controllers for scientific equipment. Languages provided included LISP (yes, LISP!) and FOCAL, DEC's alternative to BASIC.

In the early 1970s IBM was the totally dominant computer company and it had an office in Northbourne Avenue known as the IBM Systems Institute. IBM was bigger than all of its competitors combined – Snow White versus the Seven Dwarfs. The main US competitors were known as the BUNCH (Burroughs, Univac, NCR, Control Data and Honeywell). However, the balance changed with the advent of minicomputers, an area in which IBM never achieved dominance. Digital Equipment Corporation (DEC), with its range of minicomputers including 12-bit PDP-8s, 16-bit PDP-11s, 18-bit PDP-15s, 36-bit DEC-10s, and 32-bit VAXes, strongly targeted the university market and achieved great success at ANU over a period of 15 or 20 years.

Although IBM released the IBM PC in August 1981, many in the IBM world saw it as a toy, and predicted that for every dollar IBM earned from a PC they would earn ten from the mainframe it needed to connect to, to do serious work.[35]

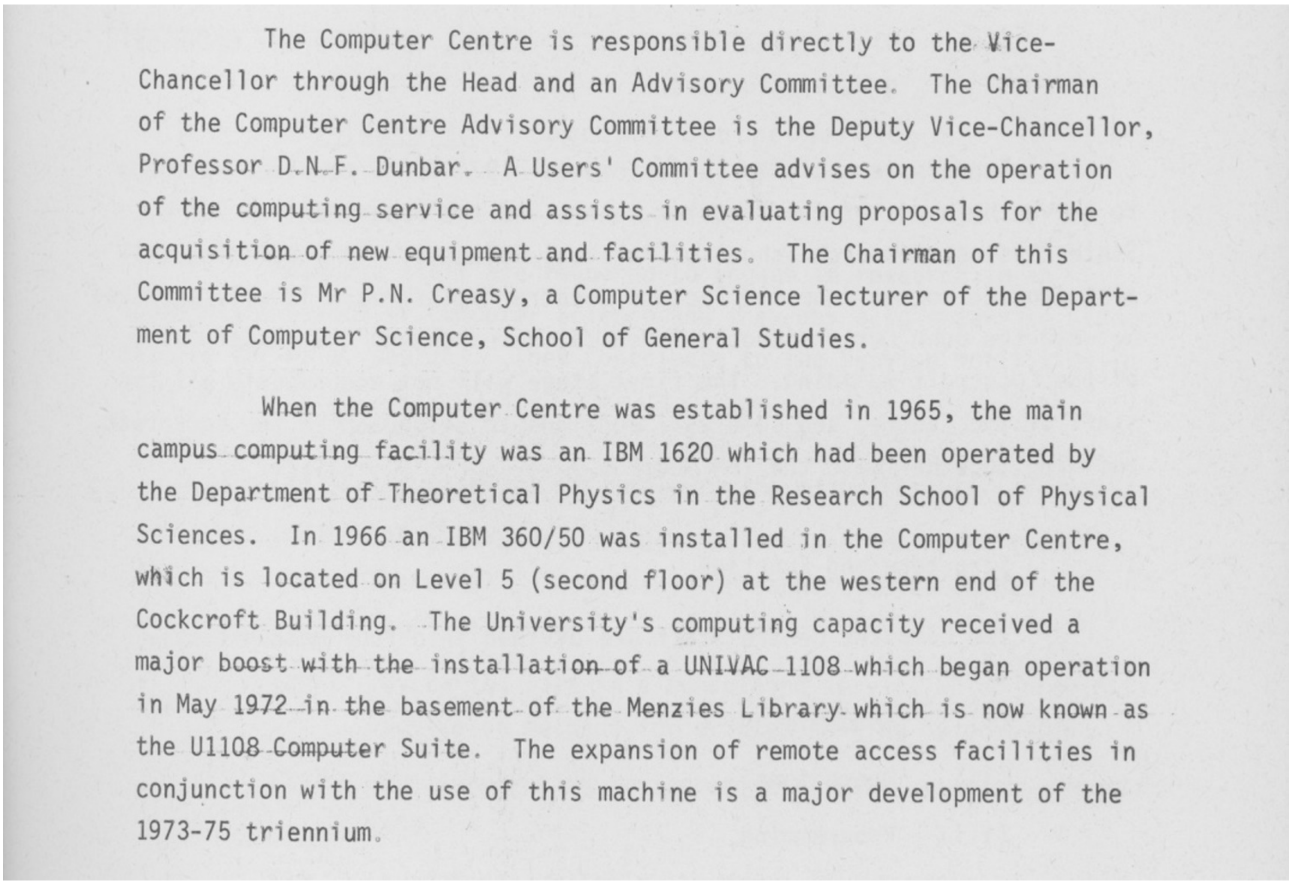

The Computer Centre is responsible directly to the Vice-Chancellor through the Head and an Advisory Committee. The Chairman of the Computer Centre Advisory Committee is the Deputy Vice-Chancellor, Professor D.N.F. Dunbar. A Users' Committee advises on the operation of the computing service and assists in evaluating proposals for the acquisition of new equipment and facilities. The Chairman of this Committee is Mr P.N. Creasy, a Computer Science lecturer of the Department of Computer Science, School of General Studies.

When the Computer Centre was established in 1965, the main campus computing facility was an IBM 1620 which had been operated by the Department of Theoretical Physics in the Research School of Physical Sciences. In 1966 an IBM 360/50 was installed in the Computer Centre, which is located on Level 5 (second floor) at the western end of the Cockcroft Building. The University's computing capacity received a major boost with the installation of a UNIVAC 1108 which began operation in May 1972 in the basement of the Menzies Library which is now known as the U1108 Computer Suite. The expansion of remote access facilities in conjunction with the use of this machine is a major development of the 1973-75 triennium.

An extract from the 1975 Computer Centre General Information booklet. The IBM 1620 mentioned is the second one acquired by ANU (a Model II) rather than the Model I which went to Mount Stromlo, then to the Department of Computer Science, and then to landfill.

By the end of 1971, ANU's 360/50 was considered to be overloaded. It was supplemented around the end of that year by a PDP-11/20. I have a memory that the latter was initially located in a

[35] `https://en.wikipedia.org/wiki/IBM_Personal_Computer` includes a number of amusing quotes.

stairwell!

Early in 1972, the University purchased a Univac 1108 system intended to replace the 360/50, and provide seven times its computational power. It ran a time-sharing operating system called EXEC-8[36] and supported a number of Uniscope 100 and DCT 500 terminals.

Uniscope displays were green-phosphor cathode ray tubes (CRTs). The positioning of text on the screen was quite unstable and the keyboards had to be pounded to register a keystroke – I used to describe them as "carillon keyboards." Uniscopes were designed to reduce the number of interrupts on the Univac CPU by sending data in blocks rather than character by character. As configured at ANU, blocks were lines rather than entire screenfuls. A triangular start-of-entry (SOE) symbol marked the start of the line, and characters between it and the cursor were transmitted when you banged RETURN.

EXEC-8 would automatically log out terminals which had been inactive for more than about 10 minutes, unless you ran Bernie Elphick's `@WHIFFLE` command. It sat in the background and every now and then printed a mysterious message such as `SYMBIONT ERROR` or `AWAIT/DEACT AMBIGUITY`. As a student you felt you were part of an establishment elite if you knew about the command and were bold enough to use it in a room full of physicists.

If an operator in the computer room was bored, he (only male operators did this) would watch through the glass and wait until you were about to transmit, then quickly send you a console message which wiped out your line of typing. I once responded by sending "Leave me alone" messages to the operator in an EXEC-8 loop, stopping it only when I saw him diving to switch off the printer which logged all console traffic. Normally, we would just make sure that the SOE was on a line (above the bottom of the screen) which would not be zapped. The DCT 500s were screenless terminals which printed everything on to 132-column fan-fold lineprinter paper.

A new computer room had been built in the basement of the Menzies Library to house the Univac and also the PDP-11. These premises were affectionately known as "the Elephant Stables" because of the arched vaults used in the undercroft of the library.

The basement "Elephant Stables" under the Menzies Library which housed the first Univac installation.
Photo: David Hawking

[36] Ray Jarvis used to joke that time-sharing was a convincing illusion that all the other users had the machine to themselves.

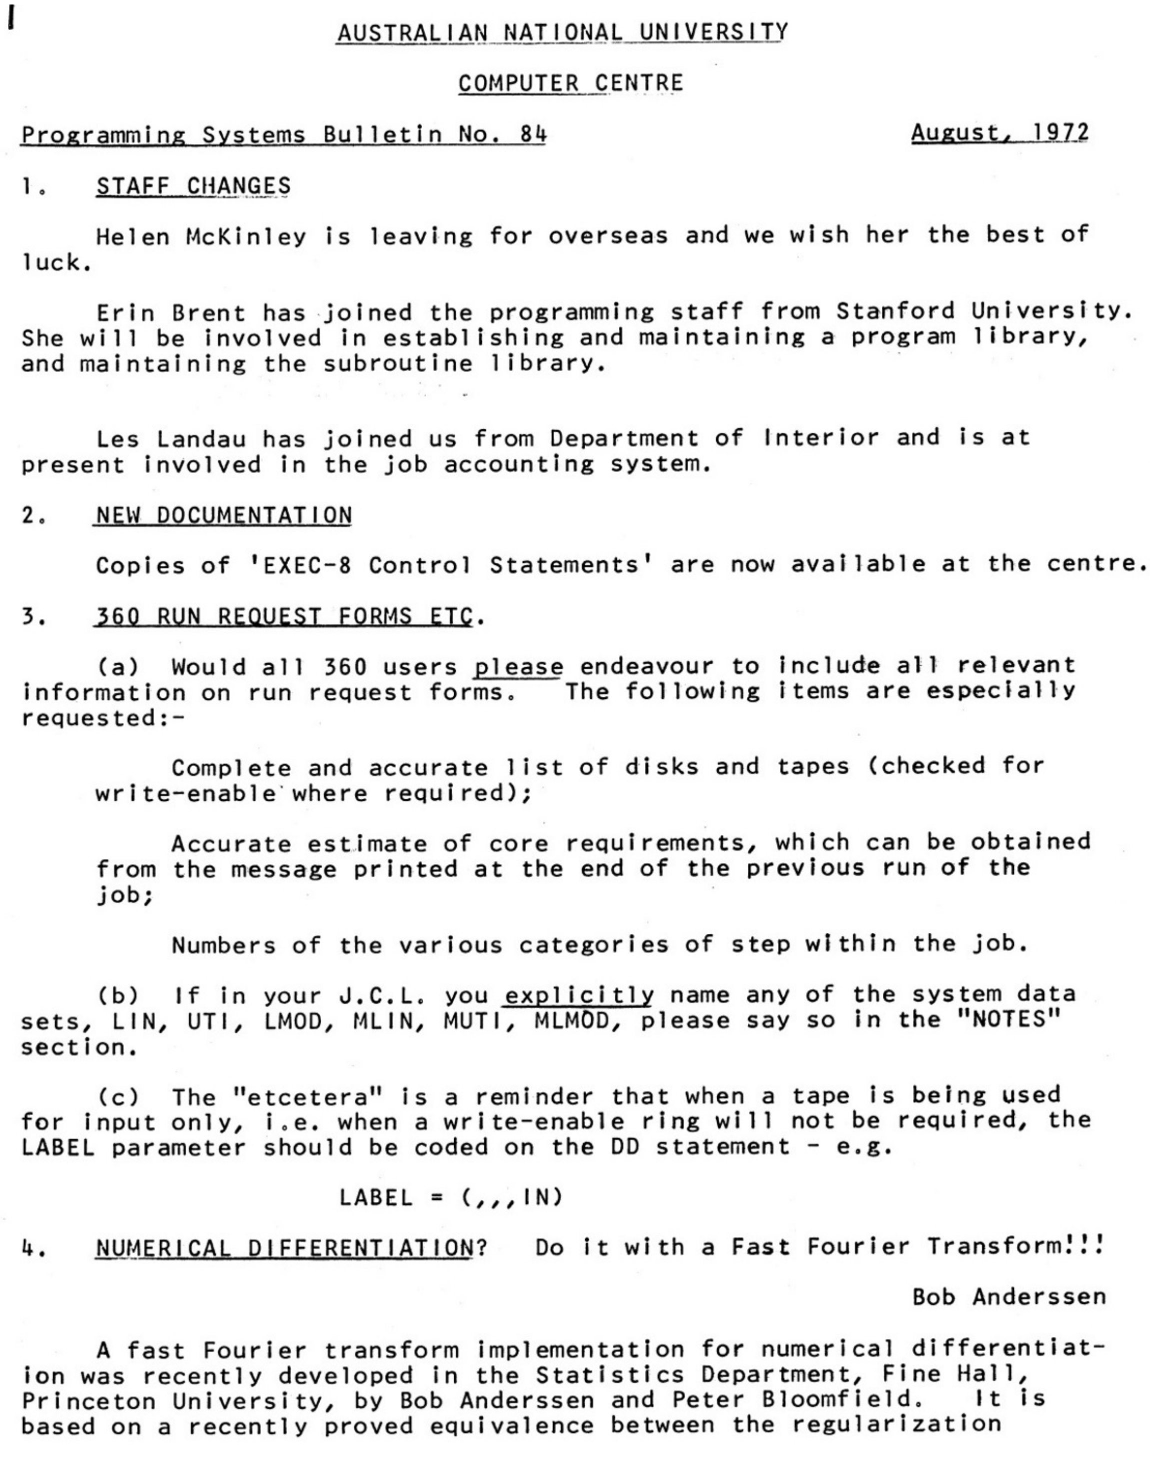

AUSTRALIAN NATIONAL UNIVERSITY

COMPUTER CENTRE

Programming Systems Bulletin No. 84 August, 1972

1. STAFF CHANGES

Helen McKinley is leaving for overseas and we wish her the best of luck.

Erin Brent has joined the programming staff from Stanford University. She will be involved in establishing and maintaining a program library, and maintaining the subroutine library.

Les Landau has joined us from Department of Interior and is at present involved in the job accounting system.

2. NEW DOCUMENTATION

Copies of 'EXEC-8 Control Statements' are now available at the centre.

3. 360 RUN REQUEST FORMS ETC.

(a) Would all 360 users please endeavour to include all relevant information on run request forms. The following items are especially requested:-

Complete and accurate list of disks and tapes (checked for write-enable where required);

Accurate estimate of core requirements, which can be obtained from the message printed at the end of the previous run of the job;

Numbers of the various categories of step within the job.

(b) If in your J.C.L. you explicitly name any of the system data sets, LIN, UTI, LMOD, MLIN, MUTI, MLMOD, please say so in the "NOTES" section.

(c) The "etcetera" is a reminder that when a tape is being used for input only, i.e. when a write-enable ring will not be required, the LABEL parameter should be coded on the DD statement - e.g.

LABEL = (,,,IN)

4. NUMERICAL DIFFERENTIATION? Do it with a Fast Fourier Transform!!!

Bob Anderssen

A fast Fourier transform implementation for numerical differentiation was recently developed in the Statistics Department, Fine Hall, Princeton University, by Bob Anderssen and Peter Bloomfield. It is based on a recently proved equivalence between the regularization

ANU Computer Centre: A representative extract from a *Programming Systems Bulletin*.

I think the Univac cost about a million dollars (about ten million 2019 dollars). It ran off a motor generator which delivered 60 Hz alternating current, 200 amps at 200 volts. A huge cooling system was imported from the USA to meet Univac's specifications. That had to be done twice because the first one fell off a truck and was irreparably damaged. The cooler had six compressor modules. An ANU engineer told me that it had been massively over-specified and consequently caused dangerously rapid temperature and humidity fluctuations. Compressors were progressively switched off until a stable environment was created with only one of the six compressor units operating.

The Univac was run by the Computer Centre whose director in 1972 was Mike Osborne. It comprised an academic section, including Bob Anderssen (famous for wearing shorts with sandals and socks in all weathers), D.E. Lawrence, Richard Brent, Bruce Millar, P.M Fenwick, Col Jarvis, and C.K Yuen; a team of systems and applications programmers, and a team of operators who staffed the machine 24/365. PhD students included Peter Creasy, and Frank de Hoog. Programmers included Julie Bakalor (Senior Programmer, taking over from Bruce Butterfield), Bernie Elphick, Erin Brent (the Program Librarian), Peter Tindale, Les Landau, David Seddon, and Ian Simpson. Ted Hempstead and Rex Jones were technical officers. Operations Manager was Alan Harris, Margaret Brinkley was Computer Supervisor, and operators included Keith Burrows, S. Campbell-Jones, B.J. Douglas, Keith France, David Jones, Cher Chung (then Cher Lek), and Errol Rumpf.[37] John Alsford, Val Airey, Inta Skriveris, Dorota Janiszewska, Bob Cremer, Lindsay Morrison, Bob Fraser, David Jobson were among the operators arriving later.

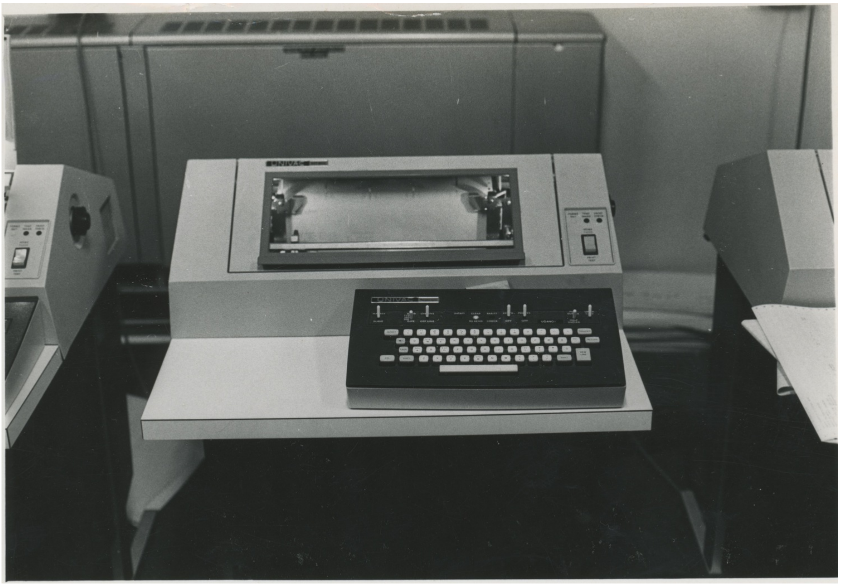

A Univac DCT 500 printing terminal, part of a cluster in G Block, across the road from the Menzies Library. *Photo: David Hawking*

Keith France has had an important connection with ANU computing over four decades. While working as an operator, he studied accountancy part-time. When Bob Watts, then Director, Computer Services Centre (CSC), refused him time off for study, he joined University Accounts. He later returned as Executive Officer at the CSC and managed the finance for the purchase of the Fujitsu VP100. In 1988, he became Business Manager of the ANU Library and was instrumental in the purchase of a McDonnell Douglas server and new library system. He enticed Erin Brent to come to the Library to lead the Library Computer Unit, implement the new system, and connect it to the ANU network. The new server was powerful enough to run the latest library system, Innopac, and enabled the library catalogue to go online for the first time, including the extensive Asian collection. Next they installed a CD-ROM server, which gave online access to many scientific journals. Then, in an Australian first, they collaborated with Reuters to download newsfeeds and made them available to academics at their desktops. In return, the Library server provided Reuters with a backup of their database in Australia. In 1993, Keith set up the business systems for the ACSys CRC.

An elephant in the room: Keith came to Canberra in the late 1960s to work in the Public Service. He lived in Reid House (one of Canberra's several public service hostels) which was then separated from Civic by a large area of open ground. One night in 1969 a group of friends returning from the Civic Pub found a circus set up on the open ground. Tied up there was a baby elephant which was discovered to have a taste for beer. It was persuaded to follow them back to Reid House and to take up residence in the room of a resident who was away for the night. A degree of fuss and consternation was evident in the hostel the next morning! (And no doubt at the circus.) [I didn't enquire whether the elephant was pink.]

Another figure often to be found around the Univac was Bill Craig, a programmer from the Coombs Computing Section and member of the CSC User Committee. He was very short sighted,

[37] According to 1972 Computer Centre Annual Report, accessible from `https://openresearch-repository.anu.edu.au/handle/1885/149187`

wearing glasses with thick glass lenses. One night when he set off to travel home from the staff centre, conditions were rather foggy (or he was?) He unfortunately crashed his bicycle into a ute parked on the side of the road and finished up on the tray top.

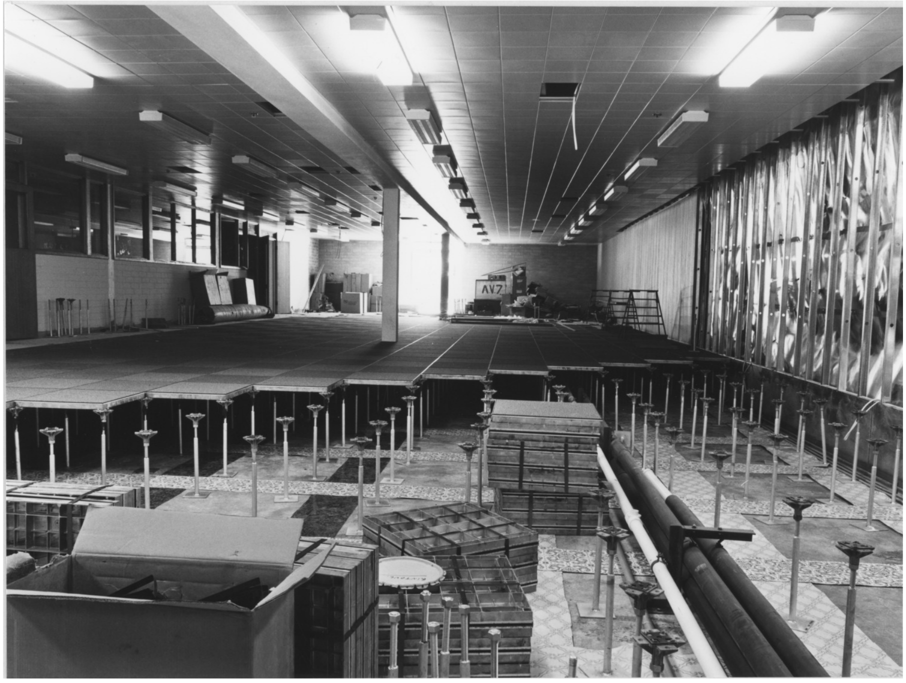

1975: Installing the false floor in the new Huxley computer room in preparation for the Univac 1100/42. (The building wasn't officially called the Leonard Huxley Building until 1987.) ***Photo from A Pictorial History of the CSC.***

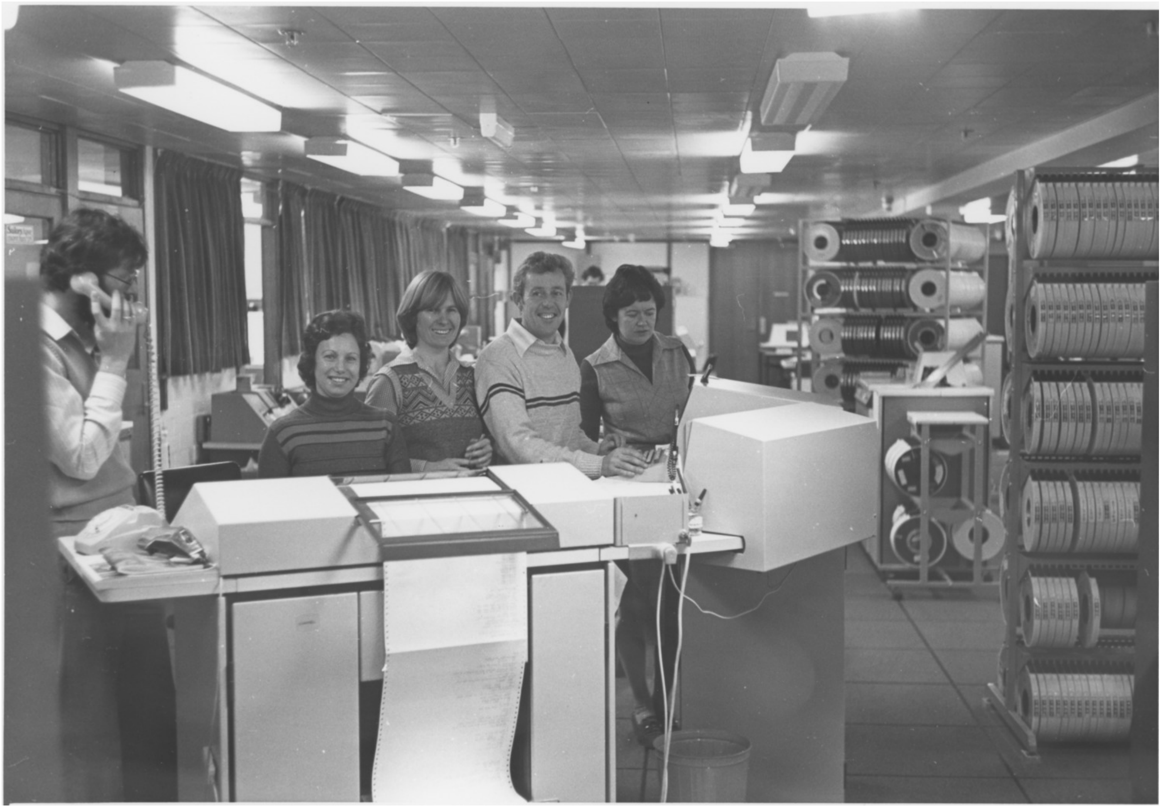

1975: Clustered around the Univac 1100/42 console. Les Landau (on phone), Nancy Imbriano, Dale Minchenton, Roger Harris, Margaret Brinkley. Note the racks of magnetic tapes. ***Photo from A Pictorial History of the CSC.***

1982: Univac salesman Walt Heuer in exuberant mood at the 10 year celebration of the Univac ANU relationship. David Jobson at left. *Photo from A Pictorial History of the CSC.*

1987: Keith France and Harriet Michell at the water boiler on the occasion of Martina Landsmann's farewell. *Photo from A Pictorial History of the CSC.*

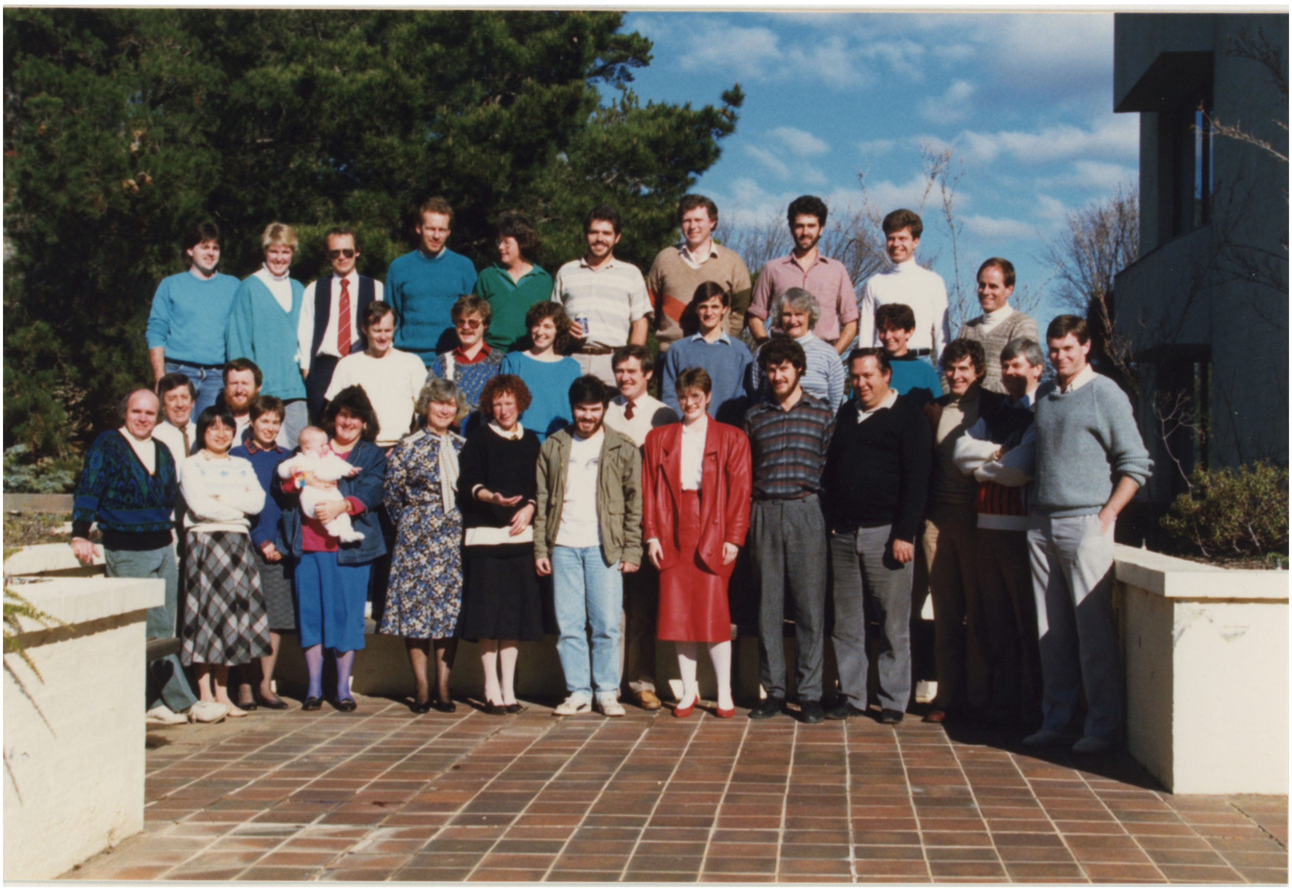

1987?: Group photo at Martina Landsmann's farewell: Back row: Brian O'Rourke, Sonya Rutter, Keith France, Rob Ewin, Harriet Michell, Peter Elford, Geoff Huston, Rob Davy, Jeremy Webber. Middle: Kevin Dawes, Peter Keast, Geoff Barlow, Jeff Brindle, Kathy Handel, Robin Erskine, Andrew Purdam, Gloria Robbins, Virginia Woodland, Peter Olsen. Front: Bob Cremer, Shu Ung, Heidi Lehtis, X + baby, Colleen Harper, Pat Lisle, Steve Ball, Martina Landsmann, Eric White, Brian Jaggers, David Hawking, David Jobson, David Kenny. *Photo from A Pictorial History of the CSC.*

Melissa Waterford and Virginia Woodland, Microcomputer Information Unit, at the 1987 CSC Xmas Party.
Photo from A Pictorial History of the CSC. Speech bubble: Anonymous.

1975: Resembling a laundrette, a row of removable disk drives for the Univac 1100/42. ***Photo from A Pictorial History of the CSC.***

Dean Spire was an on-site engineer for Univac. He was an American and wore flamboyantly stitched leather cowboy boots. He frequently fixed problems with Uniscope terminals by thumping them mightily on the side, and will be ever associated with the statement, "Ah'm a goin' fer ma hammer!". He later worked as a technician in DCS.

Bernie Elphick gained a formidable reputation as a systems programmer. When the installed version of EXEC-8 proved unstable, he would stride to the console and fix the problem with a masterly "C-keyin". Eventually people understood that he was using the C-keyin to cause the machine to skip over a recently applied patch, ... recently applied by Bernie. Whenever EXEC-8 crashed, Bernie insisted that there should be a full post-mortem dump, printed on masses of line printer paper. Computer Centre staff once gained access to his room and moved all the core dump listings to make walls which reached from floor to ceiling. When Bernie opened the door he found himself in an "office" only just big enough to stand in.

Computer Centre staff were fond of pranks. After programming staff had been exhorted to maintain tidy offices, and Ian Simpson had wilfully refused to comply, the operators covered all the items on his desk with a huge sheet of brown paper. Unperturbed, Ian left the paper there and started a new layer on top. Ian drove a tiny Fiat Bambino and one night parked it at the Staff Centre (Canberra House on Acton Peninsula). Other Computer Centre staff lifted it up to the top of a flight of stairs. When questioned the next day about how he got home, Ian insisted that he'd gotten into his car and driven off.

Sometimes school tours were shown through the Huxley computer room. The drama of the tour was often spiced up when Greg Preston, dressed up as Count Dracula, would suddenly rise up from beneath the false floor.

Users of the Univac were given a user code consisting of three letters (your initials) and three digits (the accounting code for your area of the university). The life of these codes extended way beyond the life of the Univacs. All regular users knew who GSL101 and RAG123 were. I've recently become aware that CWJ659 and GIH902 are in current use in email addresses.

One regular user drove an old Toyota Land Cruiser and normally parked it near the hall of residence, where he was a sub-warden. A week after reporting it stolen, he found it at the Computer Centre from where he had walked home after a night on the Uniscopes.

Computer programming was taught as a service course for ANU staff and students prior to 1971. Brian Robson and Iain MacLeod both taught such courses in the 1960s. The Computer Centre continued to run FORTRAN courses for some years after undergraduate teaching of computer science commenced.

When I arrived at ANU in 1971, Pure Maths and Applied Maths were separate departments.

Friends who enrolled in Applied Maths were required to write a FORTRAN program to perform integration of a function by Simpson's Rule. I had chosen Pure rather than Applied because it was the pre-requisite for Computer Science B01 the next year. We had the good fortune to be lectured to by Hanna Neumann, after whom a building was named upon her untimely death.

Another Pure Maths lecturer was Martin Ward[38] who was a computing evangelist. There were no computing assignments but Martin encouraged us to program. We had a competition to find the greatest number of prime numbers within the ten second time limit. Martin also formed an out-of-hours club to study Conway's Game of Life, and we all set to work programming gliders and glider guns. Martin provided us with very strong instant coffee from an urn. When Hanna Neumann walked past he offered her a cup. She said, "Aach, you know I like it strong!". We thought she would be very happy with the strength we were given but instead she shovelled four more heaped teaspoons of instant into her small cup. Martin Ward and Neville Smythe were later instrumental in the signing by ANU of an Apple Consortium Agreement which led to the arrival of thousands of Apple Macintoshes on campus.

There were Pure Maths Department card punches and job submission trays on the top floor of the building which for several decades was known as the Hanna Neumann Building and which is now part of the P.A.P Moran Building. When the original Hanna Neumann Building was refurbished in the early eighties, the mathematicians demanded that the coffee tables in their tea room have surfaces which could be easily written on with a pencil and easily erased with a damp cloth.

Once a day, one of the building superintendents would take the submission trays to the IBM 360/50 in Cockroft and bring back the previous lot, along with the fan-fold line printer output. I recently received the confession of a friend, Philip Allnutt, who used to repeatedly submit a card deck whose purpose was to cause the printing of twenty pages of blank paper for writing, on the grounds that a ream of foolscap[39] was quite expensive in those days!

John was the Building Superintendent for Economics Faculty and there were two superintendents for the Arts Faculty. John was close to retirement and, in the afternoons, could sometimes be found asleep in an armchair in front of *Days of Our Lives* on a small black and white television on top of a filing cabinet. In the Arts Faculty, Jeff was a communist who gave out leaflets in Garema Place during late night shopping on a Friday night, and Tom was a former London Bobby.

Our programs were run through MFFT, the Mississippi Fast FORTRAN Translator. We didn't have to worry about the very arcane OS/360 job control language (JCL). Each of our jobs started with a green `$JOB` card, and ended with a red `$END` card. In some cases there was an orange `$DATA` card to separate the program from its data. The computer couldn't tell the colour of the cards, of course, but colours were needed by the operators and the students.

When you made syntax errors, MFFT printed error codes such as NA031 (Syntax error) and, if you triggered a runtime exception, printed an ABEND (Abnormal End) code such as OC5 (address bounds error). There were multiple volumes in the 360/50 room listing all the possible ABEND codes. If you knew the trick you could get MFFT to print out a list of all its error codes. For NA031 there were about 20 possible explanations, such as "too many left parentheses", and "too many right parentheses". The last two in the list were "more than one of the above", and "none of the above".

I should point out that EVERYTHING WAS WRITTEN IN UPPER CASE – the line printer and card punches did not support lower case. The Univac 1108 represented everything in the 6-bit FIELDATA code. There were no lower-case letters in FIELDATA, but on the plus side you could fit six FIELDATA characters into a 36-bit word. DEC-10s also used 36 bit words, but packed five seven-bit ASCII characters into the word, leaving an extra bit as a flag. Both Univac 1108 and DEC-10 provided hardware instructions to access the relevant word fractions. IBM computers represented character information in the EBCDIC code, in which each character was represented in an 8-bit byte. The EBCDIC codes for letters were not contiguous, which sometimes caused annoyance.

[38] Even in his seventies, Martin was still leading Himalayan expeditions.

[39] Yes, in 1971 foolscap and quarto paper sizes were in common use, ahead of the ISO 216 paper standards such as A4.

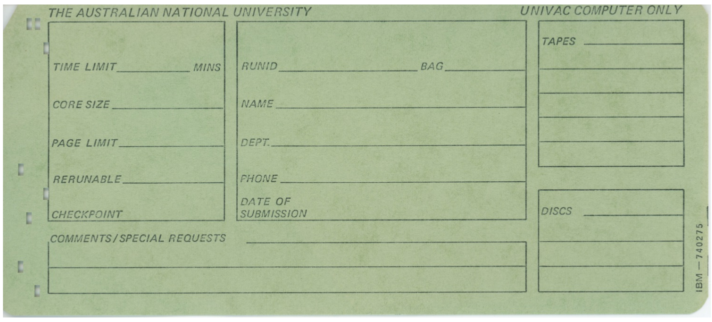

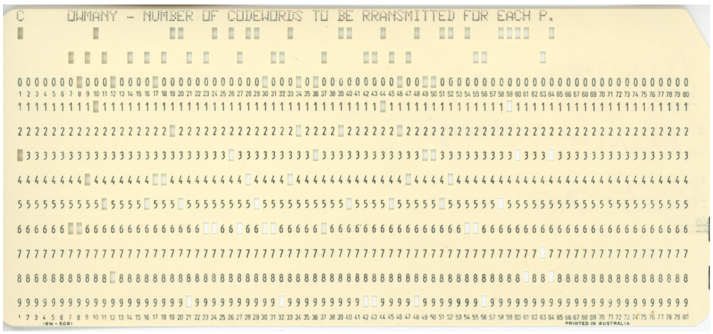
C DWMANY - NUMBER OF CODEWORDS TO BE RRANSMITTED FOR EACH P.
IBM - 5081
PRINTED IN AUSTRALIA

Top: A pre-punched `@RUN` card for the Univac. Every valid deck was required to start with a green `@RUN` and to finish with a red `@FIN`. Bottom: A FORTRAN comment card (a 'C' in column 1). Note that the card is "interpreted", i.e. the content is printed along the top of the card. Richard Brent recalls that it was possible to bypass the system check for @RUN and @FIN cards. If you didn't like someone whose cards followed yours in the input trays, you could get the system to skip through their cards. This wasn't done very often, because Mike Osborne took a dim view of it.

2.4.1 1970s: Evolution of ANU Computing Facilities

The Computer Centre Bulletin of Nov-Dec 1975, reports the replacement of the Univac 1108 by a dual-processor 1100/42. Theoretically, this was a substantial upgrade, but performance was choked by a lack of primary (semi-conductor) memory. Univac struggled for months to pass the acceptance benchmark. Rumour has it that the benchmark only passed months after installation when the SPSS licence expired.

The reason ANU didn't rush out and buy more primary memory was that in 2019 dollars, it cost approximately $5 per byte! On that basis, the laptop on which I am typing this would be worth $320 billion dollars!

By the second half of the 1970s many parts of ANU had their own "school computer".

Research Schools of Social Sciences and Pacific Studies The Joint Schools maintained a DEC KA-10 computer in the Coombs Building. It was operated by Les Dohnt and Catherine Thorpe. The programming team, headed by Mary Rose, included Anne Sandilands, Monica Berko and Brian Pearce, with Yvonne Pittelkow as Statistical Consultant. In 1979, the KA-10 was moved to the School of General Studies and replaced by a KL-10 which was roughly four times faster.

Around this time an academic in Joint Schools imported a program written in Burroughs Algol. Brian Pearce semi-automatically translated it into Simula 67 for the KA-10. Once compiled it took 8 hours to run. After using the Simula 67 code profiler, Brian found that 99% of the runtime was taken by a single block of the code. After rewriting that in assembler, the runtime dropped to 4 minutes!

Research School of Phyical Sciences Another DEC KA-10 was located in the Oliphant building, supported by a team comprising Roger Brown, (later replaced by Peter Cohn, and eventually Shiu Tin), Susan Murray, and Julie Dalco. I think Iain MacLeod still ran a PDP-15 in Engineering Physics with Arthur McGuffin and Joe Elso in support.

Printing technology was very primitive in the seventies. Most printers on campus were only capable of printing black letters on white paper. Iain MacLeod wrote a program for printing greyscale images on a lineprinter. The value (in 4 or 5 bits I think) of each pixel was represented by different combinations of overprinted characters. I modified it to suit the different characters available on the printer I used and used it for all the images in my 1974 Honours thesis on Image Processing.

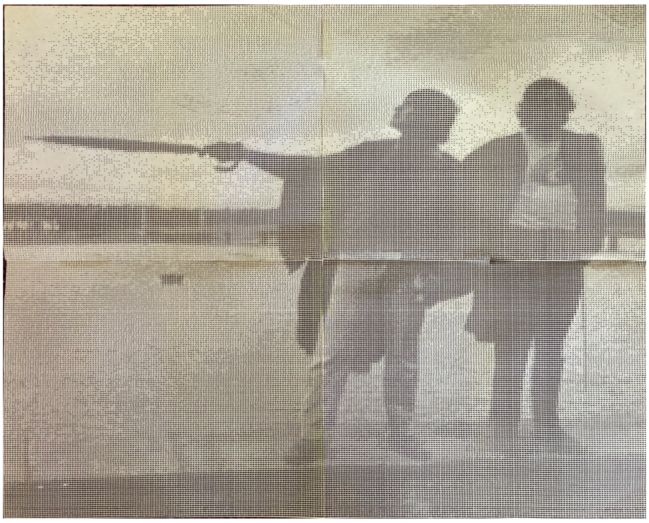

A "high-resolution" (approx. 350 x 250) image rendered on four line printer sheets by Iain MacLeod's greyscale renderer, and used in my 1974 Honours thesis. ***Original photo: Huw Price.***

Bill Caelli tells an amusing story[40] of the IBM 1800 DACS (Data Acquisition and Control System) installed in the Department of Nuclear Physics in 1968, and with a high-speed "channel-connect" to the IBM 360/50. A departing post graduate student inserted code to print a funny message into the program that the research relied on. Bill asks whether this could have been the first malware.

Mount Stromlo and Siding Spring Observatories In 1974, the 3.9m Anglo-Australian Telescope was commissioned at Siding Spring. It is said to be among the first telescopes designed to be controlled by computer.[41] The computer used as controller was an Interdata 70, and the control program was a deck of 30,000 cards (15 boxes!). In 1978, the Mount Stromlo Observatory installed Australia's first VAX/VMS system, a VAX 11/780. I was not as familiar with the Stromlo team but I believe Peter Quinn was the site manager there, succeeded by Roger Brown.

Research School of Chemistry Malcolm Bruce ran a PDP/11 installation for RSC. At some stage he was joined by Patrick Keogh, who studied computer science at ANU. Patrick installed a version of Unix on a machine in RSC; I think that was a first for ANU. At some point Pam Cohen joined the RSC team and eventually took over from Malcolm Bruce.

Research School of Biological Sciences RSBS ran a PDP-11 system managed by David Sandilands, which was later upgraded to a VAX.

Research School of Earth Sciences In 1973, Anton Hales, newly appointed Director of RSES, had purchased a Datacraft 6024, serial number 5, for use in seismic analysis. He saw it as offering superior bang-per-buck because it had a 24-bit word length, and 48-bit double precision, sufficient for the calculations he wanted to perform – and you only had to pay for 3/4 of the memory of an equivalent 32-bit machine. The problem was that the nearest Datacraft office was in Fort Lauderdale, Florida. Ken Muirhead, an engineer employed by RSES, was responsible for installing and maintaining the hardware, while I, as a vacation student, was for a while responsible for the software. I wrote a plotting library to support a Varian electrostatic printer/plotter.

[40] https://ia.acs.org.au/article/2018/australian-first-trojan-horse.html

[41] https://maas.museum/observations/2012/10/02/a-visit-to-siding-spring-observatory/

> The Datacraft included a lot of discrete componentry. When the computer failed to boot (off magnetic tape) Ken was eventually able to diagnose a fault in the BNK (Branch on Negative K register) instruction and fixed it by soldering in a new transistor.
>
> Tom Fitch, an RSES researcher, used to regularly take a 2000-card earthquake location program over to the IBM 360/50 with a couple of cards worth of seismic data, ready to be able to answer questions from the Canberra Times following any seismic disturbance during the night. Tom himself was badly shaken when, one day, his tray of cards came back with a scribbled note, "Sorry, we dropped these."

Acquisition of the Univac 1108 brought support for interactive terminals (in addition to the system console), but initially you had to go to a Computer Centre terminal room to use them. Chapter 11, on ANU networks, describes how interactive access gradually spread across campus, and how machines were gradually networked to each other and to machines across the world.

2.5 1970s: Establishment of ANU Computer Science

Computer Science started life as a sub-department of the Department of Statistics. Another Statistics sub-department, that of Econometrics, started at around the same time with Deane Terrell as Chair.

From the report by the Dean of the Faculty of Economics in the 1970 ANU Annual Report:

> "In 1970 it had been hoped to introduce courses in computer science, for which there were indications of a strong student demand. However, it was found shortly before the beginning of First Term that no satisfactory teaching appointment could be made. Following readvertisement, two posts have since been filled; funds have also been made available for the remote access equipment needed for teaching. There now seems little doubt that Computer Science I will be taught in 1971. To this, Computer Science II will be added in 1972. The Department of Statistics is at present administratively responsible for both of these units."

Enrolments in computer science at ANU in the 1970s were substantially boosted by public servants taking advantage of the three hours a week of paid study leave to which they were entitled. There was virtually no outsourcing of IT in the public service at that time, and Commonwealth departments maintained large in-house IT sections. Many IT staff in the public service were trained through the Programmer-In-Training (PIT) course at the Canberra College of Advanced Education (CCAE), later to become the University of Canberra (UC). Some public servants were keen to upgrade this diploma qualification to a university degree.

As previously noted, I was an ANU undergraduate from 1971 but had no interaction with the Computer Science sub-department until 1972.

2.6 1972: Computer Science Computing Facilities

When Ray Jarvis arrived at ANU in 1971 as a Senior Lecturer in computer science he obtained funding to purchase a Data General Supernova computer and to fit out a computer room for it in one of the internal divisions of Copland G10.[42]

Data General was a spin-off from Fairchild Semiconductor and the Supernova was a mini-computer competing with PDP-11s on value-for-money. It had 16KWords (32KBytes) of core memory, a head-per-track disk with capacity of half a million words (1 MB), plus analog-digital interface, card-reader, line-printer, paper-tape reader, and Sykes cassette drive. The Supernova was distinguished from lesser Nova models by hardware multiply and divide instructions (integer of course), and by a Channel Start feature. The latter allowed you to deposit a single 6-digit octal number via the front panel

[42] G10 was on the ground floor at the South-East corner of the Copland quadrangle. It's now labeled G46, and houses ANU's *Respectful Relationships Unit*.

switches into a special address and use the `Channel Start` switch to start a bootstrap loader on disk. Without channel start, 14 octal numbers were needed to boot from paper tape.

The furniture for the Supernova room was a garish orange and the carpet an ultra-cheap sickly green. Static generated from the carpet would make you yelp, and make the line printer to print a line, whenever you touched one of its switches.

In 1972, student card decks were run through the Supernova card-reader by Maxine Esau, the department's only technical officer. Ray gleefully reported that the new Univac was only twice as fast as the Supernova because its optimising FORTRAN compiler was so big and slow. Students doing third year projects could book dedicated time on the SuperNova.

There were two or three IBM 029 card punches in an ante room to the Supernova room. That was woefully inadequate for 260 second-years plus many third-years. Many of us used punches in the Pure or Applied Maths departments and I remember being part of a delegation to Chip Heathcote to complain about this and other perceived shortcomings. Eventually a hand-held rotary card punch was provided to allow students to more quickly correct their decks.

The facility for processing student jobs was quite unreliable. The Supernova card reader was reasonably fast but apparently not designed for the very heavy loads. The Supernova itself failed often enough and there were some spectacular crashes. In each case, Ray, Peter or Maxine would sticky-tape a green computer card on the door with a hand-written "System Down" notice. In one case the motor of the disk drive burned out and Ray Jarvis had it re-wound by a business in Queanbeyan. There was a period of great uncertainty when the motor came back while Ray found out which way the disk was designed to rotate. On another occasion the Supernova blew up in such dramatic style that the service company elected to replace it with a new Nova in lieu of repairing it. Despite this an engineer from the company eventually repaired the Supernova in his own time. He kindly returned it and 8Kwords of operational RAM (out of the original 16K) to Computer Science.

An IBM 029 Card punch. (Unfortunately, not an ANU one.) ***Photo courtesy of IBM, ©International Business Machines Corporation.***

2.7 1972: Computer Science Academic Staff

Ray Jarvis was a very good lecturer and a true character. Decades later, he told me that his algorithm for convex hulls,[43] published in 1973, was the academic achievement of which he was most proud.

He was born in Burma and studied at the University of Western Australia, completing a PhD in robotics in 1968, after which he spent two years at Purdue University in the USA, before joining ANU. He drove a slightly decrepit E-type Jaguar sports car,[44] which he always parked at the East end

[43] The gift-wrapping algorithm. `https://en.wikipedia.org/wiki/Gift_wrapping_algorithm`

[44] According to Wikipedia, Enzo Ferrari called it "the most beautiful car ever made."

of the Copland Building, only 20 metres or so from G10. He routinely arrived at about 10.20am in order to miss the peak of the traffic. ... and in those days he always got a park!!!! He smoked a pipe. (Smoking was permitted in ANU buildings until about 1990.) He was an inveterate DIY engineer and in later years built a telephone answering robot from FischerTechnik. Ray took up a Chair in Electrical and Computer Systems Engineering at Monash University in 1985 and gradually migrated more into Engineering.

Tragically, Ray died in 2013 from mesothelioma contracted working in the asbestos mine at Wittenoom while a student. Readers are referred to an obituary[45] and a late-life media interview.[46]

Ray let me look after his E-type while he went on a sabbatical year. This was very exciting as I am in full agreement with Enzo Ferrari. I got to run the engine every month, and he let me ride in it once or twice, but I never drove the car.

Ray bought all sorts of bargains from local electronics shops including a Tandy TRS-80 computer, and a speech input/output system, also from Tandy/Radio Shack. He connected the latter up to the Supernova and developed some software to recognize spoken digits. He said that he was training the speech recogniser but I thought that the system was training him. Amusingly, his speech recogniser was active one day when I came into the inner sanctum. It said, "Six!", at the sound of the door slamming behind me. Geoff Huston reminds me that one year Ray bought a large supply of cheap steel which he thought would be useful for building equipment racks. It was stored underneath Haydon-Allen Tank. Perhaps it's still there.

In 1974 Ray ordered a hardware device called a Megatek which was a vector graphics generator. I wrote an interrupt handler and other software for it, but the device was, as Brian Molinari kept ribbing us, not very successful.

Ray's penchant for interesting gadgets extended to his home life. When his wife Irene sent him on a small mission to a Trash and Treasure market, he returned with a dentist's chair! Presumably not in the E-type.

Peter Creasy, soon after arrival in Brisbane, quickly adjusting to Queensland conditions including giant avocados and mud-crabs. *Photo: David Hawking*

Peter Creasy was also a character. He was employed as a Lecturer and enrolled for a PhD based around implementing a time sharing operating system for the IBM 360 with (I think) Mike Osborne as supervisor. He wrote a software system, called the CreasyLink, which was used by thousands

[45] https://www.monash.edu/__data/assets/pdf_file/0005/710996/vale-ray-jarvis.pdf

[46] https://www.smh.com.au/opinion/of-god-family-mortality-and-the-meaning-of-life-20130913-2tq5c.html

of people. It gave Univac access to users near the Copland Building via Teletype ASR-33 printing terminals and Tektronix 4013 storage screen terminals.

Peter claimed that he could never start a program from scratch and that the CreasyLink program had actually morphed from the time-sharing operating system. This was odd because one was written in assembler for the Nova and the other for the IBM! Each new version of Peter's program, once debugged, was saved as a card-deck or paper tape wrapped in the source code listing and labeled, "VITAL. Use This One".

After a sabbatical in Nivelles, Belgium, Peter embarked upon a new thesis topic in semantic modelling languages and eventually completed it. In the 1980s he moved to an academic position at UQ.

Tony Cantoni. Apart from teaching numerical methods, Tony wrote a utility program which allowed Supernova files to be backed up and restored on the Sykes cassette drive. Tony only stayed one year as a Lecturer at ANU, going on to a stellar career at the University of Newcastle and the University of Western Australia. To quote from UWA News,[47] "'Wherever Tony Cantoni goes, that's the place to be. He is a legend in Australian telecommunications, both within academia and industry,' Professor Fynn said." He has founded a number of start-up companies.

Left: Vicki Peterson, after her return to the department in the 1990s. ***Photo: Departmental archive.*** **Right: Bill Gibson, from the first Fourth Year Honours cohort, on the Great Ocean Walk, after his retirement as CIO of the ATO.** ***Photo: Deborah Gibson.***

Vicki Peterson. Vicki Peterson came to ANU with a degree from MIT and an Astronomy doctorate from CalTech. She had an unenviable work load as Tutor due to large enrolment and huge tutorial sizes in B01. She held the position as Tutor until the end of 1974. (I applied for the vacancy and started in the role in early January 1975.) Vicki returned to the department as a Senior Tutor in 1981 and later became a Lecturer, teaching courses on databases and networks.

> Although I never got a degree in IT or computing or programming, I had used computers a lot in my PhD analysing data and we used FORTRAN.
>
> I don't remember what year that was but at some point I did think that I should start to take this [computer science] seriously instead of thinking about doing Astronomy. I remember that I was spending a lot of time on it so I better start taking it seriously.

[47] https://www.publishing.uwa.edu.au/uwanews/2004/uwanews20040322.pdf

Vicki's background was Armenian, and she contributed delicious cakes to Computer Science morning teas.

Academics in the sub-department were able to call on typing and administrative services from the Department of Statistics. Joyce Radley was the Departmental Administrator for many years. When I became a tutor, Helmi Patrikka did the typing for my courses. Assignments, course notes, and technical reports were typed using an IBM golfball typewriter onto stencils for the Gestetner duplicating machine.

2.8 1972: Launch of Computer Science B01 and B02

As noted, Computer Science was a sub-department of the Department of Statistics in the Faculty of Economics. Ray Jarvis was head of the sub-department, Chip Heathcote was Head of the Department of Statistics, and Peter Creasy was one of the first two Computer Science appointments. Peter was a lovely person but perhaps understandably appeared less than inspired by the content of his own FORTRAN syntax lectures. The tutor was Vicki Peterson. More than 260 students were enrolled and there were only four tutorial groups, giving an average tutorial size of more than 65! I did my best to assist by attending tutorials infrequently.

Algorithms meant flow charts. Flow charting templates,[48] with rectangles, diamonds, triangles and arrows, etc. were on sale in the University Coop Bookshop and good students used them to make neat visual representations of their algorithms. The text described FORTRAN IV which included some extensions to ANSI FORTRAN which we were not allowed to use.

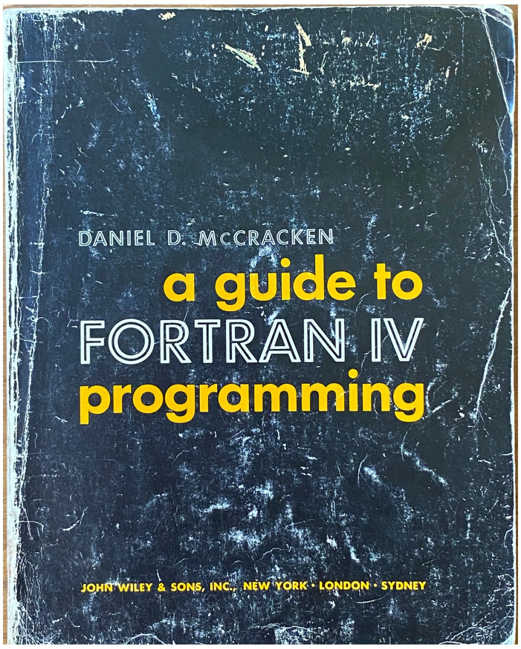

1971–1974: The FORTRAN textbook for Computer Science B01.

I was not a good student. I didn't put in much time, I completed assignments at the last minute (in one case during a large party while drinking cider), I resented the requirement to only use ANSI FORTRAN (e.g. to use Hollerith constants, e.g. `7HFORTRAN` rather than strings e.g. `"FORTRAN"`, and I thought myself superior to the simple programs required. On one occasion I showed my solution to another student in Garran Hall. Without my knowledge, at least four copies of it ended up being submitted and I was one of the two plagiarists penalised.

When I became a tutor, my attitudes changed immediately and I reflected upon how silly I had been as a student. In my first year at ANU I only discovered in Term 3 that Chemistry A01 actually offered tutorials. By that late stage I decided that since I'd managed without them for two terms, ...

[48] See Page 53.

The Joy of FORTRAN

FORTRAN was designed in 1954 by an IBM team led by John Backus, and commercially released in 1957. IBM claims that FORTRAN “may well be the most influential software product in history”.[a] Its longevity in the scientific side of computing was helped by the creation and maintenance of FORTRAN libraries for a host of mathematical, statistical, and physical calculations. Unfortunately two early design decisions led to adverse outcomes:

- In many situations, embedded blanks were ignored. E.g. the command `WRITE` may validly be written `W RI T E`.
- Variable names did not need to be declared. An undeclared variable starting with I, J, K, L, M, or N was implicitly assumed to be INTEGER; All others were assumed to be REAL.

As students we were told that the NASA Venus space probe Mariner 1 was lost because “FORTRAN doesn’t require declaration of variables”, but various sources on the web disagree on the specific nature of the software error. However, the above two decisions did lead to an error I once saw in a student program.

The FORTRAN looping command is `DO`. An example is:

```
DO 10 K=1,100
```

which means, execute the statements between this one and the one labeled `10` one hundred times, first assigning `K` the value 1, then incrementing it at each iteration. The unfortunate student wrote something like the following:

```
   DO 10 K=1.100
   WRITE(3,20) K
10 CONTINUE
```

and asked me for help, because the loop wasn’t working. Only one number was printed.

The explanation is that because of the typo in the DO statement (fullstop rather than comma) the syntax rules mean that it is actually an assignment statement, assigning the value 1.100 to the undeclared REAL variable `DO10K`. The value printed for K is indeterminate because the code fragment shown never assigns a value to it.

Another odd feature of many early FORTRAN compilers, including `@FOR` on the Univac 1108, was the ability to change the value of a constant, as in the following example.

```
CALL MAKE99(7)
J = 7
WRITE(3,20) J
...
END

SUBROUTINE MAKE99(M)
M = 99
RETURN
END
```

The `WRITE` statement should print 7, but using those compilers, the value printed is 99! The call to the subroutine has changed the value of the “constant” 7.

This happened because FORTRAN passed parameters by reference, i.e the memory address of the location containing the value of the constant was passed and the value of 99 was written into it. As a student I was once bitten by this.

Over the past 66 years since its creation, the definition of the FORTRAN language changed radically. Around 1982, Tony Hoare (one of the giants of computer science) was quoted as saying, “I don’t know what the programming language of the year 2000 will look like, but I know it will be called FORTRAN.”[b]

[a] `https://www.ibm.com/ibm/history/ibm100/us/en/icons/fortran/`

[b] Reported in `https://history.computer.org/pioneers/hoare.html`

Second year psychology included a small unit of statistics whose lectures and tutorials I completely missed because of my job as a projectionist. I borrowed the textbook the day before the exam but fell asleep in the Chifley Library. Fortunately, the exam was open book and I got the highest grade possible, leaving after 5 minutes, having found the answers during the reading time, by looking up the index of the textbook.

I became a tutor in early January 1975, shortly before Rob Ewin. The job I applied for was advertised at an annual salary of $6000, but by the time I started, it had risen to $8000. According to the RBA inflation calculator, the latter corresponds to about $58,000 in 2019 dollars. Currently, the range of salary at ANU for Level A academics is $68,000 to $87,000. So, I was hard done by, even after a 33% pay increase!

The merits of "desk checking" were proclaimed by experienced people both within and outside the department. You were expected to be very familiar with the syntax of FORTRAN, and the card column conventions, and to carefully scrutinize your card deck, correcting errors before submitting for processing. It made sense, when turn-arounds were so slow, and machine resources so scarce, but it didn't suit my mental style. I came into my own later, when I could use computer resources to undertake the checking for me. LightSpeed Pascal for the Apple Macintosh reached a peak with its syntax directed editing, automatic indentation, flagging of syntax errors as you typed, and a choice of compiled or interpretive execution.

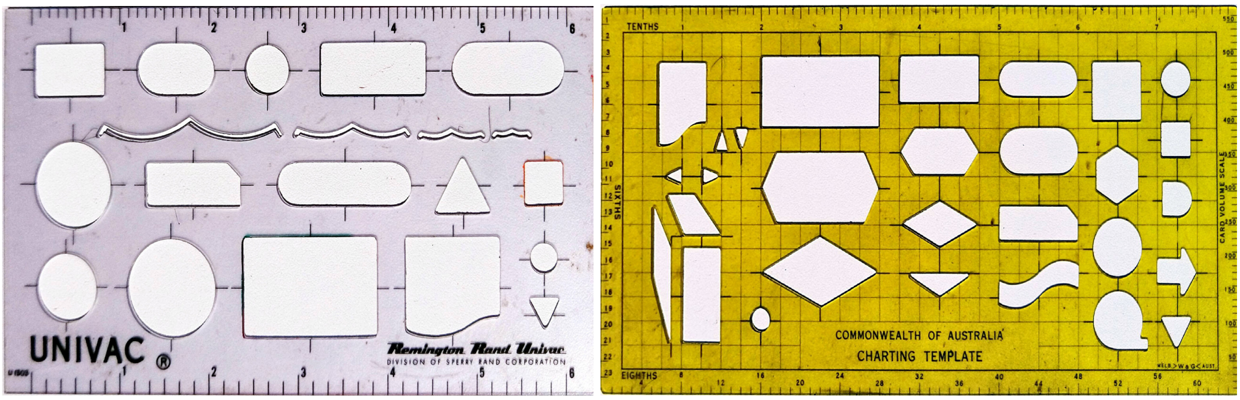

2020: Two interesting examples from Robin Stanton's collection of historical flowcharting templates. The scale on the right edge of the Commonwealth one allows you to estimate the number of cards in a deck. I assume that Robin acquired the Univac one when working for Remington Rand Univac from 1959.

Computer Science B02 was taught by Ray Jarvis and was a lovely mixture of logic design, computer architecture, Boolean algebra, and finite state machines. It also included a small section on Nova assembler programming. It was fairly clear that hardly any of the class "got" assembler programming. Furthermore, the running of the assignments on the Supernova was fairly disastrous. A high proportion of the small proportion of student card decks which assembled without error, caused the system to crash because there was no memory protection. Maxine Esau spent ages endlessly restarting the machine. After I became a tutor in 1975 I wrote a Nova simulator (in Nova assembler) which was used to run the student programs. I knew that the released version of the Nova simulator contained a couple of tiny bugs but it supported courses for several years without once crashing and no-one ever complained. The fact that it ran 80 times slower than the real machine was totally irrelevant.

2.9 1973: Launch of Computer Science C01, C02, and C03

Robin Stanton started the first lecture of C01 with the declaration that, "Algol 60 is the language of the course, but it won't be covered in lectures. There are quite a few good primers in the library." That came as something of a shock when half of the lectures in B01 had been devoted to the syntax of FORTRAN. A considerable burden passed to Chris Johnson as Tutor. The two-hour C01 tutorials became mini-lectures.

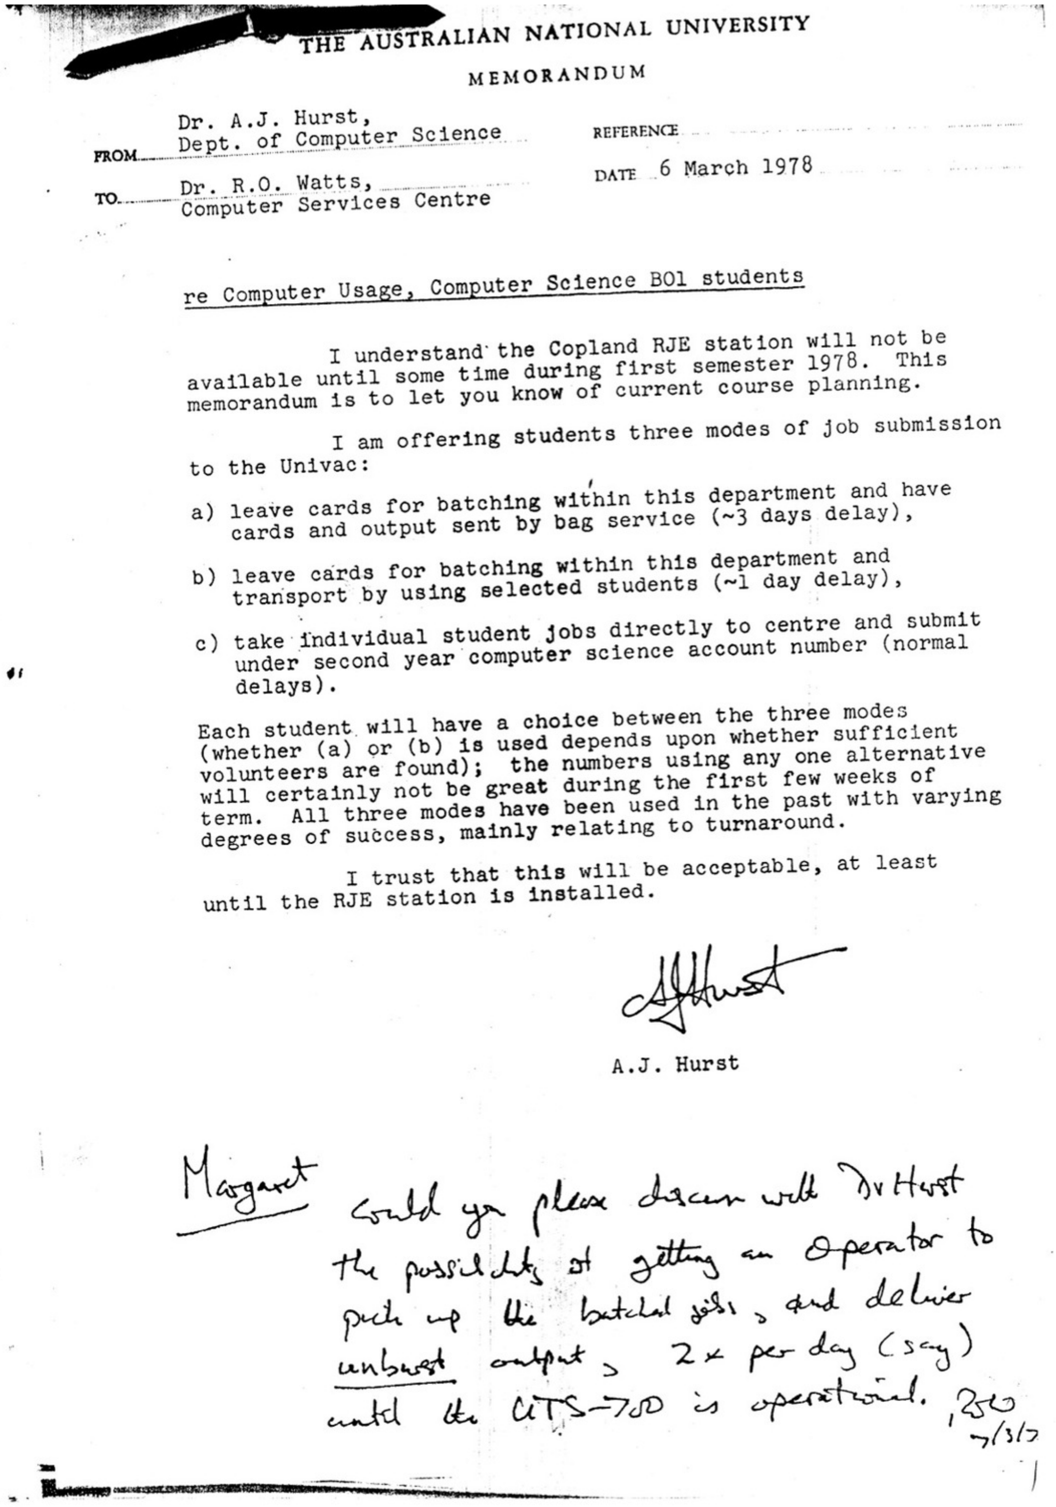

THE AUSTRALIAN NATIONAL UNIVERSITY

MEMORANDUM

FROM Dr. A.J. Hurst,
Dept. of Computer Science

REFERENCE

TO Dr. R.O. Watts,
Computer Services Centre

DATE 6 March 1978

re Computer Usage, Computer Science B01 students

I understand the Copland RJE station will not be available until some time during first semester 1978. This memorandum is to let you know of current course planning.

I am offering students three modes of job submission to the Univac:

a) leave cards for batching within this department and have cards and output sent by bag service (~3 days delay),

b) leave cards for batching within this department and transport by using selected students (~1 day delay),

c) take individual student jobs directly to centre and submit under second year computer science account number (normal delays).

Each student will have a choice between the three modes (whether (a) or (b) is used depends upon whether sufficient volunteers are found); the numbers using any one alternative will certainly not be great during the first few weeks of term. All three modes have been used in the past with varying degrees of success, mainly relating to turnaround.

I trust that this will be acceptable, at least until the RJE station is installed.

AJHurst

A.J. Hurst

Margaret

Could you please discuss with Dr Hurst the possibility of getting an Operator to pick up the batched jobs, and deliver unburst output, 2× per day (say) until the UTS-700 is operational. RW 7/3/7

1978 correspondence between John Hurst and Bob Watts about the Univac RJE for Copland. An earlier draft of the memo had included a couple of more confrontational paragraphs.

The focus of the course was on programming language and runtime structures. The textbook for the course was *Programming Languages, Information Structures, and Machine Organization* by Peter Wegner, and published by McGraw-Hill, which, I recently learned, earned Wegner his PhD from the University of London. The first C01 assignment was to implement a Free Store Allocator (FSA) using NUALG (Norwegian University Algol) on the Univac 1108.

During the course of the FSA assignment, a massive failure of the Univac 1108 occurred, which took weeks to rectify. From memory, the computer room flooded (the drain hole was on the uphill side) and the copper earth mat under the machine got soaked and had to be removed.

Efforts were made to organise processing of student jobs on a Univac in the Department of Defence, but the service wasn't satisfactory and eventually jobs were run instead on the Algol compiler on the Coombs DEC-10, later the SGS DEC-10. This necessitated a change of Job Control Language (JCL) cards and some adaptation of syntax.

The second assignment was to implement Christopher Strachey's General Purpose Macrogenerator (GPM). At our level, this was a very challenging assignment. If implemented correctly, you could use GPM macros to perform arithmetic operations, but to achieve that you had to understand a lot of fundamental concepts and implement them correctly.

In about 1976, the main assignment in C01 changed and required students to implement the DOCTOR version of Joseph Weizenbaum's conversational system Eliza, which used pattern matching to respond to conversational input from a human, specifically a Rogerian psychotherapist.

– "Really, a Rogerian psychotherapist?"
– "Yes, he helped clients by reflecting back what they said to him."
– " That's interesting. He reflected back what they said to him?"
– "This conversation is getting very annoying!"
– "Tell me about your mother."

We third-year students had access to terminals, but were frequently evicted from the Uniscope room in the Computer Suite by members of staff. I remember sitting at a Uniscope next to a Research Assistant from RSPhysS, when he looked over at my screen and angrily said,

> – "A logical IF!!! Don't you realise that arithmetic IFs are 15% faster!"

Imagine my shame, 150 nanoseconds wasted!!

Computer Science C02 was a numerical methods course taught by Brian Molinari. I believe that the text book may have been *Introduction to Numerical Methods* by Peter A. Stark, which was also prescribed for the numerical part of B01. I suspect that the language of C01 was FORTRAN, which in that era was almost universally used for scientific computing. I didn't do that course so details are lacking.

Computer Science C03 was taught by Ray Jarvis and covered picture processing, graphics and scene analysis. There was a choose-your-own-but-get-it-approved programming project, which in my case was a FORTRAN program to perform bidding in the game of bridge. Computer Science C03 Honours was a one-hour per week unit taught by Robin Stanton. It was an introduction to artificial intelligence and required us to become familiar enough with LISP to solve problems like "the eight queens." One topic was, "the function of `function` in LISP." Sorry Robin, I've forgotten what it is.

2.10 1974: Launch of Fourth Year Honours

One student, Peter Kelo, completed a joint Honours degree between computer science and pure mathematics in 1973. Peter worked for Fujitsu at the time of ANU's VP2200 purchase and appears in the photo on Page 201. He subsequently worked for IBM and Telstra.

1974 was the first purely Computer Science Honours cohort. There were four full-time students: Graeme Williams (not Graham Williams who came to the department later), David Hawking, Bill Gibson and Liz Barker. Les Landau, from ANUCC, did his honours part-time, over two years, eventually working for Sperry-Univac (later UniSys) between 1980 and 1996. A contractor after that, he worked for Erin Brent in the ANU Library, and later in other parts of ANU, retiring in 2006. Bill Gibson subsequently went on to a distinguished career in the Australian Public Service, retiring in 2014 after an 11-year stint as Chief Information Officer at the Australian Taxation Office. His photo is on Page 50. Graeme Williams graduated with a PhD from Rochester University, NY in 1979 and worked for five years at Bolt, Beranek & Newman. According to Graeme's LinkedIn profile he:

> Implemented tools for the simulation, performance analysis, design and optimization of ARPANET (Internet) networks, including design of a worldwide network for the DoD.

Graeme later pursued a career in industry in the USA and retired in 2017 in Las Vegas, Nevada. I have regrettably lost track of Liz Barker.

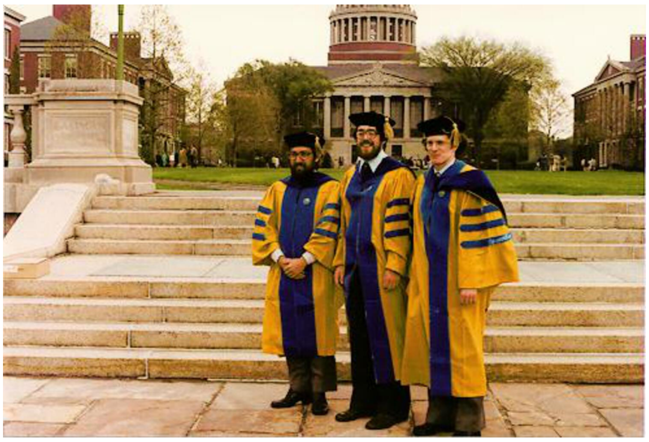

The first three computer science PhD graduates at the University of Rochester, NY, with Graeme Williams at left. His thesis was on Program Checking. Graeme was in the first ANU Computer Science Honours cohort, 1974 Photo from `https://www.cs.rochester.edu/dept/25th/gallery/3_Grads.jpg`

Fourth year Honours consisted of a sub-thesis, counting one-third of the total marks, and two courses each semester, each worth one-sixth of the marks. In 1974, the courses were:

- Artificial Intelligence (Robin Stanton)
- Communication theory (Brian Molinari)
- Compiler Construction (Peter Creasy)
- Computer Architecture (John Hurst)

In 1974, the Honours group formed a Computer Science Club and invited third years. We held lunchtime meetings in a lecture room. In one of the meetings, a member gave a presentation on a programming language which was unfamiliar to others. I have a feeling that it might have been Steve Edwards, presenting on Snobol. Whoever it was showed a program in that language to solve a simple problem. Over the next few meetings, members presented programs to solve the same problem using COBOL, FORTRAN, Algol 60, LISP, PL/1, APL, and probably others such as Conniver, and microPlanner since we studied those in Robin's AI course.

2.11 1975: Re-launch of B01 with ALGOL-W

When John Hurst took over the teaching of B01 in 1975 with Rob Ewin and me as tutors, the course was totally revolutionised. Numerical methods was dropped, and ALGOL W was adopted as the language of the course. Structured programming and top-down design were emphasised. With two tutors and John running at least one tutorial, we were able to dramatically reduce tutorial sizes. An attempt was made to make assignments interesting, and sometimes a choice of problems was allowed.

Students punched programs on cards and they were taken to the IBM 360/50. The mail delivery service which occasionally took the cards and brought back the output was way too infrequent. This led to use of the departmental bicycle, a sturdy Raleigh, ridden by students. It was quite worrying to watch a slightly built person wobbling along with three metal trays of cards (up to 6000 of them) balanced in the basket at the front.

A feature of the B01 punched cards era was the "Rubbish Bin Algorithm." Poor students (I remember but won't reveal the name of the archetypal practitioner!) were believed to pick up a handful of discarded cards from a rubbish bin near the card punches, then add job control cards, run it through the computer and generate a listing full of syntax errors. They would then take it to one of the tutors and ask for an explanation of one of the dozens of errors. Aghast, the tutor would start at the beginning by outlining the algorithm needed to solve the problem. “Yes, yes, I know that but I’m just having difficulty writing it in Algol W.” A limited amount of Algol W assistance later the student would run again with a somewhat improved deck, and repeat the process with the same, or another tutor, or even a lecturer, hoping that the number of iterations required to obtain a working solution didn’t exhaust the patience of all the available victims.

In the first year of the new-style B01, students were given access at the beginning of the course to a time-shared BASIC service run on an 8K Nova in G10. That machine was able to support a small number of Teletypes. The only way that students could save a program from one run to the next was by punching to paper tape. Students who already thought that they knew BASIC were shocked to discover that there were huge differences between one manufacturer’s BASIC and another’s. The BASIC experiment wasn’t considered a success because, after interactive, interpretive BASIC, some students seemed to find it even more difficult to get going with Algol W in batch.

Algol W was the language of the course for three years. Then, in 1978, the introductory course switched initially to Copenhagen Pascal, run on the Univac, and then to a Pascal-P cafeteria supported on a new Nova computer, I think running RDOS rather than DOS.

2.11.1 Consequences for Forestry FORTRAN

The Forestry department were very concerned that B01 was no longer going to teach FORTRAN. I think B01 may have been a pre-requisite within the degree of BSc (Forestry). In consequence Foresters attended FORTRAN programming courses offered by the Computer Services Centre, initially taught by Les Landau. Forestry students could no longer avail themselves of the full range of grades Fail, Pass, ... High Distinction. Instead, Les awarded an elephant stamp to satisfactory programs and an upside-down sheep for unsatisfactory ones.[49] Forestry employed Joe Miles as a programmer, but he wasn’t involved in teaching.

John Hurst’s memory of the Nova Pascal Cafeteria system, in conversation with Chris Johnson.

John: What I do remember about Pascal P is that it was developed in about 1977. I wrote a paper on the experience using it in 1980 when I was on study leave. And one of the nice things about Pascal P system was that it actually tracked the operation of a program by source line number. So students would get information about where the error happened in terms of the source line. This was often not possible with a compiler version. We also got a lot of data on what sort of things students were doing. That allowed us to improve the interpreter because we could see where the bottlenecks were. The paper I wrote detailed how we dealt with all this data and the improvements we had. [The paper is online.[a] It describes the system in a fair amount of detail.]

Chris: Yes, I also worked on some of that Pascal P system on the Nova. I remember writing some of the interpretative routine.

John: Yes, it was recursive descent compiler. And the interpreter we wrote ourselves. It was a bootstrap system that the compiler provided… written in Pascal and we only had to provide a cross-compiler. We had to write an interpreter for the Data General Nova end but the whole content was already provided and then we boot strapped a Pascal P compiler on a Pascal P compiler. And then subsequent improvements were made directly on the Data General.

Chris: Do you remember if that was a Nova Eclipse or Supernova?

John: I’m pretty sure that was a Nova. Just a Nova.

[a] A.J. Hurst, Pascal-P, program structure and program behaviour, *Software Practice and Experience*, 10(12), 1980, pp. 1029–1036. `https://doi.org/10.1002/spe.4380101208`

[49] Kathy Handel, who taught the FORTRAN course after Les, says that an elephant meant Excellent, a sheep was Pass, and an upside-down sheep was Fail.

2.12 1976: Computer Science Becomes a Department

It was in March 1976 that Computer Science achieved Department-hood, with Ray Jarvis as Head. Bev Johnstone was appointed as Departmental Administrator, equipped with the ubiquitous IBM golfball typewriter. Unlike other staff members and students, Bev's typing was not only very fast and very accurate, but as constant as a speeded-up metronome.

Quoting from the 1976 Annual Report:

> Our WSU[a] figure for 1976 (30th April) was 90.12 which is slightly up on last year's figure of 85.46, in both cases the postgraduate component being approximately 8.[b]
>
> ...
>
> Slightly over half of our undergraduate enrolment comes from the science faculty with the remainder divided equally between Arts and Economics. Some 16% of these students are part-time and 22% female.
>
> [a] WSU = Weighted Student Units, later known as EFTSU (Effective Full Time Student Units). A student doing one standard CS unit in each semester, counted a quarter of a WSU in first year and a third in later years.
>
> [b] The department's WSU count grew to 131 in 1979, with the advent of first year teaching and an increase to 17 for the postgraduate component.

In 1977, the Board of the School of General Studies resolved that the departments of Computer Science, Pure Mathematics, and Applied Mathematics should be reassigned to the Faculty of Science from 01 January 1978. The foundation Chair in Computer Science was advertised on 22 October 1977.

I don't remember things changing dramatically in practice as a result of departmenthood, but everyone felt better.

2.13 1976: Launch of Computer Science C04

It was 1976 I think when Computer Science C04, on Operating Systems, was launched, with Peter Creasy as lecturer and me as tutor. There was an assignment designed to demonstrate the dangers of timing bugs in asynchronous programs. Students wrote a Nova assembler program including a mainline and a handler for teletype interrupts. They converted their card decks into a binary on paper tape, using the main Nova in G10, then loaded the binary onto a small Nova upstairs, which had a Teletype ASR-33 connected. When executed, the program would echo characters typed on the keyboard, but double up certain letters because of a timing issue with the code running in the mainline.

The major C04 assignment was to implement and simulate a scheduler using Simula 67 on the Univac. Many students really "got into" this assignment, and some described it as the most realistic of their whole course.

2.14 1979: Commencement of First Year Teaching

Science departments were very resistant to the idea of DCS teaching first year courses because of a fear that their own first year enrolments (and hence departmental income) would fall. In 1978, some Science lecturers asked their first year classes, "Would you have enrolled in Computer Science instead of this course if it had been available?" Reassured by the answers, Computer Science A01 and A02 were allowed to launch in 1979. A01 was taught by John Hurst and the language of the course was Pascal, supported by a student-accessible card-reader and line-printer attached to the DEC KA-10 in Copland G3A.

Early editions of A02 were taught by Vicki Peterson and Malcolm Newey, who returned to ANU in 1981. In 1983, the newly arrived Brendan McKay shared the teaching with Malcolm. After that Brendan had sole responsibility and introduced the A02SUCS simulated machine. (See the panels on Pages 59 and 151.) As Brendan says:

> I first shared A02 with Malcolm Newey, then I got it for myself. Indeed, A02SUCS was mine. Both the program (in Pascal) and the name. The students loved the name, but they didn't love the assignments. One of the first was to implement a macro facility. When I left A02, the wowsers changed the name to PeANUt (boo!) and also changed the instruction set to be less useful.

An extract from the specifications for Brendan McKay's A02SUCS system

```
                         COMPUTER SCIENCE A02

                              CSA02-88-04

             Specifications for the A02SUCS Computer (Model 2).
             --------------------------------------------------



Contents:
 0.  Introduction and Notation
 1.  Basic architecture
 2.  Machine language
 3.  Procedure conventions
 4.  Trap processing
 5.  Assembly language
 6.  Reocatable files
 7.  Image files
 8.  The linker
 9.  Utilities

---------------------------------------------------------------------------
0.  Introduction and Notation

The A02SUCS (A02 Slightly Useful Computer System) computer is an imaginary
computer.  It exists solely as a suite of programs written to run on a
Sun computer.  This document describes the basic architecture, machine
language, relocatable files, image files and assembly language.
```

2.15 1970s: Academic Staff Arriving Later in the Decade

Tony Cantoni left at the end of 1972 and Robin Stanton, Brian Molinari, and Chris Johnson arrived the year after. Rob Ewin and I started as Tutors in 1975. John Hurst was next to arrive, then Malcolm Newey and Richard Brent, both refugees from ANUCC/VCCRG.

Robin Stanton left school at the end of Year 10 (Intermediate Certificate) and completed a "pre-apprenticeship" course in electrical trades at Sydney Technical College. He told me, "I started a serious working life in '58 with a small industrial electronics company and enrolled part-time (7 yr) in an engineering degree [BE (Electrical)] at UNSW (initially NSWUT)."

In 1959, he started work for Remington Rand Univac. He wrote a program to show off the Univac SS80, officially opened by the Lord Mayor of Sydney in March, 1962.[50] According to Robin (and to Wikipedia[51]) the SS stood for "Solid State", and there were two variants: the SS80 which used 80 column cards, and the SS90 which used 90 column, round-hole cards.

Although most of the logic in the SS80 was solid state, the primary memory was a rotating magnetic drum. Robin says that the coding sheets for the SS80 were laid out in blocks in such as a way as to minimize rotational latency.

[50] http://digital.sl.nsw.gov.au/delivery/DeliveryManagerServlet?embedded=true&toolbar=false&dps_pid=IE922175&_ga=2.66674990.1265873781.1604539563-1909933466.1604539563

[51] https://en.wikipedia.org/wiki/UNIVAC_Solid_State

In the same year, Univac released the model 1004, a plug-board programmed punched card data processing system with 961 six-bit characters of core memory.

A UNIVAC 1004 plugboard. *Photo copyright Museums Victoria.* **`https://collections.museumsvictoria.com.au/items/383697`** *Used under Creative Commons licence.*

One of these machines was installed in Sydney, and Robin worked with it. He recalls, "The 1004 was an interesting machine! I wrote/plugged programs for one in Liverpool St (Remington Rand Univac) — The plugboard was bloody heavy! It took two people to install it."

The plugboard allowed for programming up to 62 steps. Some detail of how it worked is at `https://ub.fnwi.uva.nl/computermuseum/univacpatch.html` A brochure for later models of the Univac 1004 is at `https://archive.computerhistory.org/resources/text/Remington_Rand/Univac.1004.1964.102646252.pdf`.

ICT (International Computers and Tabulators, later ICL) were given the rights to market the Univac 1004 throughout the British Commonwealth because of ICT's tabulator customer base. Robin's employment was transferred to ICT. During this period he worked as an outsourced consultant to both DEC and Ferranti. He says of that time,

> Both Univac and ICT/L ran "computer centres" which offered computing services based on programs which were written within the Centres. At Univac I worked on a number of client admin systems requirements – the most prominent of which was a financial processing system for Shell (petroleum). While there was a hope that clients would invest in their own computers, most targeted clients were attracted by moving from tabulating systems to the more flexible — business process oriented — computer based systems. When I moved to ICT (in North Sydney) I ended up being responsible (Chief Programmer) for the computer centre and its operations (i.e a largish number of "self managed" programmers!).

Robin enrolled for a PhD in Electrical Engineering at UNSW in the mid-year intake in 1966 with Murray Allen and Peter Jones as supervisors. Chris Barter was in the same cohort. Robin's thesis title was *Computer Graphics – the Recovery of Descriptions in Graphical Communication*. Post graduation, he worked with Max Clowes in the Department of Experimental Psychology at the University of Sussex. He came to ANU in November, 1972.

Robin Stanton's experience of a very early computer that you're unlikely to have heard of

Dave: Did you say it was an Elliott 503[a] you programmed at Sussex?

Robin: Yes – I think somewhere I have the instruction set – the Exp Psych Dept obtained the machine from somewhere to build digital models of vision (the Dept was big on running visual discrimination experiments on octopuses – they had a lab in Italy – Naples as I recall!).

[a] `https://en.wikipedia.org/wiki/Elliott_Brothers_(computer_company)`

In 1973, Robin taught Computer Science C01 and Computer Science C03 Honours. His office usually contained an array of indoor plants, a cello, filter coffee apparatus, and piles of papers to be read, one pile for each year. He was a Fellow in the Computer Centre and had strong links to CSIRO's Division of Computing Research, in particular John O'Callaghan, Don Langridge, Barry Rosenberg, Hugh MacKenzie, and Peter Milne. He also liked sailing.

Once, as an undergraduate student, chatting to Robin, I remarked on the artificiality of computer science. Unlike physics, I said, computer science has no analogue in the real world. His response was instant:

> Nonsense. Plant a seed and the earth computes a flower!

Robin was instrumental in introducing laser printing to ANU/CSIRO (with Peter Milne), in pioneering the use of graphics workstations (from Sun Microsystems) and Ethernet. He introduced the teaching of Algol 60, Simula 67, LISP, and possibly Prolog. He was instrumental in the development of DCS over many years, in the creation of the Faculty of Engineering and Information Technology (FEIT), in the setting up of the ANU-Fujitsu CAP Project, in the establishment of the Centre for Information Science Research, and in the successful bid for the Advanced Computational Systems (ACSys) CRC.

To quote from the citation for Robin's receipt of the ANU Chancellor's Award for Outstanding Service to the Campus Community in 2013:[52]

> Professor Stanton was also Deputy Director of the Centre for Information Sciences, a Centre that spanned the IAS and Faculties at the time. With the establishment of the Faculty of Engineering and Information Technology in 1993, for which he was again instrumental, Robin became the inaugural Dean serving in that position concurrently as Executive Director of the Advanced Computing Systems CRC until 1997 when he became Acting Deputy Vice Chancellor. Professor Stanton was appointed Pro-Vice Chancellor (Academic) in 1998 and later adopted the role of PVC (E-Strategies), leading the University's strategy around information technology, e-learning and high performance computing. He also concurrently served a term as Director of the Research School of Information Sciences and Engineering from 2002 to 2003. Robin has also been key in developing and maintaining the University's relationship with Fujitsu resulting in a continuous level of collaborative research activity over 25 years.

You can read more about his distinguished career at ANU and the National Computational Infrastructure in the Pearcey Foundation Hall of Fame, at `https://pearcey.org.au/2018-hall-of-fame/`.

Among the DCS administrative staff, Robin had a reputation for feverish activity. Jeannie Haxell described him as a "whirling dervish", and Nola Whitecross says, "Of course Robin Stanton was always like a tornado, he did not believe in slow motion."

Steve Blackburn tells of giving his first lecture, while a tutor, and prior to starting his PhD. Robin unexpectedly arrived in Steve's office with a large blue stain on his shirt, due to a leaking pen.

> **Robin:** How would you like to give a lecture in my course?
>
> **Steve:** That could be good. When were you thinking?
>
> **Robin:** In five minutes time! Chemistry lecture theatre. [Now demolished.] I've got an important meeting and need to go and get a new shirt.

[52] `https://services.anu.edu.au/files/public/award-grant/professor-robin-stanton.pdf`

Brian Molinari's PhD in the Faculty of Engineering, Cambridge University

This was the classic UK research degree. The new student was given a supervisor who suggested a research topic. There was no graduate-level coursework. It was a three-year program with the objective of a PhD thesis submitted at the end of the program. The system gave you a desk and showed you where the library was, and let you get on with it. Absolutely no quality control.

The area of engineering I followed was the then-hot topic of control engineering (now called system engineering). The spectacular application area was the space program which generated a lot of interesting theoretical problems in "control theory." The Control Group in the Department of Engineering was a research group with an industrial focus. There were about 40 graduate students but only about 3 research staff. Several good researchers had recently left, to chairs at other UK universities, and hadn't been replaced. About half the students were sponsored by industry, who tended to come with an applied project. The others (like me) were assigned theoretical topics of varying interest. The industry-sponsored students generally succeeded, but the other students had a high drop-out rate. The quality of supervision was not good. I had close friends in other disciplines with the same experience.

Of course, some departments with focussed programs were excellent. In the same building was the Department of Theoretical Physics. There was a recent PhD student in a wheelchair who was reputed to be brilliant. We would see him getting into the lift. His name was Stephen Hawking.

My assigned topic didn't go very far. It involved some extensive simulation using an analogue computer that the Group had recently acquired (it must have been one of the last ever delivered). Through my own reading I discovered several problems that had some open issues (involving the Riccati equation) and I made a certain amount of progress. In the end I bundled both topics into a thesis and got it submitted. I took ten terms rather than nine, but as I was close the Department managed to find me financial support for the extra term.

I accidentally found a strategy that helped me survive. Each week I would go into the Engineering Library and take down a PhD thesis and deconstruct it. From this I worked out how to write a thesis, and calibrated what was needed. The Cambridge legend, of course, was Dirac's 21-page thesis on the relativistic theory of the electron. He got the Nobel Prize for that.

Postscript: I checked Dirac's thesis. It is 72 pages, handwritten, barely half is coherent narrative, the rest is notes and blotches. Strange – we were always told it was 21 pages. I met one mathematician who was much aggrieved his thesis was 32 pages, and hadn't beaten Dirac for brevity.

Brian Molinari came to ANU after completing a PhD at Cambridge, and then working as a lecturer at UNSW. In 1973 he taught Computer Science C02, a specialist numerical methods unit. Brian, had great influence on the computing curriculum, went on to become Head of Department, and in 2004 retired from ANU as Director of Scholarly Technology Services. Brian had a great store of humorous anecdotes and jaundiced descriptions of the hardships endured by DCS. He used to entertain departmental morning teas which were an institution and attended by all.

Brian Molinari's Experiences of 1970s University Recruiting

1970 – 1972: UNSW; Lecturer, Department of Electrical Engineering. Late in 1969 I went to a conference in Leeds (or some such) and was interviewed there by Brian Speedy of UNSW for a Lecturer job. Got an offer in the mail a few days later. It was as simple as that in those days. Worked at the Randwick campus, with a standard range of duties. Moderately successful in research following up my PhD field. Successfully supervised two PhD students. A lot of numerical computation involved.

Ran into Ray Jarvis on the stairs in the Department at UNSW one day late in 1972 (he was at a conference). I of course knew him from UWA. He was two or three years ahead of me, and he was a PhD student when I was in my last undergraduate year. He mentioned he was at ANU now, and that they had several positions in computer science going. I said I would be interested at the senior lecturer level. Canberra was much more attractive to us than Sydney, we had family and friends in Canberra but no family in Sydney. I never went to ANU for an interview, but got an offer in the mail while we were in WA on holidays over the summer. Those were the days.

As Head of Department, Brian often had to convey to members of the department, news from the University which was unfavourable or requiring of even greater efforts. He adopted the habit of signing off such emails with: `(Sigh!)`

Chris Johnson arrived from Monash and tutored Computer Science C01 from 1973 until I took over in 1976. He was also enrolled as PhD student on the subject of Runtime Structures. He wrote a technical report on the runtime structure [forked stacks] of Simula 67. After ANU, Chris spent three years at York University, working on an Ada compiler. He then moved to the Computer Science Department at the Australian Defence Force Academy (ADFA), and later rejoined DCS, teaching a whole range of different courses, and serving two terms as Head of Department. Having spent part of his youth in Birmingham, he is able to entertain Australian listeners by reverting to a Brummie accent. Chris has an abiding interest in book binding, and an array of culinary skills. When he and I were both PhD students under Robin Stanton's supervision, he used to sometimes spice up our discussion sessions with his own gourmet bakery creations. He passed on a quiche recipe which included, "approximately one bunch of spinach." Chris tells the story of his career in computer science:

> I joined DCS in 1973, fresh from an honours year in computing at Monash University. I brought a reference from my honours supervisor saying that I was good at explaining computing patiently... and learnt to do so at ANU in tutorials that quickly expanded to mini-lecture format. The task was to introduce third year students, including mature part-timers from public service administrative data processing and practical programming, to the challenges of the theoretically more complex and elegant language Algol 60, and the then radical idea of recursion. I learnt a lot from Robin and from the students, and by report so did many of the students. A success. Over the next seven years I tutored classes in first, second and third year, and started my own PhD research with Robin Stanton advising me, stimulating me to learn, and giving critical, thoughtful, challenging feedback on everything I wrote.
>
> The increasing pace of change in computing hardware and software technology carried much excitement. We rapidly went from punched cards to Teletype terminals to visual display terminals, from batch processing with punched cards carried across campus to remote batch facilities, visual display terminals in the elephant stables under the Menzies library, a handful of teletype terminals, personal desktop workstations and our first rooms of student terminals in the Copland building. Larger and faster computer hardware and their individual operating systems, IBM to Univac to DEC to Burroughs to Pyramid and on and on, seemed to mean that we spent every summer break learning a new operating system, new programming languages, or new interactive on-screen editors (character by character and then line by line, yanking data in and out of files, into editors that understood lines and paragraphs and the full screen, then interactive programming environments that understood the program itself, all with different conceptual models and command instructions), learning under pressure to use a new system in teaching later in the same year.
>
> My own research was around the design construction and use of programming languages themselves, and I learned even more languages than were needed for the teaching: Simula 67, pure Lisp, Snobol, Algol 68, Algol W, Pascal, Modula, Modula-2, EL1. I learned to be a critic of programming language design standards, alongside Malcolm Newey and John Hurst. This did not give me the required focus on my PhD however, which was not completed until 1981 as a 400 page first draft and a final shorter version a year later. The technology and science supporting the definition and implementation of programming languages — compilers, parsers, grammars, formal semantics – seemed to me to be inadequate tools, while I was searching into the smoke of why and how programs were constructed, seeking a more human outlook on the problem without realising it. This led me later to an interest in software engineering — the processes of design and construction of software at larger scale — while my PhD thesis on the transformations of data structures in programming languages prefigured the later development of design patterns and refactoring, techniques for constructing and manipulating programs that did not arrive until the wider use of more appropriate languages such as Java. I finished my teaching career with subjects including design patterns and another recent language, Python.
>
> Along the way I left ANU before my thesis was finished, going to the University of York in England for three years to do serious development of a compiler for the new language Ada (and getting in on the ground floor with the power tools of the new operating system Unix that came to dominate the academic world). I came back to Australia to teach at Monash for a year, and learned how to lecture to classes of 300 to 400 ("large lectures are a performance, not just an explanation", was the useful advice I had from my father, also an academic), then rejected larger cities and came back to Canberra to the Australian Defence Force Academy (UNSW in Canberra). I also adapted to teaching very small classes (from one semester to the next my compiler construction course had

> nearly 100 students at Monash, falling to four at ADFA). And then back to ANU as lecturer: with more variety of courses to teach, from introductory programming, the second year programming language landscape, third year computer networking, software engineering design. Compiler construction fell out of favour in the curriculum, while there were more new environments and languages: real time programming for the KRIS robot, and parallel computing with the Fujitsu CAP project.
>
> While Brian Molinari and Robin Stanton fought the battles with the university to get reasonable funding for laboratories and staffing, I lectured and researched an interesting and challenging variety of topics, and helped to run the (Australian) Computer Science Association (now CORE). Brian and Robin moved on to higher things in the university — and I volunteered to follow Brian as head of department. The immediate pressure was to recruit enough staff to meet the demands of booming numbers of students just before the dot com crash, with Brian and Robin's efforts having paid off to bring the department the income to match. One of the new recruits was Alistair Rendell, later head of the research school. I learned to be supportive and as fair to my colleagues as possible, and to shield them from the unwanted pressures and vagaries of the university's higher administration: all of the department's efforts were needed to keep the ship of computer science running.
>
> Being head of department brought me a few illusions of power, and the ability to make things happen beyond the day to day and year to year operations of a small department. I persuaded the department to volunteer to run the Australasian Computer Science Week conferences in 2000, and over several years to create a workable Software Engineering degree. By the end of my term as head the budgetary pressure was reversed, and the administration's demands for forward budget planning came with the requirement to cut costs, despite our projection of sufficient student numbers and a notional surplus in the bank — which was deemed to be needed more elsewhere. Fortunately a return to growth followed.
>
> As computing technology kept changing and yet students' challenge of learning remained, I increased my interest in the questions of how people learn computing — whether there is a sudden 'aha' moment, a radical shift in understanding, or a series of hurdles and threshold concepts, or whether it's as true as it appears that some people are unable to learn to program; and interleaved this problem with designing a curriculum with a well-structured flow of ideas and increasing complexity, and the interests and capacity of the teachers. This led me to become Associate Dean for Education in the faculty, with some research and project management in computing education. After retiring from teaching and research, this experience with designing and refining the ANU computer science and software engineering degree programs has taken me to the computing education community Australia-wide, where I am reviewing the computing curricula at many universities through the accreditation activities of the Australian Computer Society.

John Hurst was the next academic to arrive, in 1974. John was in the process of completing a University of New South Wales PhD on extensible computer architectures. He completed it in 1976.

> My PhD was on extensible machine systems and we used a home built computer at UNSW to do that work. When I moved to ANU, or course I didn't have the computer anymore. Robin Stanton suggested that I write an interpreter. Hence I wrote an interpreter for the architecture we had used and ran that on the NOVA. So a lot of the experimental work I did for the PhD was actually written for a research machine called the AR16 running on an interpreter on the NOVA.

On arrival, John found himself (reluctantly) teaching the numerical methods part of B01. He also wrote a better version of the utility to control the Sykes cassette drive on the Supernova, and a program called AJAX to remove temporary files from its hard disk. As noted above, he soon took charge of B01, and completely changed the curriculum, introducing Structured Programming and teaching with first Algol W and then Pascal, with assistance from me and Rob Ewin as tutors. In 1976, he was successful in obtaining the funding to purchase a Burroughs B1700. This was a machine whose micro-architecture could be reprogrammed to produce new instructions and new data types. This was upgraded in 1983 to a B1860.

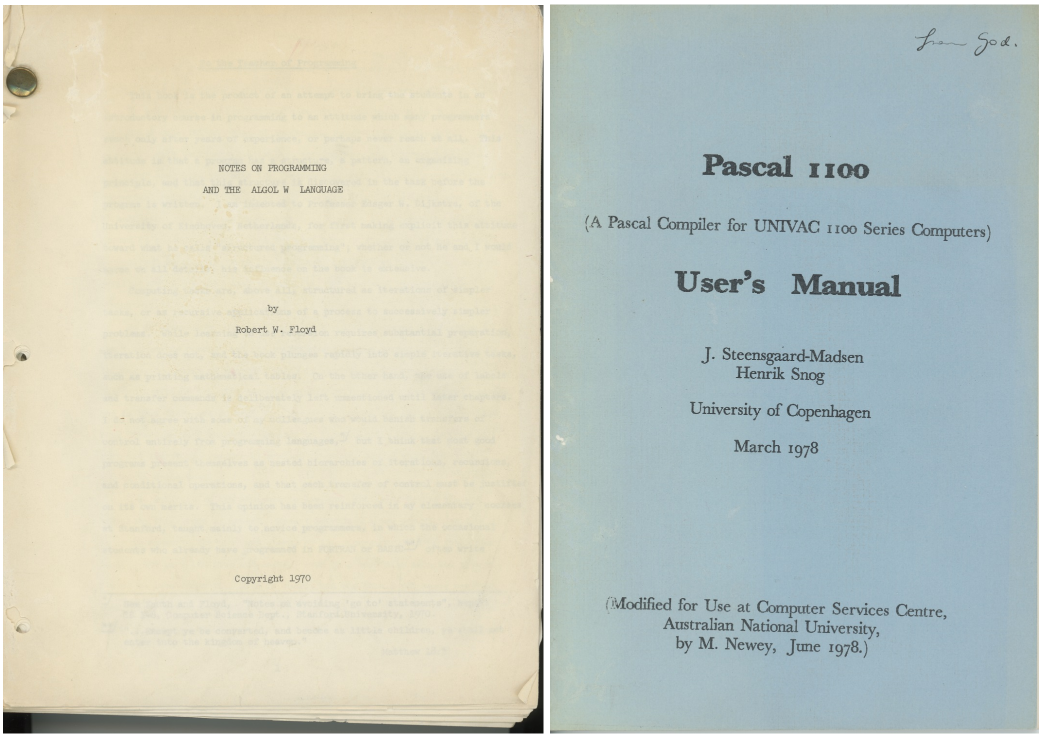

NOTES ON PROGRAMMING
AND THE ALGOL W LANGUAGE

by
Robert W. Floyd

Copyright 1970

from God.

Pascal 1100

(A Pascal Compiler for UNIVAC 1100 Series Computers)

User's Manual

J. Steensgaard-Madsen
Henrik Snog

University of Copenhagen

March 1978

(Modified for Use at Computer Services Centre,
Australian National University,
by M. Newey, June 1978.)

Manuals for IBM 360 ALGOL W and Univac 1108 Pascal.

John Hurst's account of the Burroughs B1700

We got an internal grant I believe. I don't recall there being any external funding but we also had a lot of help from Burroughs who were looking to get a foothold in ANU. Burroughs already had a B6700 installed at Canberra CAE and they wanted to advertise that they had computers on various Canberra campuses. And that was part of the deal. I recall Ray Jarvis who was Head of the Department at the time being involved in the wheeling and dealing with Burroughs about that process.

As to how much it exactly cost, it was in the order of $10-15,000, I think. Which was a fair amount of money in those days. It was very much a research machine. It was not used for any teaching purposes. We had a couple of masters students involved with it. One called Neil Justusson who basically migrated Unix to it and I wrote a micro coded interpreter both for Modula[a] and for the Unix intermediate code that we developed, that Neil developed, and we actually had a Unix system running on it, Unix light system by the mid 80's. It was incredibly slow. I remember being able to do an edit but it would take several seconds between typing it and the key stroke getting response to it. But it did actually run a full Unix light system and we had an editor and a command line interpreter.

[a] Neil's thesis is online at https://openresearch-repository.anu.edu.au/handle/1885/138340.

In the early days, quite a lot of computing equipment was purchased through the SGS Budget Review. ANU was not allowed to carry over funds from one year to the next, leading to a desire to find out how much money was left and make sure it was spent. This may or may not have been the source of funds for the Burroughs, but I remember Deane Terrell coming down the corridor and telling Ray, "There's $xxxK available for equipment purchases if you can spend it by next Tuesday". Ray always had a list!

Among the items purchased by this means were a number of Tektronix storage-screen terminals, several 4013s, a large (4096 x 3120) 4015, and a programmable 4051. All of them featured an APL[53] keyboard, vector graphics hardware, and a green screen with highly persistent phosphor. You drew your graphics and text on the screen. When finished you pressed the erase button and started again. This meant that Textronix 401x screens used for interaction tended to burn out the phosphor on the top left corner. A life extending measure was to rotate the CRT tube by 180 degrees.

At least one of the Tektronixes had a thermal printer which could preserve a monochrome version of what was on the screen. The 4051 came with a Lunar Lander game.

[53] https://en.wikipedia.org/wiki/APL_(programming_language)

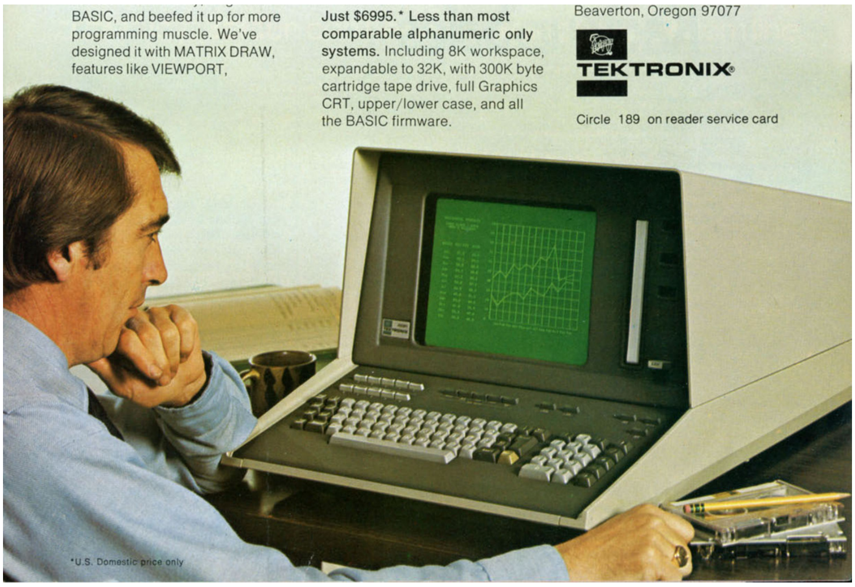

BASIC, and beefed it up for more programming muscle. We've designed it with MATRIX DRAW, features like VIEWPORT,

Just $6995.* Less than most comparable alphanumeric only systems. Including 8K workspace, expandable to 32K, with 300K byte cartridge tape drive, full Graphics CRT, upper/lower case, and all the BASIC firmware.

Beaverton, Oregon 97077

TEKTRONIX®

Circle 189 on reader service card

*U.S. Domestic price only

A Tektronix 4051 like the one which used to sit in the tiny room at the end of the 2nd floor corridor in the Crisp Bldg. I crawled a couple of hundred metres on a hot day, sandwiched between the sizzling copper roof and 30cm of fibreglass insulation to pull the cable needed to connect this room to the patchboard down in the Copland end of the corridor. *Photo public domain. from:* `https://commons.wikimedia.org/wiki/File:Tektronix_4051_ad_April_1976.jpg`

John fondly recalls meetings in the Department of Statistics: "And [Chip Heathcote] used to hold staff meetings that were characterised by half a case of red, from his own cellar. ... There were bottles of red that would come out during the meeting and people would imbibe, not too ostentatiously but enough to loosen tongues a bit and often the debate was quite relaxed and vigorous. There was talk of continuing the tradition once we became a separate department, but at that stage none of us had adequate cellars!"

John's second consuming interest was in steam trains. He was an ardent member of the Railway Historical Society and very excited about Canberra's acquisition of an enormous Bayer-Garrett locomotive.

John habitually sported a black beard. He shaved it off once and when I passed him in the corridor I failed to recognize him, though he looked vaguely familiar. He eventually left Canberra for Monash University in September 1987.

Rob Ewin did his degree at Monash University and had, I think, done a PIT course at the CCAE. He wrote a portable text editor (MiniEd?) and, with John Hurst and Chris Johnson, developed a cafeteria Pascal facility based on a Data General Nova. It was a huge advance that students were able to walk up to the card-reader at any time, load their card deck and "immediately" collect their output from the printer. Rob also implemented TEX in the department. Rob later worked on networking for the Computer Services Centre, eventually becoming Head of all the CSC programmers. After that he spent three years working for an Olympic drug testing facility in Barcelona, and later for Optus, then AARNet. Despite being very fit and very athletic, playing tennis and squash, swimming regularly, and cycling, he unfortunately died of a heart attack in 2015 in a hotel during a cycling tour in Spain. His son later completed the ride to Gibraltar on his father's bike.

While tutors, Rob and I built fibreglass kayaks and had a fine adventure on the swollen Murrumbidgee between Casuarina Sands and Uriarra Crossing.

Malcolm Newey lived for a time, as an undergraduate, in a University of Sydney hostel in which Peter Jones,[54] was a sub-warden. Peter was doing a PhD in aeronautical engineering, writing software in A9 for SILLIAC. Peter went on to work with Seymour Cray at the Control Data Corporation in Minnesota. He gave an inspiring seminar at DCS on the architecture of the CDC STAR-100.[55]

[54] `https://www.smh.com.au/national/peter-jones-pioneer-of-computing-networks-in-australia-20160919-grj7dc.html`

[55] `https://en.wikipedia.org/wiki/CDC_STAR-100`

Malcolm Newey's path to a 28-year lecturing appointment in DCS

1960-67	University of Sydney, degrees in EE, Physics & Pure Math, MSc in computing – extensible types for Algol like languages. My first program was in A9 for Sydney Uni's SILLIAC in 1962. .. Hooked! Wrote various compilers 64–67.
1967-69	'Instructor' in EE at Univ of Colorado – taught graduate courses in programming languages and their compilers.
1969-73	PhD student at Stanford Comp Sci Dept., including full time stints as research programmer at Stanford AI Project & SRI PhD thesis was Formal Semantics of Lisp in LCF[a] including proof of correctness of the Pure Lisp Interpreter
1974	Research Programmer at Edinburgh Univ – part of team, headed by Robin Milner, developing first LCF system & ML language
1975-79	Research Fellow at ANU (ANUCC, VCCRG, DCS.) Taught half of a couple of courses at DCS (while Jarvis & Stanton on study leave)
1979-81	Back to Boulder for 2 years as visiting Associate Prof teaching Programming Languages, AI, Operating Systems
1981-2009	DCS, ANU

[a] Logic of Computable Functions. https://en.wikipedia.org/wiki/Logic_for_Computable_Functions

Among Malcolm's lecturers in the Basser Department of Computer Science were Peter Poole and Chris Wallace, who went on to head departments of computer science at the University of Melbourne and Monash University respectively. Malcolm rates Peter and Chris's courses very highly. "Peter was superb on operating systems and Chris was great on Hardware. The other impressive lecturers in Basser were Earl (Buz) Hunt who has written extensively in psychology and AI, John Bennett (the prof), and Jan Hext who supervised my MSc thesis." In Electrical Engineering, Malcolm was taught by David Wong:

> David was the lecturer on computers in Elec. Eng. III and supervisor of my final year's honours project which was designing an assembly language for ARCTURUS, the computer that he and Kevin Rosolen were building in the EE Dept. I also did a cross assembler on the KDF-9 for this language. His claim to fame is that he was the principal builder of SNOCOM, assisted by Murray Allen.

SNOCOM[56] was the first semiconductor-based general-purpose computer to be built in Australia. It was built for the Snowy Mountains Hydro-Electric Authority, hence the name SNOCOM (SNOwy COMputer), and after completion in 1960 it was moved to Cooma.

During his Masters at USyd, Malcolm worked as a part-time programmer for Charles Hamblin[57] at the New South Wales University of Technology, who he rates as the most outstanding computer science person in Australia at the time. Hamblin was actually a philospher but he designed a programming language GEORGE and built a GEORGE compiler for the English Electric DEUCE. Malcolm worked on implementing PolyGEORGE, arithmetic with polynomials, in which real numbers were approximated as ratios.

GEORGE is considered to be the first language to use reverse Polish notation and its runtime system was based on a stack. Hamblin presented his stack concept and compiler at the Australian Computer Conference in June 1957. English Electric engineers who were present took the ideas back to England and implemented them in the company's next product, the English Electric KDF-9.

Also at the University of Sydney, Malcolm wrote a compiler for Minigol, a subset of ALGOL-60, which was used in teaching for "decades."

Malcolm's wife Marie was an operator on the Basser KDF-9 at the University of Sydney. The first program that Marie wrote for it (in KDF-9 autocode) ran first time, an unusual occurrence!

[56] See https://heritage.pearcey.org.au/snocom-and-cirrus/

[57] https://en.wikipedia.org/wiki/Charles_Leonard_Hamblin

1979: Mike Allen (DEC salesman, L) and Malcolm Newey inspect the hand-me-down SGS DEC-10, recently installed in Copland G3A just off the Copland Lecture Theatre foyer. *Photographer unknown.*

While Malcolm was completing an MSc at USyd, Bill Waite spent a year there on a Fulbright Fellowship after completing a PhD at Columbia University. Malcolm followed Bill to the University of Colorado at Boulder and lectured there for two years. Following that, he completed a PhD under John McCarthy and Robin Milner at Stanford, and later worked for several years with Milner in Edinburgh. Other famous people he encountered at Stanford were Robert Floyd who ran a first-year PhD course called "Problems Seminar" and Ken Colby, a disillusioned psychoanalyst who was active in artificial intelligence and wrote the chatterbot *PARRY*. Colby worked in the Menlo Park veterans hospital which inspired Ken Kesey (a former inmate) to write *One Flew Over the Cuckoo's Nest*. Colby was a pioneer of the use of computers in psychiatry. As well as PARRY, and with Malcolm's assistance, he built games for people with autism, featuring sounds, voice and keyboard interactions. One morning a week the clients would line up to use the games.

At Boulder, Malcolm worked on a Librascope L3505[58] command and control system, obtained by the University of Colorado as military surplus from the US AirForce, after other Librascopes supplied to the military failed, allegedly due to sabotage.

Malcolm and Richard Brent were at Stanford at the same time, one in artificial intelligence and the other in numerical analysis. Malcolm remembers seeing Richard and his wife Erin attending seminars together. Malcolm also attended seminars in Donald Knuth's house on the Stanford campus.[59] Knuth's house featured a very large room with a high vaulted ceiling and a huge pipe organ.[60] The floor was covered in a large mosaic depicting a map of San Francisco Bay. (For non-CS people reading this history, Knuth is a giant of computer science and digital typography. In 1968, he started writing a seven-volume book titled, *The Art of Computer Programming*. So far only three full volumes and parts of Volume 4 have appeared, but it has long been an essential reference. In 1978 Knuth released TEX, using which this book was prepared.)

Malcolm's research interests related to higher order logics and program correctness. At ANU, he first worked at the Computer Centre. When it imploded (not his fault!) he was briefly a member of the resulting VCCRG fission product and then joined DCS in 1978, retaining his research fellowship.

[58] http://ed-thelen.org/comp-hist/BRL64-l.html#LIBRASCOPE-L3055

[59] https://www.youtube.com/watch?v=_PTjOiRAS0A

[60] https://www.youtube.com/watch?v=OBO613Q8hAw

Dec 1975: Robin Stanton (L) and Brian Molinari at a departmental barbeque. ***Photos: David Hawking***

At the Dec 1975 barbeque: John Hurst (L) and Chris Johnson. ***Photos: David Hawking***

Photos from the mid-1980s, in the Copland Building. Top to bottom: Ray Jarvis, Rob Ewin, and Malcolm Newey ***Photos: Departmental archive***

Having experienced the world of printing with the Xerox Graphics Printer at Stanford AI Lab in the early 1970s, Malcolm wanted something similar for documents at ANU. Coming across Bruce Millar's setup of PDP-11 and Versatec printer/plotter he set about digitising some common fonts and writing the software to take annotated text to the desired hardcopy on the Versatec. This system was used from 1976 for a number of reports, lecture notes and papers, for example the customised *Pascal 1100 User's Manual*, the cover of which appears on Page 65.

Malcolm had a short period of study leave at Edinburgh in early 1979 and returned with a mag-tape containing Brian Reid's document formatting system. By means of pre and post processing programs he wrote, along with judicious poking of the binary, was able to provide it as a service with output to the wet-process Canon laser printer / Imagen controller. This was revolutionary – prior to this the department had relied on Diablo daisy-wheel printers. In 1979, Malcolm organised a highly influential, well-attended, do-it-yourself typesetting symposium which presented ANU authors with tools (TeX and Scribe) and devices (laser printers and typesetters at ANU and CSIRO).

Bob McKay of CSIRO Division of Computing Research, made us jealous by describing the Xerox 9700 laser printer and phototypesetter installed at DCR. The Xerox was capable of printing 120 pages per minute which, at the time, seemed impossibly fast, while the phototypesetter was capable of far higher resolution than state-of-the-art laser printers.

In late 1978, Malcolm and Marie organised a highly memorable dinner party for the department and for the four short-listed candidates for the newly created Chair of Computer Science. The candidates (Richard Brent, Ray Jarvis, Mike Osborne, and Vaughan Pratt) sat, like like-charged particles in minimal energy configuration, at the four corners of the space, and chatted with members of the department and their partners.

Malcolm taught and organised seminars in DCS long after retiring from it.

Mike Robson: An unexpected 18 years at ANU

I came to ANU quite by accident. After a few years in my first job at Lancaster I was ready to move on for various reasons and after a while, I had outstanding applications at ANU, Victoria (BC) and Utrecht. I heard that I had not got the lecturer job at ANU, so I was down to two possibilities when, out of the blue, I received a suggestion to consider a three year appointment as a lecturing fellow at ANU. After a while worrying about the wisdom of giving up a tenured position and taking a temporary one, I decided to take the chance and head off for a country I had never even imagined myself visiting. Little did I imagine that I would not leave ANU until nearly eighteen years later, with two ANU degrees, an Australian passport and many happy memories of friendships with Australian colleagues, academic and technical!

Academically, I appreciated working alongside Richard Brent and later Brendan McKay whose research interests were close to mine but even colleagues in quite different areas were always able to explain their work in an understandable and interesting way at regular seminars (normally followed by a barbeque at Black Mountain Peninsula!).

Mike Robson arrived from the University of Lancaster, UK in June 1977, bringing with him a strong English accent, to a degree that caused his students to create and circulate a "Phonetic Guide to Understanding Mike". in 1981 he successfully supplicated for a PhD by published work, *Worst-case fragmentation of dynamic storage allocation algorithms*. Among his interests were Algol 68, for which he obtained a DEC-10 compiler, tree structures, and the challenge of writing the shortest possible program to print its own source code. He achieved impressive brevity in many languages, and I think he managed to create a program which could be represented in the 80 characters of a single punched card, and which, when run, printed itself.

At his house in Mawson he had a small above-ground swimming pool, in which he swam every day when the water temperature reached 8C or more. He enjoyed a glass of wine and at one Departmental barbeque at Black Mountain Penisula managed to knock himself unconscious by diving into a metal rubbish bin after tripping on a low brick wall. When the likes of Peter Wishart and David Poole locked him into the Departmental Centre after a heavy-drinking celebration of the end of semester, he smashed down a glass door using a squash racket.

Mike left DCS in 1995 to take up a chair at the Université de Bordeaux, where he still works.

Lead-up to Richard Brent's appointment to the Chair of Computer Science

My first visit to ANU was in the 1965/66 summer vacation after my second year as an undergraduate at Monash (where CS was not yet an option, as Chris Wallace had not yet arrived). I got an ANU vacation scholarship in Astronomy and went to Mt Stromlo, where I learned how to program the IBM 1620 that was there at the time. That was really my first experience of programming (though I later did a bit at Monash in Ferranti Sirius autocode and by 1967 on a CDC 3200).

The title of my thesis at Stanford was "Algorithms for Minimization without Derivatives." You could class it as numerical analysis (NA), which was part of CS at Stanford in the early days. In fact one of my two supervisors, George Forsythe, was a numerical analyst who left the Maths Dept at Stanford to found the CS Department. Forsythe died in April 1972, so I was his last PhD student.

When I went to Stanford in 1968 I actually wanted to study AI (I had vague ideas of writing a chess-playing program). For the first year of a PhD at Stanford (like most US universities) you have to do coursework and pass qualifying exams, so I had to pass exams in AI, NA, Systems, and Programming Languages.

However, my stay at Stanford was funded by a CSIRO Overseas Postgraduate Scholarship (such things don't exist anymore), and one had to study a field that was "of interest" to CSIRO. AI was not on the list. However, NA did qualify, as the head of the CSIRO Division of Computing Research (which was the organisation on Black Mountain that ran the CSIRO mainframe and primitive network) was Godfrey Lance, a numerical analyst. So I figured that I could do AI later, and found a supervisor (Gene Golub) in NA. Forsythe got involved later, when Golub went on sabbatical (there was no internet at the time – ARPANET was just starting to be used), so I couldn't easily communicate with him while he was away).

So that's how I did a thesis in NA (finished April 1971). I always kept an interest in AI, but never really got involved in research in that area. I didn't want to do NA all my life, so I got involved in algorithms and complexity theory (with people like Joe Traub at CMU, and his student HT Kung [now William Gates Professor of CS & EE at Harvard]). From Kung I got an interest in parallel computing, which turned out useful later for the CAP project. I also had an interest in computational number theory, which I started at Stanford and IBM, and have continued.

After 12 months at Stanford, I made a quick trip back to Australia in order to marry Erin, and we both went to California after the wedding. Erin worked in various jobs while in USA – as TA for George Forsythe, and as a programmer in a startup company. However, she wasn't able to work after we moved to NY.

After my PhD, I worked at IBM Research, Yorktown Heights, NY, from April 1971 to (about) August 1972. I would have stayed longer, but my CSIRO Scholarship had the condition that I should return to Australia. I got (somewhat reluctant) permission to delay my return for 18 months while I worked at IBM Research. There was an unwritten understanding that you might be offered a job in the relevant CSIRO Division. But when I enquired about this, Godfrey Lance wasn't interested. That suited me fine, as it turned out.

In any case, I could not extend my "student" visa longer than 18 months, and on a different visa I might have been drafted and sent to the Vietnam war. (That was also possible in Australia, but I was lucky in the "birthday lottery" so was not called up.) Initially I was going to go to Monash to help Chris Wallace start up CS there, but Bob Anderssen visited me in NY and persuaded me that I would have too heavy a teaching load at Monash (probably true) so instead I got a "Research Fellow" position in the ANU Computer Centre. I had my first two PhD students there – Jim Abbott and Andreas Griewank.

At some point the service and academic sections of the Computer Centre came together in the Leonard Huxley Building but I think things worked better when both groups were separate! Eventually the Uni (or Deane Terrell, who was head of the committee that recommended it) split it up, with one half becoming Computer Services with Bob Watts as Director, and the other half becoming the "Vice-Chancellor's Computing Research Group". This was a strange name, as the VC had no interest in computing research, and it was clear that the group had no future, so we all had to find new jobs.

When an advertisement went out for the Chair in Computer Science, Mike Osborne suggested that I should apply for it. (Thanks to Mike for that, as I might not have done so if he hadn't encouraged me. After all, I was only 32 and had no undergrad teaching experience and very little administrative experience at the time!) In the end, the short list consisted of me and Vaughan Pratt. Vaughan was an Australian who had also studied at Stanford (his advisor was Don Knuth), just a year or two older than me. Somehow they decided in my favor. Vaughan went on to be one of the founders of Sun and became quite wealthy, but essentially stayed at Stanford.

Richard Brent was the successful candidate for the position of Foundation Professor of Computer Science, joining the department early in 1979. He received an "endowment" of $50,000 while rumour had it that when Vance Gledhill took up a computer science chair at the Australian Defence Force Academy (ADFA) at the same time, he was granted $850,000. Ray Jarvis continued on as Head of Department for an overlap period. As Richard says: "This seemed sensible to everyone concerned, since I had a lot to learn and was in no great hurry to become Head with all the paperwork, committees, and responsibility for courses, students, staff etc!"

Richard Brent at the wheel of a 1936 Plymouth coupe, purchased new by Richard's father. The Plymouth has a "dickie-seat". I.e. the boot lid folds partly open and two people can sit in it, only able to communicate with the driver by yelling through the open back window. *Photo: Brent family.*

Richard is quoted in several panels (on Pages 35, 72, 73, 84). A nice biography appears in the proceedings of *Number Theory Down Under, 2016. (Dedicated to Richard Brent on his 70th Birthday.* – `https://www.austms.org.au/wp-content/uploads/Gazette/2017/May17/NTDU.pdf`

Ray Jarvis held an ARGS (Australian Research Grant Scheme) grant. Combined with some of Richard's endowment, this permitted the acquisition of a VAX 11/750. Ray used it for real-time robotics experiments and Richard used it for research on dynamic storage allocation algorithms on machines with virtual memory. Richard's work was published in Australian Computer Science Communications (1981) and ACM TOPLAS (1989). Says Richard, "The long delay was because the editor forgot about the paper and I forgot to hassle her – it was only when a new editor came in that the paper was 'rediscovered'."

The VAX 11/750 was located in or near Copland G10 and connected by thickwire Ethernet to a cluster of Sun workstations set up by Robin Stanton in a second-floor room (Copland 2046?) on the other side of the quadrangle. I think that this was ANU's first Ethernet. The Suns were an example of high resolution, bit-mapped graphics workstations, running Unix. At the time so-called 3M machines (one MIPS, one Megabyte of RAM, and one Megapixel display)[61] were considered by some as ideal machines for computer science. One of the suns ran diskless, using NFS (Network File System) off the other.

Soon after Richard's arrival, (coincidentally!) I left the department to become site manager of the SGS DEC-10, relocated to the Copland Building from the Joint Schools (RSSS and RSPacS), specifically to provide a computing resource for undergraduate teaching. The DEC KA-10 was a 36-bit machine housed in a string of cabinets the size of tall refrigerators. One such cabinet, big enough to comfortably get inside, contained 16Kwords of memory, i.e. about 80Kcharacters. Three shorter cabinets, like washing machines, housed a total of 100MB of removable disk storage.

[61] `https://en.wikipedia.org/wiki/3M_computer`

Software written by Hawking and Huston, in the era of the Faculties DEC-10s

SNAP Sensible New Accounting Package, allowing usage reporting to the Management Committee.

CRUD Controller of Resource Usage on the DEC-10.

DIRT Disposer of Irritating Remnant TTY jobs.

SHARK Search for Hobos (students who'd exhausted their daily or weekly time quota) and Ruthlessly Kick Them Off.

KANGEL TOPS-10 only allowed two levels of privilege: God (whose PPN (Project, Programmer Number) was [1,2]) and mortals. Course managers could gain godly powers over their students by running the command `r KANGEL` (i.e. archangel).

SPOTTY Spooler of Output to a TTY. This allowed students to print to the serial printer (A DEC LS 120) in the terminal room. It avoided paper wastage by checking whether the originating account was still logged in, and sending a message as the output was about to appear. In conjunction I had managed to reduce the memory requirements for the standard TOPS-10 line printer spooler `lptspl`from more than 50% of the available RAM on the DEC KL-10 to about 1%.

PASTIE PAScal Terminal Interactive Environment (?). This was the only way that students in introductory computer science could be supported on terminals. PASTIE restricted students to a tightly controlled Edit-Compile-Execute cycle using a very lightweight editor (perhaps Rob Ewin's MiniEd) and a Pascal compiler which was kept memory resident.

MINITAB I can't remember the name of the system but this was a version of PASTIE for Statistics students which enabled 400 students to run the very memory-hungry Minitab package by running it as a memory-resident daemon.

TFAULT A program which allowed students to report faults in terminals whether connected to the Univac or the DEC-10. Dean Spire received the notifications and organised the repairs.

Geoff Huston took over my position as a tutor. Like me, Geoff was a graduate of ANU Computer Science. He had spent six frustrating months in the commonwealth public service and was very pleased to return to an academic environment. However, he relatively quickly tired of giving repeated tutorials. Like me, he found that the first time you gave a particular tutorial, there were usually a few things you didn't explain perfectly. The first repetition was the best because you ironed out the bugs, and after that you got bored, and lacked sparkle. After a year or so, with the departure of Chung, he joined me in the Faculties Computer Unit (FCU, formerly the SGS DEC-10). Between us we wrote a lot of software, mostly in DEC-10 Simula 67 and DEC-10 assembler, designed to maximise the number of students which could be squeezed onto what were, by today's standards, very weak machines.

We were successful to the extent that the DEC KL-10, which replaced the KA-10 in 1982 or 1983, was capable of providing quite reasonable interactive response to 90 simultaneous interactive users, even though it had only about 3MB of RAM, and a processor rated at about one MIPS (Million Instructions Per Second). The VAX 11/780 system in Computer Science at the University of Sydney supported similar numbers on even weaker hardware but gave very low priority and poor service to students once their very limited interactive time quotas were exhausted.

After a long holiday, among other things, "sailing down the Nile on a dhow" and getting married, Geoff returned to the Computer Services Centre, and worked on CSC support for VAX/VMS systems. Some time after he left on holiday, "Gone but not forgotten, Geoff" messages started appearing on DCS staff terminals, infrequently at first but becoming a flood as a course deadline approached. Geoff had modified the `SPOTTY` program so that when a student printed a listing, the system would check for logged-in DCS staff accounts and send the message to their terminal.

When I left the Faculties Computer Unit in 1984, Geoff replaced me as Head. Geoff was a pioneer and evangelist of networking. From the CSC he was seconded to the Australian Vice Chancellor's Committee (AVCC) to work on building the Australian Academic Research Network (AARnet), which marked the beginning of the Internet in Australia. After AARnet, he worked for Telstra, and is currently Chief Scientist at the Asia Pacific Network Information Centre (APNIC). He has written several books on the Internet. He features prominently in Chapter 11.

Which DCS staff member wrote the first program?

Malcolm Newey wrote a program in 1962, in A9 for Sydney University's SILLIAC computer.

Robin Stanton wrote one in March 1962, for a Remington Rand Univac SS80.

I think that counts as a draw!

2.16 1970s: Sabbatical Leave & Tenure

Up until 1978, many academics held tenured positions, and were entitled to spend one year in seven on sabbatical at another institution. Universities would typically pay fares for an academic on sabbatical, and their family, and continue to pay the academic's salary. During the second half of the 1970s, there was a concerted attack on academic conditions, led by Senator John Carrick as Education Minister and Malcolm Fraser as Prime Minister, with the aim of reducing the cost of running universities. I remember an inflammatory article appearing in a leading newspaper or magazine, which accused academics of featherbedding, and of managing to extract conditions far more generous than those obtained by militant unions like the Waterside Workers Federation. I've searched for this article, but failed to find it. However, a Bulletin article, whose lead paragraphs appear below, documents a real-terms cut of about 12% in recurrent funding for the university sector, and reports the steps being taken or considered by universities to adjust to the reductions. These include an increasing reliance on contract rather than tenured staff, and significant reductions in the level of support for study leave (sabbatical). Calling it study leave was intended to weaken the connotation of one-year-in-seven.

A working group, reporting to the Tertiary Education Commission, recommended that study leave should not be an entitlement for all staff, but should be granted to those most able to benefit. It also recommended that study leave should be restricted to a maximum of six months. However, the scope for savings was small; The working group found that the direct cost of study leave in 1975 was $1.58 million – less than $2000 per academic on leave, and less than 0.2% of the higher ed. budget.

Bulletin article documenting erosion of academic conditions. (18 July, 1978)

Ivory towers under siege[a]

By DAVID ARMSTRONG

ACADEMICS ARE the envy of other workers. They get good pay more than $31,000 a year for a professor – they have their jobs for life and they get one year off in every seven as study leave.

But, for many, the cushy conditions may soon end. Career-long tenure and the almost automatic right to study leave are now under threat.

The main spur for the attack is the squeeze on tertiary finances. Next year, funding to meet the running costs of universities and colleges will rise by only $7.6 million from (at 1977 prices) $1086.4 million to $1094 million.[b] Grants for buildings and equipment will fall by $33.6 million from $85.6 million this year to $52 million in 1979.

...

[a] https://nla.gov.au/nla.obj-1606056974/view?sectionId=nla.obj-1766748237&partId=nla.obj-1606167073#page/n39/mode/1up

[b] Inflation was 12.3% in 1977, so the "rise" was in fact a cut by $126 million in real terms.

The case for sabbatical was weakened at ANU, when an academic (whose name I think I remember) returning from sabbatical, failed to deliver the study leave report required by the university. When pressured, he wrote, "I visited Italy. I had a good time. I returned to ANU."

Despite, or perhaps because of this, several DCS academics took study leave in the following couple of years. Ray Jarvis went on sabbatical in 1978. Robin Stanton went to the University of Essex and the MIT AI Lab in 1979. John Hurst went to the University of Manchester in 1980. In the panel below John explains what that involved in the changed environment, for him and his wife Barbara, a computer science lecturer from the Canberra College of Advanced Education (CCAE).

Living it up. John Hurst's description of his sabbatical year in 1980.

I spent my sabbatical at the University of Manchester, which was big in computer systems at the time. I did some work on MUSS, the system software for the MU5, as well as looking at virtual machines, such as the Pascal-P interpreter. That prompted a paper on the work that Rob Ewin and I did on the Pascal NOVA system which was published that year.

I do recall Gordon Frank (one of the MUSS people) bursting into my room one morning waving a copy of *Software Practice and Experience*, and saying, "Hey, your paper's been published", in great excitement! Gordon, Colin Theaker and I used to spend evenings working on the MU5 (evenings, because the machine was in heavy use during the day), but since we could only take over the machine after about 6:30, Colin, Gordon and I would repair to the pub and have a few pints of Guinness to "put ourselves in the mood!!!" before starting work. I have to admit that I was not a patch on their ability to down several pints and still be lucid enough to work through the evening!

While there I made contact with Ian Somerville from St Andrews, who was doing work on micro coding. I think it was by Royal Mail, back in the days when letters took less than a day to reach a destination in the UK. Ian replied, and sent me a few of his papers on work he had done on the B1700.

I went for 12 months, but ANU only paid my salary for 6. The other 6 was leave without pay. Barb negotiated the same package from CCAE, but the other way round, so the plan was that we would have an income all year. Her deal was conditional upon finding "suitable employment" when she got to Manchester. Unfortunately, she did not have a work permit, and the Immigration Officer at Heathrow gave us the third degree about our income support.

Barb did find the "employment", but couldn't be paid for it. We got round the problem by negotiating that MU would pay for Nathan's childcare – it was enough to satisfy CCAE, at least.

It was tight, especially as the dollar was sitting at 33p for most of the year. Fortunately, we had a forgiving landlady who extended credit to us when the money transfers came late.

2.17 1970s: Computer Science General Staff

The first member of general staff was Maxine Esau, whose role in supporting student computing has previously been mentioned. I think she was responsible for purchasing a Departmental Bicycle which played a role in delivering card decks from Copland to the Cockroft building. Maxine was later replaced by Dean Spire, an American formerly employed by Sperry Univac. He wore cowboy boots and had a standard response to intermittent faults in terminals and other computing equipment, "Ya cain't fix it if it ain't broke."

Bev Johnstone was the first Departmental Administrator once DCS attained departmental status. She was competent, efficient, energetic, and loved by everyone. She made a significant contribution to the development of the department during her tenure from 1976 – 198?

With the installation of the SGS DEC-10 in 1979, DCS and Statistics shared the cost of employing a programmer (Y.C. Chung) to support the operation. He was located in Crisp LG3 with me. My position was funded by the Computer Services Centre, and I was required to go over to the Huxley building once or twice a week to attend low-value meetings.

Another technical appointee toward the end of the 1970s was Anthony Glenn, who was employed half-and-half between DCS and Statistics.

Outside DCS, Peter Scardoni succeeded Harvey Jones (before him Val Brett) as SGS Business Manager at some point. He was a master wheeler and dealer, and maintainer of hollow logs. However, from my vantage point it seemed that Peter's efforts were aimed at achieving good outcomes for the areas for which he was responsible.

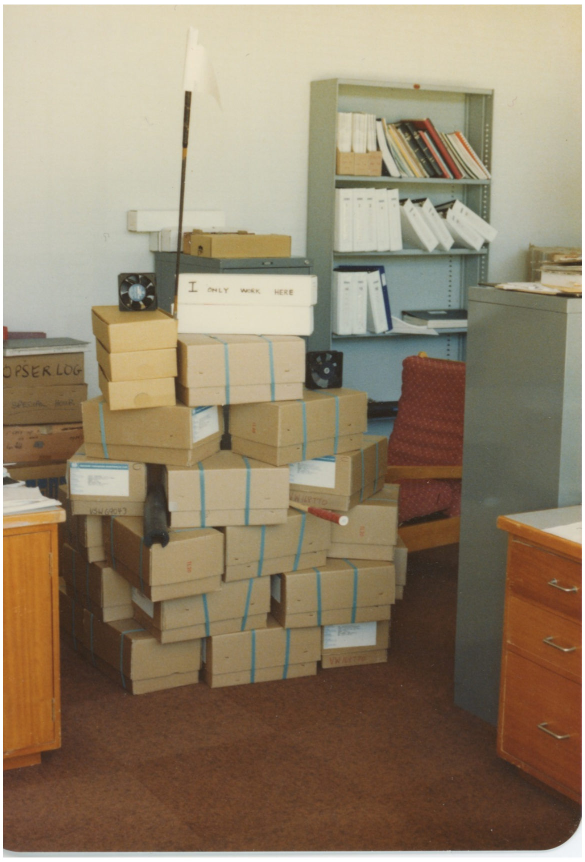

Chung and I shared an office in Crisp LG3. I like to think the barricade of line printer paper was to shield him from my constant stream of visitors, rather than from me! The style of furniture you see was typical of the ANU over several decades. *Photo: David Hawking*

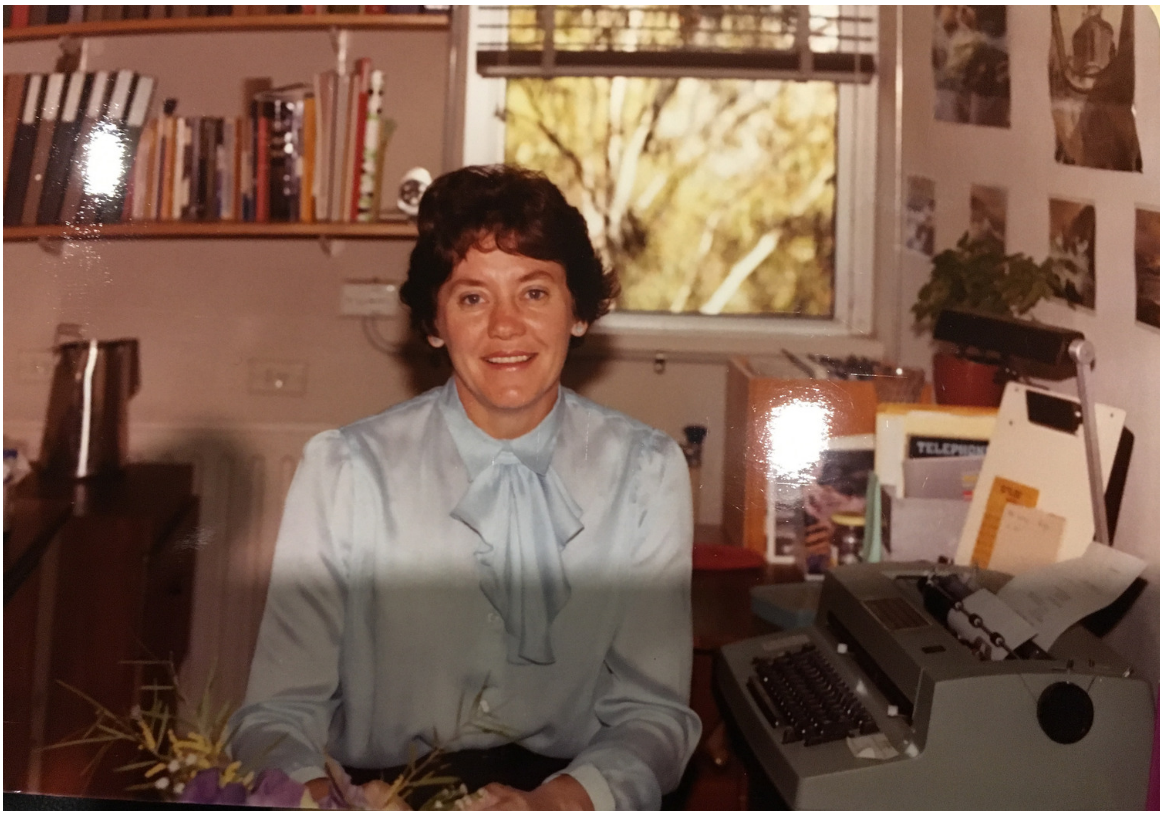

Bev Johnstone, Departmental Administrator. Note the IBM golfball typewriter, and the view of trees in the Copland quadrangle. *Photo: Departmental archive*

2.18 1970s: Esprit de Corps

In the 1970s, Computer Scientists were a very social lot. Ends of semester were marked by lunches in town or by barbeques on the lake or up Black Mountain. Bev Johnstone, Maxine Esau, and Rob Ewin were enthusiastic organisers of spontaneous barbeques. Just before 12 noon, one of them would say, "Lovely day for a barbeque". There would then ensue a dash around offices to find out who was "in", and a quick trip to the shops for bread and meat. By 12.30 the meat would be sizzling on one of the gas barbeques on Black Mountain Peninsula, or on Black Mountain itself.

Robin Stanton recently told me that Ray Jarvis also organised occasional golf days. I wasn't invited, possibly because I'd told Ray about my best score of 72 off the stick.[62]

We were also regular attendees at morning tea (10.30) and afternoon tea (15.30) in the large Arts Common Room, located on the ground floor of the Crisp Building. Tea, coffee and biscuits were served by the "tea ladies" from a central kitchen for very insignificant amounts of money. John Mulvaney (eminent pre-historian) was Chair of the Common Room Committee and, when inflation hit the finances, promulgated a sign lamenting that, "The age of the Iced VoVo is dead."

The esprit de corps of the department continued to be very strong, really until the move into the CS&IT Building in early 1995. The panel on Page 91 records Peter Bailey's view of it.

Even in CS&IT, Nola Whitecross remembers a positive atmosphere: "Having worked in ANU Departments which were not very collegiate, CS was, as I recall, a breath of fresh air, and tolerant of my left leaning political views."

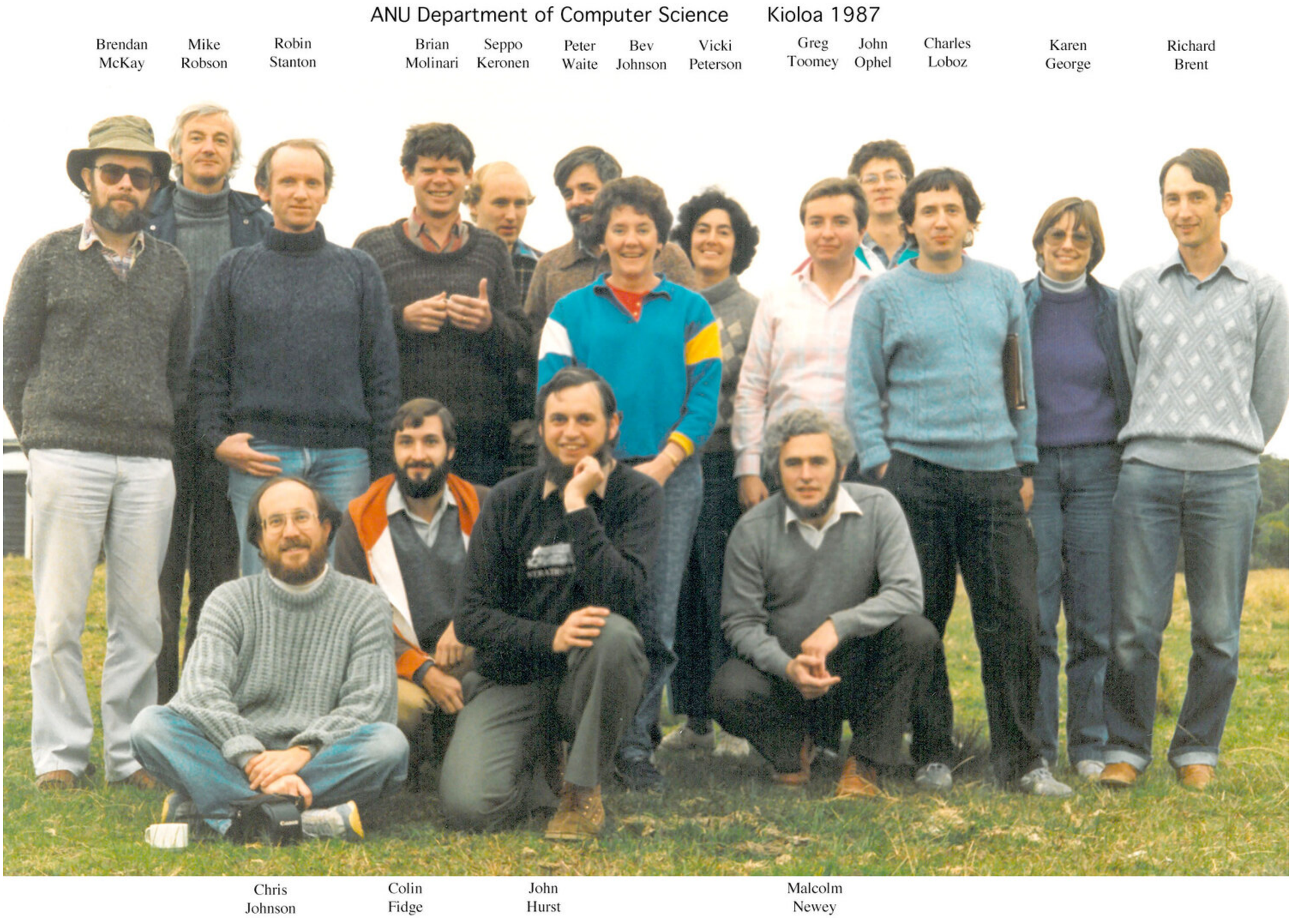

Departmental Advance: Kioloa, 1987 ***Photo: Brendan McKay***

[62] For nine holes.

2.19 1970s: Nars and Conferences

In 1977, a departmental seminar by a Masters student (Bill Ginn, I think) started with the student anouncing the topic of his talk. There immediately ensued a vigorous debate between the academics about the meaning and validity of the topic, which used up the entire duration of the seminar, to which the speaker contributed just one sentence!

In 1978, using powers not vested in me, I convened a series of Departmental Nars (I didn't want to do things by halves) and maintained a talk-first-questions-later format. At exactly 12.30, staff would walk to the Student Union to buy lunch, then return to the Departmental Centre where the talk would start despite the sound of munching. In those days, our standards were so low that we drank instant coffee using boiling water from an urn in the Departmental Centre. I have a terrible memory of "coffee whitener" but it seems too terrible to be true. The first nar in the 1978 series was given by Mike Robson, on *The height of binary search trees*.

In 1974, Robin Stanton, John O'Callaghan, and Steve Kaneff organised a workshop on Artificial Intelligence. Chris Barter from the University of Adelaide attended and may have been an organiser. Keynote speakers included Edward Feigenbaum from Stanford and Rangaswamy Narasimhan from IIT Bombay. Marvin Minsky from MIT had accepted an invitation to attend but didn't make it.

In March 1977, Malcolm Newey, Robin Stanton, and Garth Wolfendale organised a workshop on Programming Language Systems.[63] The keynote speaker was Edsger W. Dijkstra and contributors included Lloyd Allison, David Wilde, Chris Barter, Robert Elz, Peter Poole, Jan Hext, Barbara Kidman, Cliff Mason, Peter Milne, Malcolm Newey, Paul Pritchard, Ken Robinson, Rodney Topor, and Garth Wolfendale.

The comedic highlight of the workshop occurred after a talk *not* given by Robin Stanton. The instant questions were invited, Col Jarvis leapt out of his seat:

> "I have a question! ..."
>
> (Turns sideways)
>
> "A question for Robin Stanton – Robin, what did you think of that load of rubbish?"

Col Jarvis spent some time in Canberra as an academic in the Computer Centre. Prior to that he had been recruited in the very early days by the University of Western Australia, and according to an oral history from Dennis Moore[64] was UWA's first computing PhD graduate. He worked with a PDP-6,[65] ancestor of the DEC-10, and claimed by the same oral history to be the first commercially available timesharing computer. At ANU, he popularised several colourful expressions including "snow job", "kludge" (pronounced "clooj", and meaning "unprincipled coding to get something working") and "tarmac professor" (an itinerant academic who gives seminars at one institution before quickly flying off to the next).

Col became a Senior Lecturer at UWA before officially applying for a demotion to Lecturer, on the grounds that he didn't want the responsibility of a Senior Lecturer position. This caused consternation for Jeff Rohl who was head of the Computer Science Department at UWA and, in an odd connection to ANU, Julie Bakalor's brother.

Col was a skilled artist and entered the Tooheys Paint-a-Pub competition, migh even have won. One of his large oil paintings, which normally hung in Robin Stanton's house, filled up a wall in ours when Robin went on sabbatical.

The conference dinner was a bush meal with damper, at the Gundaroo Pub. As part of the "entertainment," the bus was "held-up" by a "bushranger." Dijkstra didn't enjoy this at all.

Dijkstra was a very interesting person. He allegedly prefered to fly on Fokker Friendship aircraft, knowing that they were safe because he had done the stability calculations for the wings. He argued

[63] The proceedings are online at `https://openresearch-repository.anu.edu.au/bitstream/1885/114711/2/b11749143.pdf`

[64] `https://oralhistories.arts.uwa.edu.au/items/show/1`

[65] The same one used by Ray Jarvis. As far as I know Ray and Col were not related.

against teaching programming to undergraduates, asserting that they should master mathematics first. Peter Milne went to considerable lengths to obtain a record player and suitable records to play in his room in University House. Dijkstra's entertaining trip report is worth a read.[66]

Although my PhD topic was on Computer Animation, my supervisor Robin Stanton insisted that I should make an appointment to talk to Dijkstra. I met him in the tea room at the Leonard Huxley building and learned a couple of good Belgian jokes. To Dijkstra's rising annoyance, a group of Americans at a table on the far side of the room were talking very loudly. He sashayed over to them and said, "Is it necessary that the volume with which you speak be proportional to the size of your country?" I feared violence, but they angrily left.

Dijkstra, along with Tony Hoare, and David Gries, came to Australia again in summer 1983 for "The First Wollongong Summer School on the Science of Programming" and also visited ANU, where Tony Hoare gave a lecture. Dijkstra's very entertaining trip report is online at `https://www.cs.utexas.edu/users/EWD/transcriptions/EWD08xx/EWD847.html`. It contains gems such as,

> I went to Thredbo to climb the Kosciusko, which was precisely the sad mistake I knew it would be: even the first 600m in the chair lift I disliked. At the end of the day I was so tired that next morning I slept soundly through most of Tony's lecture (but every time I woke up, it sounded so nice that afterwards I felt I could compliment him).

Chris Barter tells me that the Programming Language Systems Workshop is now regarded as the first Australian Computer Science Conference.[67] He says that the sequence numbers are controversial, but there is a middle ground that says the ordering is:

ACSC-0	ANU	1977
ACSC-1	UNSW	1978
ACSC-2	Tasmania	1979

> The conferences were called the Australian (later Australasian) Computer Science Conferences ACSC until 1997, when other related interest conferences (such as database research) were amalgamated under the Computer Science Association (CSA), and the conferences were renamed ACSW (ACS Week) at the persuasion of Jan Hext who hosted the 1997 conference at Macquarie.

The proceedings of the 1980 ACSC, edited by Ray Jarvis and Richard Brent, clearly show "Third" on the cover.

A delegation from ANU drove up to UNSW for ACSC-1. There was a conference dinner on the first evening, but the next evening no-one was up for socialising with ANU people. We loaded up an airways bag (remember them?) with cans of beer and took the ferry to Manly, later walking to the 24 hour Pancake Parlour (Lovely!) at Circular Quay. Mike Robson remembers that this was the night when Ray Jarvis suggested that the possibility of obtaining a PhD by published work.

On Friday, Ray Jarvis, Brian Molinari and I left after lunch to dodge the Sydney traffic, in a Holden Kingswood that Brian had on loan from another academic who was on sabbatical. The route we chose turned out to be sub-optimal and the cooling system of the Kingswood wasn't up to the crawling and idling required. Counting time off for a boiled radiator, it took us four hours to reach the outskirts of Sydney! Fortunately a spectacularly good meal at a restaurant in Berrima provided compensation.

By the end of the decade, DCS had taken over from the Computer Centre, the ANU end of organising joint ANU-CSIRO computing seminars.

2.19.1 CORE and CSA

Although occurring after the 1970s, a few notes on the formation of the Computer Science Association (CSA), which was eventually renamed CORE (Computing Research and Education Association of Australia), seem to belong here better than anywhere else.

[66] `https://www.cs.utexas.edu/users/EWD/transcriptions/EWD06xx/EWD613.html`

[67] Malcolm Newey agrees.

Gopal Gupta, who has held senior computer science positions at Monash, James Cook, and Bond Universities, recalls:

> First thing I remember is a meeting of a bunch of CS Professors in Chris Wallace's office. I think it included John Bennett, Peter Poole, Arthur Sale, and a few more. It was either early 1985 or early 1986. I assume there was an ACSC in Melbourne that year. Chris Wallace asked me to present the idea of establishing CSA which I did. There were some questions but most people were supportive. The idea was then presented to the meeting of Profs + Heads at the ACSC and was approved. By this time Chris Barter and John Rosenberg also got involved. I assume Chris Barter was elected President. Rao Kotagiri suggested we needed an election for the President. Rao was PC Chair for a number of years. He passed the job on to me. Later I was elected CSA Chair.

Chris Barter, who was at the time Professor of Computer Science at the University of Adelaide, passed on copies of the 1987 CSA Constitution, drafted by Jan Hext and Arthur Sale and put to the 1988 meeting, and also minutes of the 1997 CSA AGM. He tells me:

> On the formation of the Computer Science Association, I have emails in 1987 involving a steering committee to prepare for a Special/Inaugural General Meeting at the 1988 ACSC. The SC members were: Chris Barter, U of Adelaide; Jan Hext, Macquarie; John Rosenberg, Monash; Ron Sacks-Davis, RMIT; Jennie Seberry, Sydney U; Robin Stanton, ANU.
>
> They (mostly?) were members of a "Ginger Group" which advised the "Professors' Association" that the annual conference ACSC needed to be managed differently. It included Gopal Gupta, (and probably David Abramson).

John Hurst, formerly of DCS, was secretary of CSA for several years. Richard Brent was secretary of a predecessor organisation, "Association of Australian Professors of Computer Science", between 1981 and 1985.

2.20 The Rest of This Book

Having covered the formative years of DCS and introduced you to its founding characters, I'm going to organise the rest of the book by theme, rather than in chronological order. The first three themes are from a DCS perspective. The first of these chronicles the battle waged by DCS over decades to achieve: sufficient staff; sufficient accommodation for staff, research projects and students; and sufficient computing resources for research and for students. The second records some prominent academic members of the department from 1980 onward, and the third outlines the evolution of the CS curriculum over the years.

A second group of themes covers the evolution at ANU over the decades of:

- Computer science research,
- Academic computing resources,
- Supercomputing,
- Student computing faciliies,
- Administrative computing,
- Networks

The penultimate chapter recounts the story of women in computing and the book concludes by considering external views of the contributions of ANU computing.

Chapter 3

An Ongoing DCS Battle for Resources

ANU teaching of computer science began during the conjunction of two major upheavals. The first was the revolution in access to tertiary education. A far higher percentage of the population were attending university, and many students were the first in their family to do so. As a minister with responsibility for Education, John Gorton had significantly increased both student numbers, and the number of Commonwealth Scholarships to support them. I was awarded one of those and was very grateful. In 1974, the Whitlam government abolished university fees.

The second upheaval was, of course, the IT revolution. Tabulating machines were being replaced by general purpose computers, and increasing numbers of computers were being purchased, or leased, for new applications. Programmers and operators were needed in large numbers and they needed to be trained. Initially this training took place in technical colleges, institutes of technology or colleges of education rather than universities but many students were keen to gain university degrees rather than certificates or diplomas.

University enrolments continued to grow faster than the funds which the government was prepared to spend. Universities found it difficult to accommodate the needs of new disciplines like computer science, because they were not provided with extra funds, and because established departments strongly resisted having their resources reduced to accommodate newcomers.

In the Canberra Tmes, columnist Ian Warden and cartoonist Geoff Pryor lampooned the resistance to change of the "Faculty of Inconsequential Studies." I remember one cartoon which depicted bearded and gowned academics from this faculty looking with horror at the monstrosity of a new computer science building. Beneath their gothic revival lead-light window, the fuse on a pile of explosives burned steadily.

In this generally resource-constrained context it is not surprising that an ongoing theme in DCS annual reports, starting in the 1970s, was the lack of resources:

- Insufficient teaching staf;f
- Lack of offices for staff;
- Lack of appropriate computing facilities for students; and
- Lack of space and funding for research.

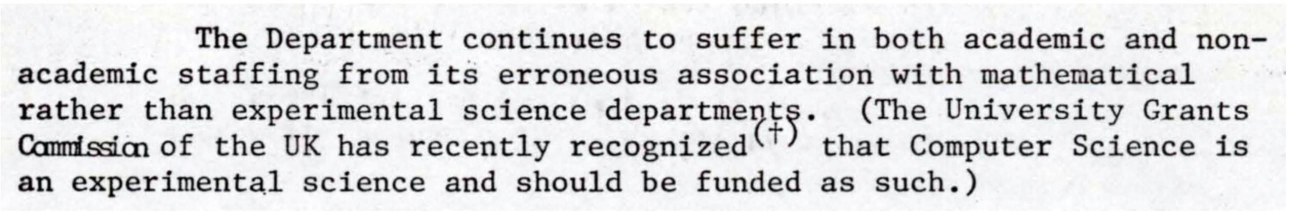

The Department continues to suffer in both academic and non-academic staffing from its erroneous association with mathematical rather than experimental science departments. (The University Grants Commission of the UK has recently recognized(†) that Computer Science is an experimental science and should be funded as such.)

Extract from the 1979 DCS Annual Report, expressing concern about the basis on which the department was funded.

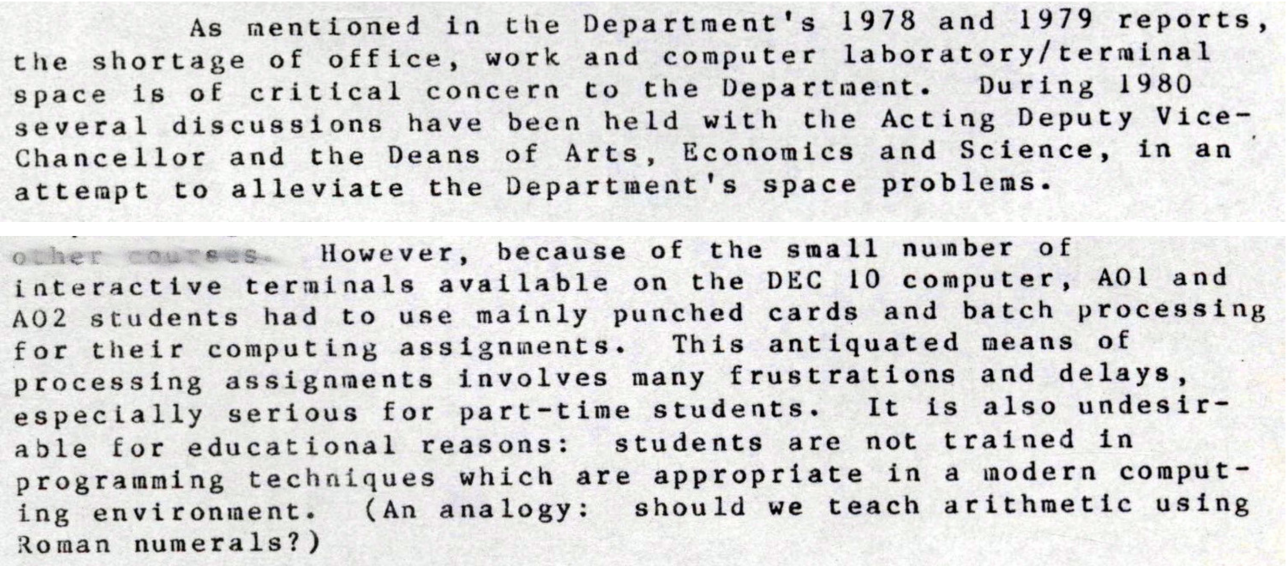

As mentioned in the Department's 1978 and 1979 reports, the shortage of office, work and computer laboratory/terminal space is of critical concern to the Department. During 1980 several discussions have been held with the Acting Deputy Vice-Chancellor and the Deans of Arts, Economics and Science, in an attempt to alleviate the Department's space problems.

other courses. However, because of the small number of interactive terminals available on the DEC 10 computer, A01 and A02 students had to use mainly punched cards and batch processing for their computing assignments. This antiquated means of processing assignments involves many frustrations and delays, especially serious for part-time students. It is also undesirable for educational reasons: students are not trained in programming techniques which are appropriate in a modern computing environment. (An analogy: should we teach arithmetic using Roman numerals?)

Extracts from the 1980 DCS Annual Report, expressing concerns about space for academics and about the suitability of facilities for teaching.

The 1981 DCS Annual Report shows that these pleas didn't go entirely unheeded. It reports that A01 and A02 students were using interactive terminals on the Faculties DEC-10 computer and that card punches were dispensed with. This was by no means due to a deluge of funding, but rather due to the purchase of hand-me-down memory from RSPhysS, to progress in reducing the memory requirements of the operating system, and to the use of locally developed software, such as PASTIE, to rigidly constrain student activities on the terminals.

Teaching of university computer science in the early days, took a lot more work than teaching established subjects. In subjects like physics and chemistry, there was a wide choice of suitable textbooks and a large set of tried-and-true assignments, laboratory sessions, and quizzes. By contrast, computer science assignments needed to be devised and debugged and staff needed to judge how long students would take to complete them, if indeed they could complete them at all. Robin Stanton says that he ran weekend tutorials to help students who feared that computer programming was beyond them.

For the Head of Department, trying to equalise teaching loads across academics, it was necessary to assess the amount of effort required to teach a unit. That in turn depended upon whether the unit had been taught before. I.e. was there an established curriculum? Were there pre-existing handouts and assignments?

3.1 The Relative Funding Model

Over the decades, government funding of Australian universities has attempted to take into account the greater per-student cost of teaching some subjects compared with others. When the Unified National System (see Page 98) began in 1989, a sample of universities provided data on the actual costs of running programs in various disciplines. This data was used to determine a set of funding weights for undergraduate subjects, with a full-time humanities enrolment weighted at 1.0, mathematics at 1.3 and engineering at 2.2. Computer science was at the same level as mathematics. This was known as the Relative Funding Model.

As noted by Richard Brent in the panel overleaf, these calculations went into determining the overall budget for the university but not how the funding was distributed internally, among departments.

To improve its lot, the computer science community needed to successfully argue with the government for a higher discipline weighting, and the ANU department needed to successfully argue with the university to have the full amount (or even more!) passed on to it.

1979–1983: Richard Brent's Recollections of Battles for Resources

Yes, there were some battles. One problem was that computer science was very new. ANU was not the first uni in Australia to start teaching undergraduate CS, but it was one of the first.

There was pressure to expand undergraduate teaching because the students wanted it – CS was in favour at the time, and there were plenty of jobs available. The academic staff were also (in most cases) keen to expand, teach more courses and have first-year courses. However, we didn't have the staff to do it properly. The student/staff ratio was much higher than in other subjects like economics, maths, physics etc.

Also, at that time CS was funded by the Tertiary Education Commission (TEC) at a level below engineering and science subjects. (Maths was at the same level, and was initially in the Faculty of Arts.) At some point both CS and Maths moved to the Science Faculty. This helped Science because it reduced the average cost per student in the faculty, but it didn't directly help us much.

Deane Terrell was an important character. He was chairman of the University Computing Committee (which had responsibility for the Computer Centre, then CSC and VCCRG). He also became the Dean of the Faculty of Economics and Commerce, where the CS group was located. (Later, 1994–2001, he became the VC.) I don't recall all the details, but we had various battles over staffing in the CS Dept. I think that at one point TEC decided to fund CS at an intermediate level (above maths but below physics) in recognition of the fact that CS students needed labs (with equipment) unlike purely paper-and-pencil subjects. (Of course maths tried to claim that they too needed computing equipment, which was partly true, but not to the same extent as CS.) However, it wasn't compulsory for the universities to pass this extra money on to their CS departments. Some did and some didn't. Deane Terrell was certainly in the camp of those who didn't want to. In those days the School of General Studies (later called The Faculties) was run (in theory, anyway) by the Board of of SGS (there was a separate Board of the IAS, or BIAS), and I remember that Deane and I had some public clashes at the BSGS meetings. I more or less threatened to resign if things didn't improve for CS.

I didn't have the right personality for such political battles. Around 1982 Neil Trudinger and others (including me) managed to persuade the Australian Research Council (ARC) to fund a "Centre for Mathematical Analysis" (CMA). Since I wanted a break from being head of CS, but didn't really want to resign from ANU, I took the opportunity to go on secondment to the CMA, from July 1983 to August 1985. I then became Head of the Computer Sciences Laboratory in RSPhysS.

3.2 1983 Bid and Petitions for Increased DCS Staffing

On 12 May, 1983, Richard Brent wrote to the Chairman of the Resources Committee, bidding for:

- One academic post at the level of tutor / senior tutor;
- Three academic posts at the level of lecturer / senior lecturer / reader;
- Two nonacademic posts.

Later in 1983 a *Petition to the Board of the Faculties from the Department of Computer Science* was prepared and submitted. My understanding is that John Hurst submitted a more strongly worded petition at the same time, but I haven't found a copy.

A subsequent report of the subcommittee of the Board which dealt with staffing formulae provides an insight into why Computer Science struggled to increase its complement of technical staff. In its report on General Staffing in The Faculties by Sub-Committee on the Staffing Formulae (ANU paper 222B/1984) this subcommittee, chaired by Deane Terrell, says:

> The Sub-Committee has identified two areas, physics and chemistry, which it urges should be given some degree of special consideration. They are the least viable in size when compared with ... other Australian universities. The larger universities have four to five times the student load in physics and chemistry. ... The Sub-Committee believes that physics and chemistry are fundamental to the Faculty of Science and that both are close to being non-viable. ... the Sub-Committee recommends that physics and chemistry ... should make their case to the Resources Committee for support staff in excess of their share of the 123 posts available to experimental science departments.

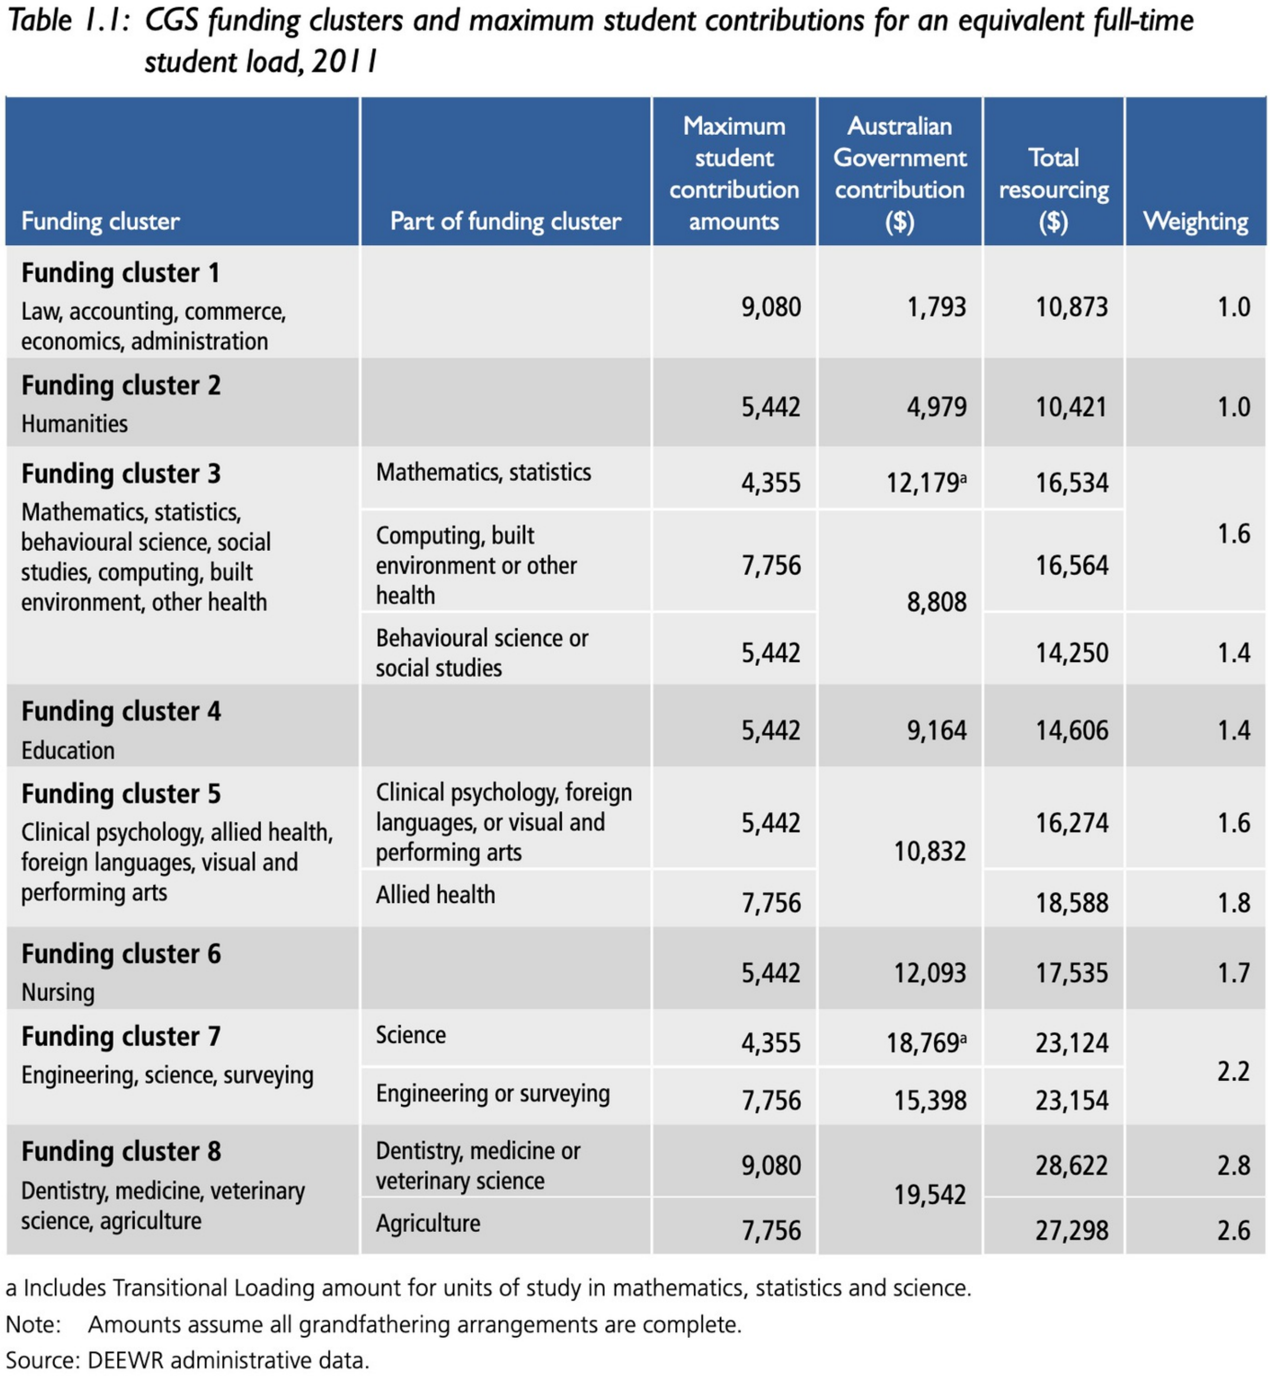

Table 1.1: CGS funding clusters and maximum student contributions for an equivalent full-time student load, 2011

Funding cluster	Part of funding cluster	Maximum student contribution amounts	Australian Government contribution ($)	Total resourcing ($)	Weighting
Funding cluster 1 Law, accounting, commerce, economics, administration		9,080	1,793	10,873	1.0
Funding cluster 2 Humanities		5,442	4,979	10,421	1.0
Funding cluster 3 Mathematics, statistics, behavioural science, social studies, computing, built environment, other health	Mathematics, statistics	4,355	12,179[a]	16,534	1.6
	Computing, built environment or other health	7,756	8,808	16,564	
	Behavioural science or social studies	5,442		14,250	1.4
Funding cluster 4 Education		5,442	9,164	14,606	1.4
Funding cluster 5 Clinical psychology, allied health, foreign languages, visual and performing arts	Clinical psychology, foreign languages, or visual and performing arts	5,442	10,832	16,274	1.6
	Allied health	7,756		18,588	1.8
Funding cluster 6 Nursing		5,442	12,093	17,535	1.7
Funding cluster 7 Engineering, science, surveying	Science	4,355	18,769[a]	23,124	2.2
	Engineering or surveying	7,756	15,398	23,154	
Funding cluster 8 Dentistry, medicine, veterinary science, agriculture	Dentistry, medicine or veterinary science	9,080	19,542	28,622	2.8
	Agriculture	7,756		27,298	2.6

a Includes Transitional Loading amount for units of study in mathematics, statistics and science.

Note: Amounts assume all grandfathering arrangements are complete.

Source: DEEWR administrative data.

Australia's Relative Funding Model, as it operated in 2011. Note that computer science, along with mathematics and statistics, have a weight of 1.6. Taken from the final report of the Higher Education Base Funding Review, October 2011. *Screenshot from:* **`https://www.canberra.edu.au/research/archived/edinstitute/documents/HigherEd_FundingReviewReport1.pdf`**

First page of the 1983 DCS Petition to the Board of The Faculties

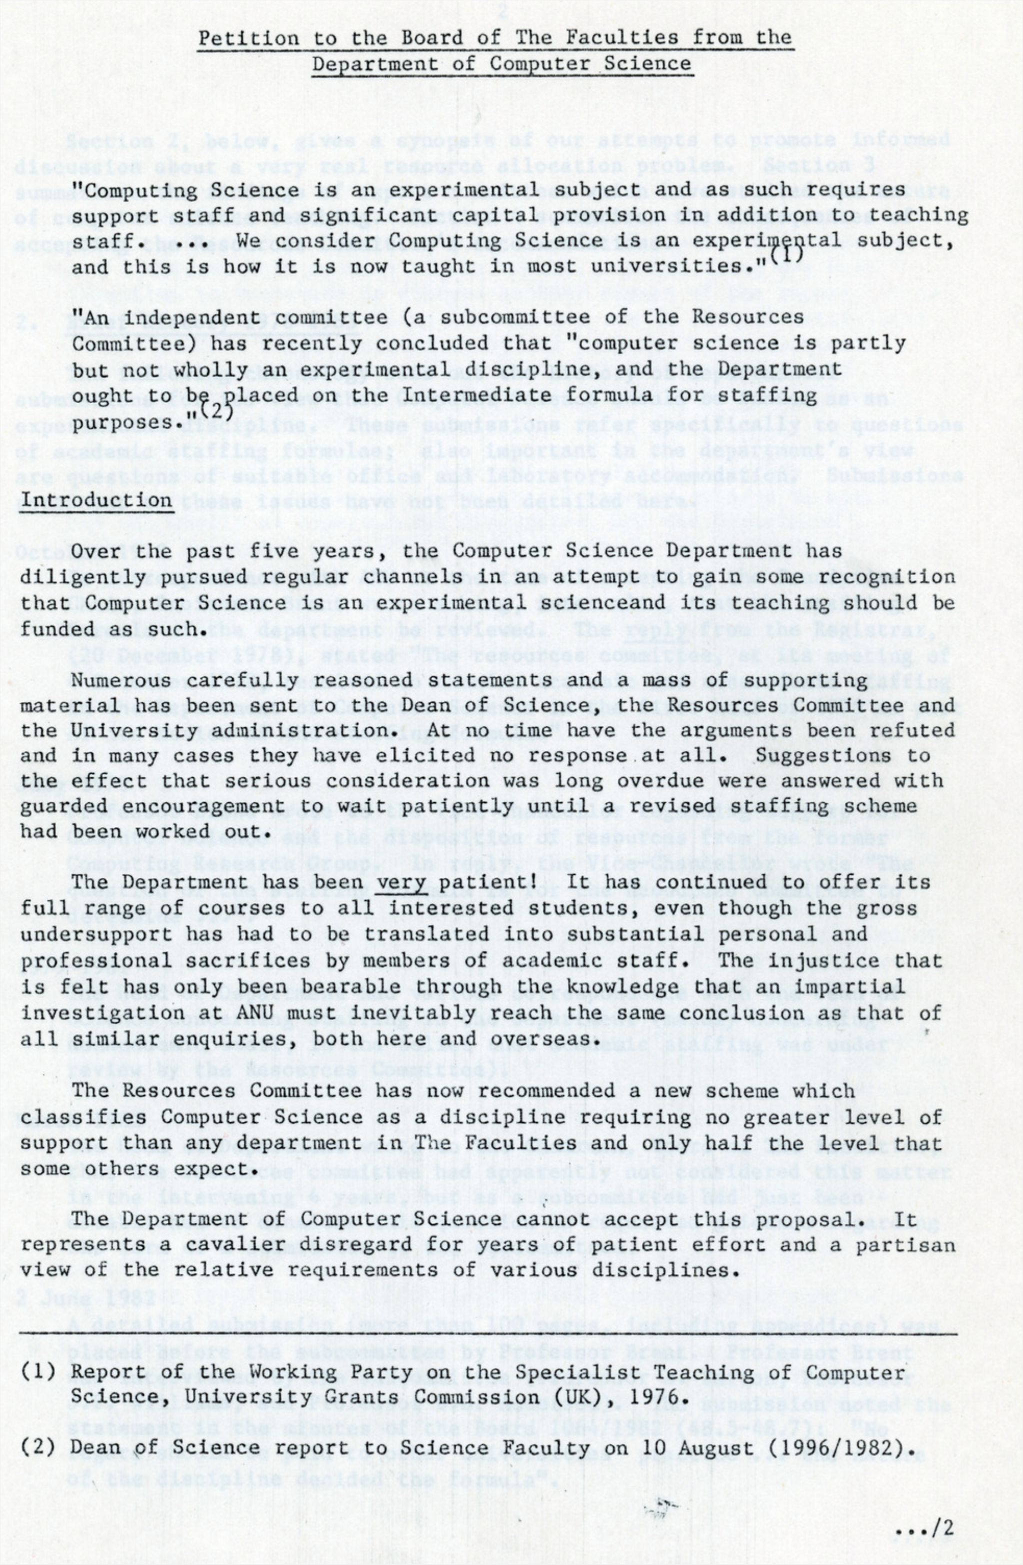

Petition to the Board of The Faculties from the Department of Computer Science

"Computing Science is an experimental subject and as such requires support staff and significant capital provision in addition to teaching staff. ... We consider Computing Science is an experimental subject, and this is how it is now taught in most universities."(1)

"An independent committee (a subcommittee of the Resources Committee) has recently concluded that "computer science is partly but not wholly an experimental discipline, and the Department ought to be placed on the Intermediate formula for staffing purposes."(2)

Introduction

Over the past five years, the Computer Science Department has diligently pursued regular channels in an attempt to gain some recognition that Computer Science is an experimental scienceand its teaching should be funded as such.

Numerous carefully reasoned statements and a mass of supporting material has been sent to the Dean of Science, the Resources Committee and the University administration. At no time have the arguments been refuted and in many cases they have elicited no response at all. Suggestions to the effect that serious consideration was long overdue were answered with guarded encouragement to wait patiently until a revised staffing scheme had been worked out.

The Department has been very patient! It has continued to offer its full range of courses to all interested students, even though the gross undersupport has had to be translated into substantial personal and professional sacrifices by members of academic staff. The injustice that is felt has only been bearable through the knowledge that an impartial investigation at ANU must inevitably reach the same conclusion as that of all similar enquiries, both here and overseas.

The Resources Committee has now recommended a new scheme which classifies Computer Science as a discipline requiring no greater level of support than any department in The Faculties and only half the level that some others expect.

The Department of Computer Science cannot accept this proposal. It represents a cavalier disregard for years of patient effort and a partisan view of the relative requirements of various disciplines.

(1) Report of the Working Party on the Specialist Teaching of Computer Science, University Grants Commission (UK), 1976.

(2) Dean of Science report to Science Faculty on 10 August (1996/1982).

.../2

The petition goes on to document more than ten attempts to communicate its needs to the Board, culminating in:

20 Jun 1983

The Department obtained access to the Subcommittee's report and was amazed to find that it contained no reference to the questions raised by the Department over the preceding five years.

The petition includes an appendix documenting eleven areas within the discipline of computer science, of which it says only *Theoretical Foundations* and part of *Numerical Computing* can be regarded as primarily mathematical in nature. A second appendix lists the nine members of academic staff, and lists their research areas.

3.3 The 1985 Review of the Department of Computer Science

In 1985, Ian Ross (DVC) asked Brian Anderson (Professor of Systems Engineering in RSPhysSE) to lead a review of DCS. The normal practice for departmental reviews at the time was that the department being reviewed would have representation on the review panel. Brian tells me that he argued that this should not be the case, because the report of a fully independent panel would have greater weight, and managed to persuade the then Head of DCS to agree to this.

The review received submissions from Robin Stanton (Head of Department), computer science students, Brian Molinari, Mike Robson, Malcolm Newey, Peter Wishart, John Hurst, and John Ophel (Postgraduate Student Representative). All argued strongly for increased resources.

A further submission from the Head of the Computer Sciences Laboratory, RSPhysS (Richard Brent) pointed out the fragmentation of computer science research at ANU, listing DCS, CSL, CMA, Systems Engineering and ARP as among the areas undertaking computer science research. Richard argued for proximity of those groups on campus and for greater cooperation, including resource sharing, involvement of Institute staff in teaching, joint seminars, and shared supervision of postgraduate students.

According to Brian Anderson, the review found it totally inappropriate that Computer Science was funded at the same level as less resource-intensive subjects like Classics. The Review Report (see Page 88) stated that inadequate resources were the major issue and recommended a substantial lift in staffing levels, to be phased in over a five year period.

Two other interesting recommendations were that:

- CS be treated as a laboratory subject and that laboratory sessions with demonstrators be timetabled in the same way as other Science Faculty subjects.
- Software engineering be taught, preferably in third year, with such teaching to include teamwork on a large project.

The first paragraph of a memo from the DVC, Ian Ross (see Page 88) shows that there was significant pushback to the recommendations of the Review. His survey of CS resourcing at other universities was naturally coloured by the fact that other CS departments were also struggling.

The 1985 Faculty of Science Annual Report shows that while computer science enrolments increased by 14%, overall Faculty enrolments dropped by 2%. The 1986 Faculty of Science Annual Report indicates that the Review recommendations were adopted at least to some extent. The Dean references a Faculty academic plan allowing for a five-year period of stability for physics, chemistry and geology, growth in computer science and some contraction in the biologically based departments. The 1986 DCS Annual Report notes that two general staff positions had been added.

3.4 Full Fee Paying International Students

Enrolling fee paying students is a way of expanding the size of the resource cake. In 1985, it became legal for Australian universities to enrol fee paying overseas students. Soon after, when fee paying students were first accepted at ANU, the Science Faculty, under Michael Barber as Dean, took a very conservative approach. Recognising the volatility of this form of income, Michael imposed a one-year delay in spending the proceeds.

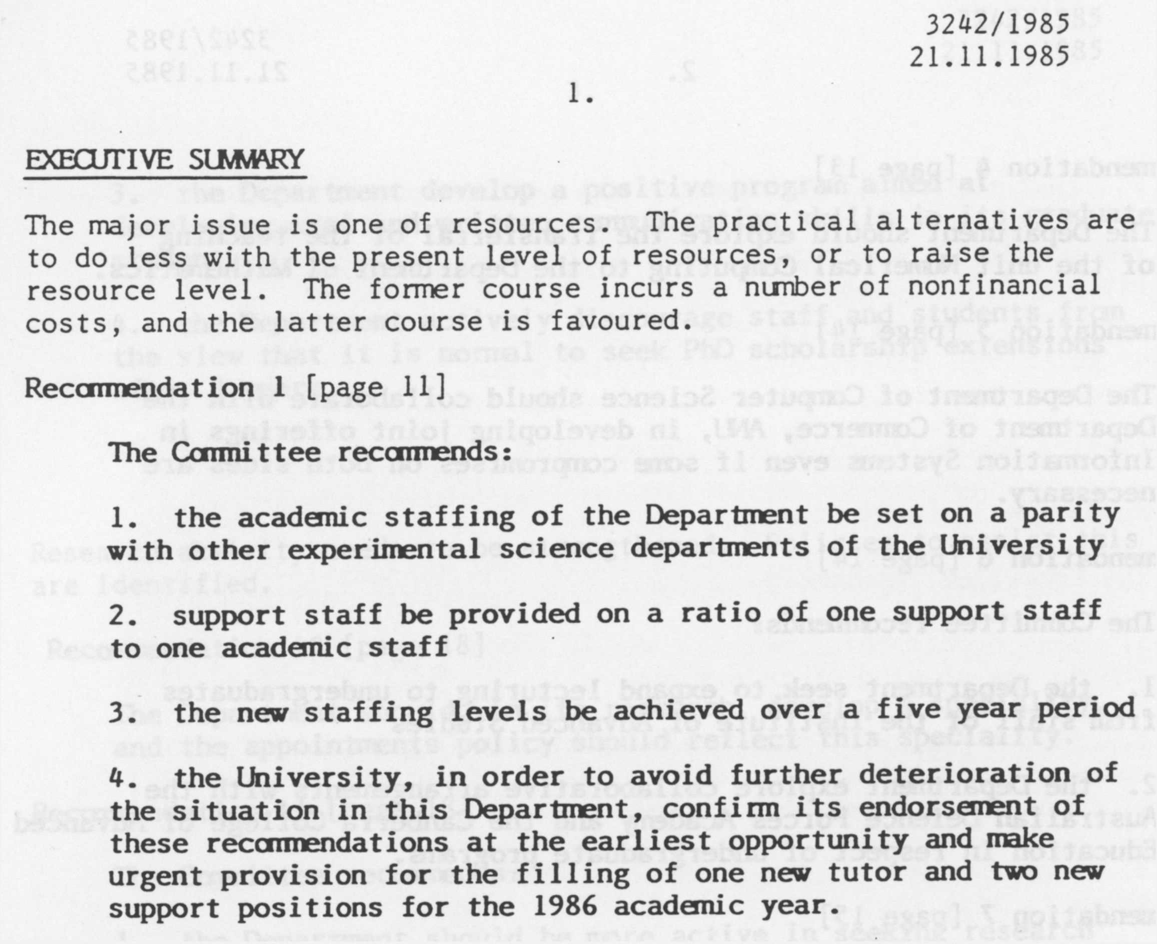

3242/1985
21.11.1985

1.

EXECUTIVE SUMMARY

The major issue is one of resources. The practical alternatives are to do less with the present level of resources, or to raise the resource level. The former course incurs a number of nonfinancial costs, and the latter course is favoured.

Recommendation 1 [page 11]

The Committee recommends:

1. the academic staffing of the Department be set on a parity with other experimental science departments of the University

2. support staff be provided on a ratio of one support staff to one academic staff

3. the new staffing levels be achieved over a five year period

4. the University, in order to avoid further deterioration of the situation in this Department, confirm its endorsement of these recommendations at the earliest opportunity and make urgent provision for the filling of one new tutor and two new support positions for the 1986 academic year.

1985 Review of DCS: Executive Summary and the first of fifteen recommendations

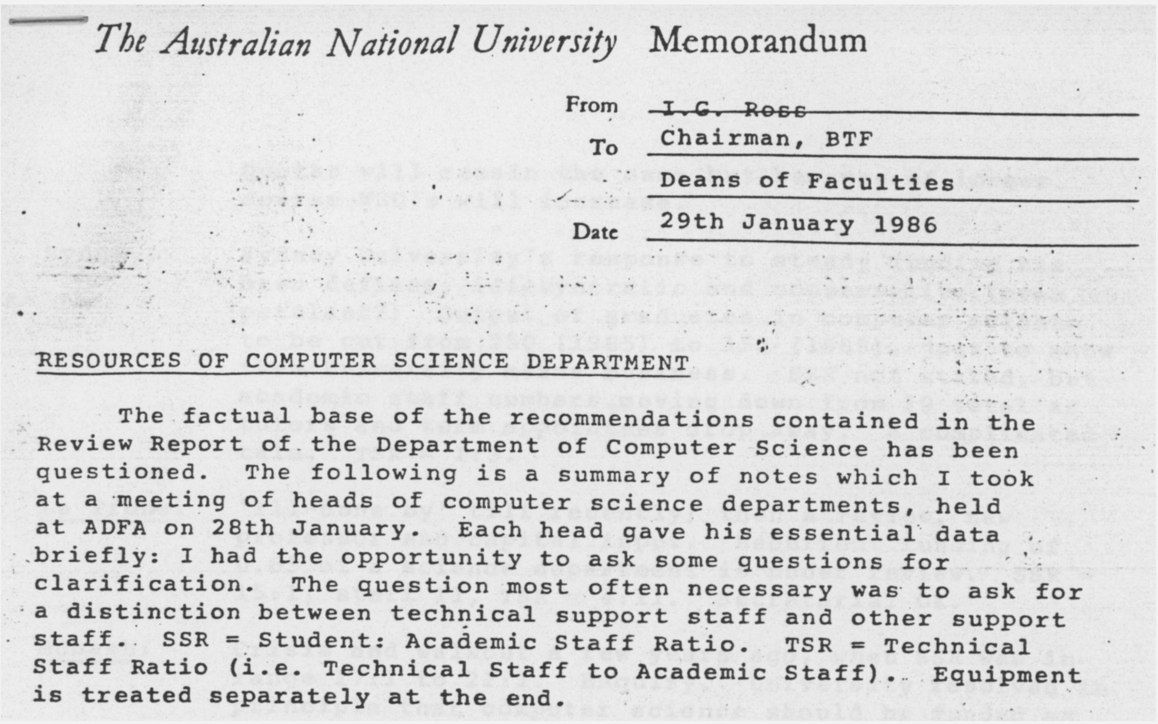

The Australian National University Memorandum

From I.C. Ross
To Chairman, BTF
Deans of Faculties
Date 29th January 1986

RESOURCES OF COMPUTER SCIENCE DEPARTMENT

The factual base of the recommendations contained in the Review Report of the Department of Computer Science has been questioned. The following is a summary of notes which I took at a meeting of heads of computer science departments, held at ADFA on 28th January. Each head gave his essential data briefly: I had the opportunity to ask some questions for clarification. The question most often necessary was to ask for a distinction between technical support staff and other support staff. SSR = Student: Academic Staff Ratio. TSR = Technical Staff Ratio (i.e. Technical Staff to Academic Staff). Equipment is treated separately at the end.

1986: Ian Ross (DVC) memo reporting resources at other Australian CS departments.

Over the decades, the numbers of international students at Australian universities, including ANU, increased dramatically, and the dependence of university activities on fee-paying international students rose in proportion. In 2018 the number of overseas students in higher education in Australia reached 398,563.[1] Details of the massive increase in ANU IT student completions (largely made up of international students) are given in panels on Page 279.

According to the Department of Education, Skills and Employment (DESE) 2019 Financial Performance spreadsheet for all higher education providers,[2] ANU earned $329 million from fee-paying overseas students, approximately 22% of its total revenue of $1509 million. The corresponding figures for the Higher Education sector as a whole were $9979 million, more than 27% of $36,519 million.

When the ability of international students to come to Australia was suddenly curtailed by the COVID-19 pandemic in 2020, the effect on ANU was dramatic.

3.5 Physical Accommodation

Space to accomodate DCS research and teaching was very tight until 1986. Since then the situation has improved in a number of stages.

In the late 1970s and early 1980s, government funding of capital works at Universities was restricted and the allocation to ANU was insufficient to pay for new buildings. Eventually, in the middle of the decade, with the promise of rent from the National Centre for Development Studies (NCDS) and the Australia Japan Research Centre (AJRC), ANU borrowed part of the money needed to construct the Crawford (now Beryl Rawson) building next to the Family Law Court. Initially, the building was to have had an undercroft carpark on the ground level, and to house NCDS and AJRC on levels 2 and 3. In the panel on Page 90 Brian Molinari explains how the design evolved to end up with all three floors built in, and DCS on the top floor. The figure on Page 92 shows how the occasion was marked.

Government funding and ANU's budgeting maintained a strict distinction between capital and recurrent funding. At one stage, when the English Department was desperate for tutors, and The Faculties capital budget was in much better shape than the recurrent one, English unsuccessfully proposed the purchase of two Rolls Royces. Their plan was to immediately sell them and use the proceeds to employ staff.

The Crawford Building provided a much-needed upgrade for DCS. For the size the department was then, there was adequate accommodation for staff and postgraduate students in offices (shared offices for postgrads) with windows. The tea room was a warm and inviting space – many productive conversations took place there.

At various times, the inner core of the Crawford floor accommodated a space for shared-use staff and postgrad workstations, a robotics laboratory, a student workstation laboratory, and a Fujitsu AP1000 supercomputer.

For the arrival of the AP1000, I made contingency plans to have windows removed from the tea room and to hire a crane, but it turned out that the biggest component would just fit into the lift. Mind you, the main unit weighed about 950kg, close to the lift limit. When it was wheeled in, the lift dropped a few centimetres, and ANU and Fujitsu staff collectively held their breath hoping that it would stop at the right level on the top floor. Fortunately it did.

3.5.1 Amusing Cross-Overs and a Bad Smell

For a little light relief from resource battles, here are a couple of amusing anecdotes, related to buildings.

After DCS moved into the top floor of the newly constructed Crawford Building, in the winter of

[1] https://www.aph.gov.au/About_Parliament/Parliamentary_Departments/Parliamentary_Library/pubs/rp/rp1819/Quick_Guides/OverseasStudents

[2] https://www.dese.gov.au/higher-education-publications/resources/2019-higher-education-providers-finance-tables

1986, it soon became apparent that there was a significant problem with the airconditioning. People on the north side of the building were roasting and people on the south were freezing. Adjusting the thermostat controls made things worse!

Yes, the thermostats were cross-wired. As the temperature rose on the north, the airconditioner on the south started cooling, causing the airconditioner on the north to start heating.

Brian Molinari tells the story of the DCS move from Copland to Crawford

In 1985 the Department was still located in a string of rooms in the Copland Building, where we started in 1971 in the Faculty of Economics. All attempts by successive Heads (Richard Brent and Robin Stanton) to negotiate appropriate accommodation had come to nothing. To be fair, there was not much creation of new space on campus at that time, and most departments in The Faculties were ensconced in their ”own” building and were not minded to share it with newcomers.

My memory of our deliverance goes as follows.

Robin would regularly visit the Chancelry – we presumed for lobbying and information-gathering purposes. He kept his cards fairly close to his chest.

After one visit he admitted that he had found out about a new building under construction, to be named the Crawford building. It was a small building on Ellery Circuit. The top two floors were assigned to research groups from the Institute of Advanced Studies while the ground floor was to be car parking. Low potential for DCS, one would have thought.

After a subsequent visit he gave us the news that the university had agreed to forgo the car parking, and fill in the ground floor as office accommodation. And, mirabile dictu, that it would be assigned to DCS. The marginal cost of the upgrade had obviously proved acceptable.

After a subsequent visit he gave us the news that the occupants of the building (now including DCS) were to be repacked into the building to make best use of the space. The nominal occupants of the top floor, the Australia-Japan Research Centre (AJRC), apparently didn’t need all of the space whereas DCS more than overfilled the smaller ground floor. And, mirabile dictu, our locations were to be reversed. We moved into our new space in mid 1986.

The Head of the AJRC, Peter Drysdale, would occasionally come upstairs and walk around admiring the view. “Bastards” he would say. But with relatively good humour, as he knew when he had been outmanoeuvred.

Robin Stanton downplays the significance of his role in securing the top floor, though he says that he, “did let the University know how difficult it would be for CS to work as a harmonised department working across two disjoint ground floor areas, with a smaller usable footprint.” It was the Vice-Chancellor Peter Karmel who decided that DCS should have the top floor. Robin says that Peter Drysdale’s initial reaction fell short of fulsome good wishes toward computer science.

He also said that AJRC’s plan for the top floor had been to use all the perimeter offices, and make the windowless interior space available to the University.

A similar commissioning problem applied to the lecture theatres in the then new Manning Clark Centre (now demolished). Robin Stanton (course convenor) was introducing Steve Blackburn’s first lecture to a very large first year CS course but found that the radio microphone in the newly constructed theatre wasn’t operational. Undaunted, he introduced the lecture by enthusiastically commending his students for their good sense in taking Computer Science, the path to the future.

Eventually, someone came in from an adjacent lecture theatre, asking if Robin would stop trying to steal his students! Yes, the radio mic in Theatre 1, was connected to the amplifier and speakers in Theatre 2.

In the late 1980s, students complained for several months that the terminal room in Copland G8 stank of urine. There was a fresh-air ventilation system but it seemed inadequate to the task. Assuming that the problem was caused by the fug associated with a room full of sweaty students, building management installed air conditioning units to cool the room. This cooled the fug, but made no difference to the smell.

Eventually, I crawled into the roof space – all part of my duties as DCS Head Programmer – and discovered that the ventilation system sucked air from a narrow void within the building. That void contained various building services including the sewer pipe from the mens' toilet next door.

Yes, the sewer pipe was leaking, and the G8 fan was giving our students the full benefit.

I wrote a detailed report on the problem to Facilities and Services and the Business Manager, including diagrams, and taped over the ventilation switch with a note saying, "DO NOT SWITCH ON". Problem solved!

Not so fast! Next time I entered the room, someone had removed the tape and modified my note to say, "DO NOT SWITCH OFF"!!!

Peter Bailey recalls coffee and other social rituals

One of my early recollections of being a staff member in the department, which happened in 1992 after my honours year, was the morning coffee ritual in the Crawford Building when pretty much everyone in the department would somehow scramble into a slightly too-small room and sit around and drink coffee for half an hour, chatting, telling stories and joking around. It was an extremely important social ritual.

I had the habit of coming into work early and would sometimes catch Robin Stanton, who was also an early starter, presumably to get stuff done before his much busier day than mine. We would periodically go and get what he referred to as an Italian coffee down at Calypso which was in the Student Union Building. He would always insist on paying for it, regardless of my ability to do so.

In later years after the move to the CS&IT building it was my perception that a lot of the closeness was lost. We had CSIRO people in the shared tea-room and they had their half and we had ours. There was a little bit of mingling, for example between Trevor Vickers and Kerry Taylor. [**Dave:** They are husband and wife!]

That was one of the reasons that I started going outside for coffee. I could then afford to buy coffee and I started going to the Gods or the Purple Pickle. I did a lot of work on my PhD there, reviewing chapters, and poring over things before taking them to Malcolm or Robin.

Then there was the saga of obtaining an expresso machine for the department. Steve Blackburn drove the rigorous reviews of machines and grinders. I think it was a Saeco machine.

A group of us who were doing PhDs at the time I came back from Edinburgh thought that we shouldn't just be sitting in isolation in our offices thinking deep thoughts but that we should get together, share experiences and present our work and ideas. The department provided some funding for a bottle of red wine, some non-alcoholic beverages and some cheese and biscuits and we met once a week late on Thursday afternoons.

3.5.2 SPAM & SPAM-2

Later in the 1980s, ANU was told that it would not receive government funding for any new buildings until its building stock was brought up to an acceptable standard. ANU was able to bid for funding to renovate, refurbish, and rehabilitate existing teaching buildings. There were two rounds of funding and the very complicated renovation programs were called SPAM and SPAM-2. Much as I love contrived acronyms, I have no inkling of what SPAM stood for. Neil White and Rod Whitty were employed by Peter Scardoni to oversee the SPAM and SPAM-2 processes.

I remember that there was a long chain of moves, in which a building was refurbished, department A moved in to it, department A's former building was refurbished, department B moved into it, department B's building was refurbished, ... It was an era of peak acronym – The Botany and Zoology departments were merged into BoZo, and Biochemistry and Molecular Biology became BaMbi.

The SPAM process caused distress in some areas. Geology moved into the former Botany Building, leaving behind the building which had been specifically designed to accomodate their needs. Andrew Cockburn, head of BoZo, railed in his annual report about the unsuitability of their new accommodation and suspected that geologists were equally unfortunate. As he said, "People who live in glasshouses shouldn't study stones."

Geology moving out of its building made room to accommodate the newly created Department of Engineering, with Darrell Williamson as Head. Since DCS had outgrown the top floor of Crawford, the University came up with the idea of housing DCS in an extension to be built on the roof of the new Engineering Building. Since this would have been a smaller space than their existing accommodation, I guess the idea was that DCS and Engineering would fight it out over space.

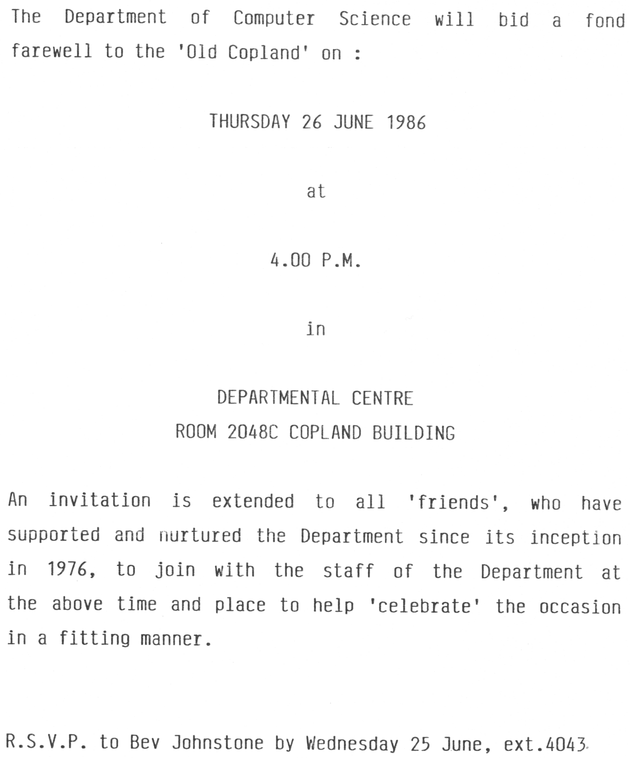

The Department of Computer Science will bid a fond farewell to the 'Old Copland' on :

THURSDAY 26 JUNE 1986

at

4.00 P.M.

in

DEPARTMENTAL CENTRE
ROOM 2048C COPLAND BUILDING

An invitation is extended to all 'friends', who have supported and nurtured the Department since its inception in 1976, to join with the staff of the Department at the above time and place to help 'celebrate' the occasion in a fitting manner.

R.S.V.P. to Bev Johnstone by Wednesday 25 June, ext.4043

1986: Marking the DCS move from Copland to the top floor of the Crawford (now Beryl Rawson) Building.

When I heard stories of the "Flaming As" ritual[3] among Darrell's Engineering cohort at UNSW, I wasn't quite sure what to expect from him. He was certainly capable of informality. One day I was standing talking to the Engineering Administrator, Joanne Perks, when I suddenly fell to the ground. Darrell had sneaked up behind me, knelt down and pushed against the back of my knees! He told me once of having to jump his bicycle over a cattle grid after descending Fitz's Hill at 90 km/hr.

In the meantime, DCS and Engineering were cooperating closely. DCS initially provided staff support for Engineering's computer laboratories and software. Peter Farmer and I wired up a laboratory of 20 Sun SLC workstations, and some facilities for Darrell, Joanne and Sharon Goddard. Ruth Jarvis, a member of the first intake of Engineering students, and niece of Ray Jarvis, reminds me that the SLC lab was located in the Physics Building and was used to teach AutoCad. The lab was called the Pub and the machines were called "Fosters", "XXXX", "Coopers" and so on. Later the Pub moved to the ground floor of the Engineering Building near the foyer, and a new lab of colour Sun workstations was set up by Bob Edwards on the other side of the foyer. Those workstations were named after colours.

This was occurring at around the same time as the mooted merger between the Faculties and the College of Advanced Education located at Bruce. Funding for a new ANU engineering and computer science building was granted, but was constrained to be spent on a building on ANU's Bruce Campus!

3.5.3 The CS&IT Building

The next step in improving DCS's space shortages and the need for more and more teaching laboratories was the construction of the CS&IT Building.

In December, 1990, Robin Stanton started a bid to gain funding for an Engineering / CS annexe next to and connected with the Engineering building. He asked me to sketch a site plan of an Engineering / CS precinct showing the close connection to the Science precinct and the Maths precinct.

[3] I understand it involves a burning copy of the Daily Telegraph and a run in bare buttocks.

He also asked me to sketch floor layouts for the two floors of the proposed building, taking some inspiration from the design of the Engineering / CS building at Bruce. His logic was that it is much easier for a committee to think about specific concrete proposals and approve them. I also calculated net and gross space, counts of offices, and a breakdown between teaching, research, postgrad and office space. I still have copies of those sketches but they are too faint to include here. I'm afraid they look a bit scrappy!

Eventually, approval was given and the process of designing of the CS&IT building commenced. It was greatly complicated by expansions of scope. The first such scope expansion was that the new building should be extended to house a considerable number of undergraduate teaching laboratories. The second was a desire to accommodate the CSIRO Division of Information Technology (DIT) along the lines of a new shared building at Macquarie University. DIT were keen on collaboration, and on access to ANU postgraduates. At the time DIT, led by John O'Callaghan, was housed on the ANU campus in transportable buildings. CSIRO floated the idea of locating the new shared building on CSIRO's Black Mountain campus, possibly re-purposing some of the old DCR, then CSIRONET buildings (now demolished). ANU rejected the idea.

By the time the building was complete, DIT had merged with CSIRO Mathematics and Statistics to form CSIRO Mathematical and Information Sciences (CMIS).

ANU Facilities and Building Services imposed several constraints on the building, designed to limit the cost. The most notable of these was that funding would in no circumstances be provided for separate seminar rooms or tea rooms. The shared tea room caused a lot of difficulties.

- CSIRO maintained a secure wing with swipe card access. This made it inconvenient for their staff to access – getting a cup of coffee to drink back in your office often resulted in coffee stains on the floor.
- CSIRO operated a tea club. Their refrigerator contained cans of softdrink, perhaps even beer, and a tin of money. Because DCS allowed open access, the CSIRO refrigerator had to be fitted with a padlock.
- When CSIRO's padlock was broken and money stolen, ANU students were of course suspected. However, Arch Brayshaw (CMIS) rigged up a webcam and discovered that the culprit was in fact a cleaner, promptly dismissed. (He should have realised that students would take the beer!)

- CSIRO wished to reserve the right to close the tea room for functions, but DCS insisted that their staff should be able to make a cup of tea at any time, regardless of CSIRO functions. That led to the need for two kitchens.
- It was my brilliant idea to fit the DCS kitchen with twin dish-washers, one clean, one dirty.
- The shared tea room was supposed to encourage collaboration between CMIS and DCS but that didn't happen much, except on the occasion of the Melbourne Cup lunch.

The funding agreement was that CSIRO would pay the cost of constructing their part of the building plus a share of the cost of shared spaces, in return for a 25-year no-cost lease, with an option to extend for a further 25 years. CSIRO insisted that they should have access to the DCS seminar room on the ground floor and that it should have a capacity to hold all-staff meetings of 100 seated people. They paid 10% of the cost of the room.

Once it was determined that the new building would be in the originally planned location, it was decided that it would be three storeys but would be designed to "read" as two storey, and that it would have two wings connected by a curved central section. At the time, buildings on campus (other than the Chancelry) were required to be no more than three storeys. In recent years, taller buildings seem to be preferred.

The CS&IT architect was John Firth of Lester Firth and Associates. I was DCS's representative on the Project Control Group (PCG) and Peter Fox and Arch Brayshaw represented CSIRO. CSIRO started planning their space needs from the South-East end, and thought they would take up as much of the central section as they needed. I eventually persuaded the PCG to agree on a DIT-DCS boundary in the central zone and for each to plan from there. Only that way would the tenants discover whether the building was the right size to accommodate their needs.

Eventually it was decided that the DCS wing was not sufficiently large, and Robin Stanton persuaded Peter Scardoni to support the extension of the building by 6 metres in length, a total of 324 square metres, gross. Don Hardman, Head of Facilities and Services, was opposed. I was a fly on the wall at a meeting between the three of them, which was about the tensest meeting I've ever experienced. Sparks flew, the air crackled and the building grew.

R&F Building Management was selected to provide Project Management services, and the Project Manager was Gary Gisik. The Project Control Group met regularly and I spent two days a week on the new building for a period of more than two years. Early on, the PCG was told that minimal cost for the concrete structure could be achieved either with a 6m or a 7.2m spacing of columns.

My considerable experience with student workstation laboratories suggested that 6m was too narrow for the then favoured layout for a lab of 20 workstations. In this layout each desk was 900 deep and 1200 wide, and there were four rows of five desks with the two middle rows back to back and the others against the side walls. You needed 1800 between the edge rows and the middle to accommodate chairs and for the tutor/demonstrator to be able to circulate. With this layout, all students were able to see the tutor, and the whiteboard and projection screen at one end. This led to an overall width of 7.2 metres.

The artwork *Computer waste bin* by Margot Zanstra which used to grace the main entry of the CS&IT building. It was purchased specifically for our building, and former DVC Ian Ross personally contributed to the purchase. Unfortunately, recent renovations did not provide a place for it. *Photo: ANU Art Collection.*

I therefore argued that the column spacing in one dimension should be increased to 7.2m. I won, but only on the ground floor. However the benefit from this was substantial. In all of the ground floor labs there were no columns impinging on the space. One lab which was, late in the piece and against my advice, moved to level one, had a column in it. That significantly changed the feel of the room and meant that three of the students were unable to see the whiteboard.

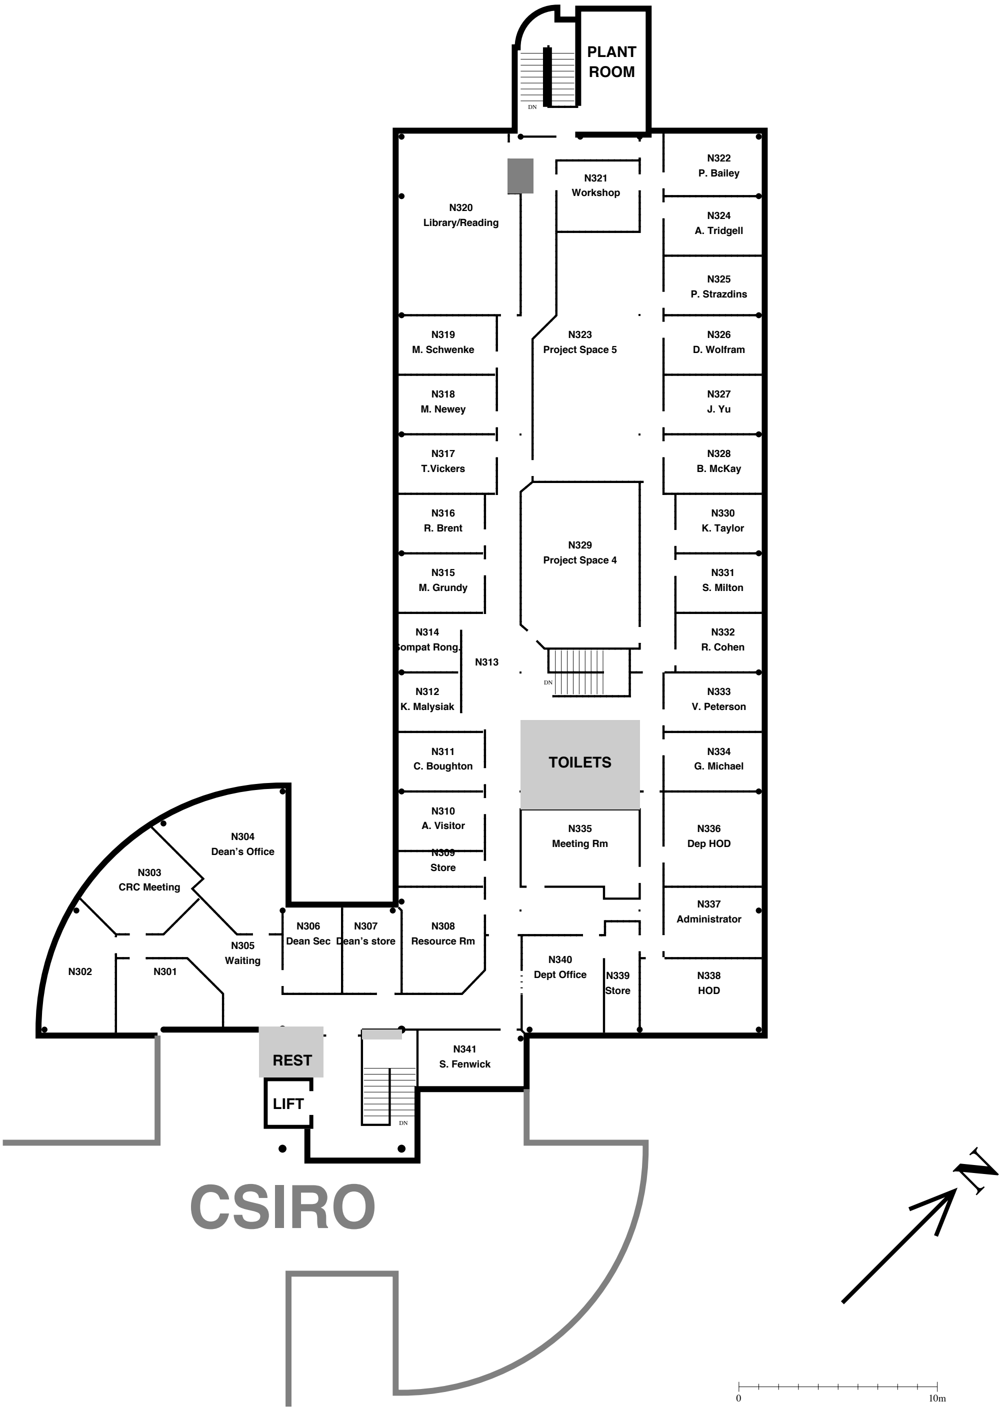

CSIT Bldg: ANU Wing, Level 3

Scale - 1:300
David Hawking
15 May 1995

My layout for Level 3, in hand-written PostScript, of the CS&IT Bldg, with names of occupants, 3 months after move in.

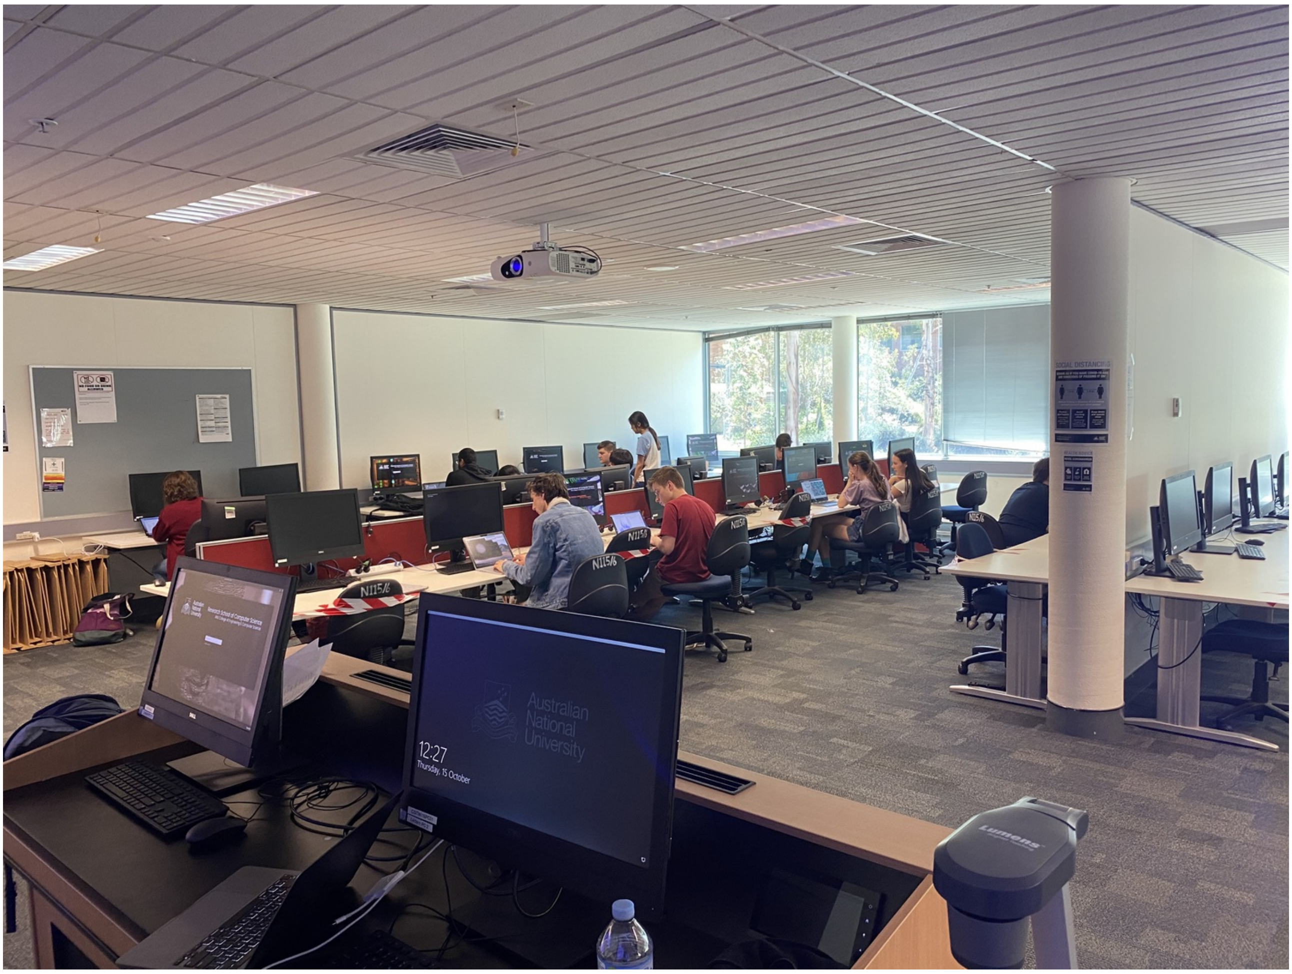

Computer labs CS&IT rooms N116/N115. The number of workstations in each lab has been increased by narrowing the desks. The dividing partition between the labs has been partially removed to allow both labs to be treated as a single one. Due to the COVID-19 pandemic, every second desk has been taped off.
Photo: David Hawking. October 2020.

Over the succeeding decades, the optimal design for student computer labs changed. Significant increases in student enrolments led to a need for greater numbers of students in a lab session. Changes in deployed technology have reduced the size of desk needed. Keyboards are considerably narrower than those in the early 1990s, and students may be less likely to bring writing pads or textbooks, meaning that the width of a desk can be safely reduced to 1000 mm. The required depth can also be reduced because, in 2021, the standard workstation is an all-in-one Linux computer and LCD panel rather than a combination of a pizza box computer and a CRT monitor. In the computer labs in the new Hanna Neumann Building, I believe that the width per student has been reduced to 900 and the depth to 650, allowing around 40 students in each lab. However, the reduced depth has led to an increase in the incidence of cable connectors being kicked out of sockets on walls.

Much as I liked John Firth, I was unable to communicate to him DCS's requirements for the layout of internal partitions. He kept reminding me that, in commercial office space, financial considerations would dictate a much denser layout – his own staff had desks in a corridor. After a couple of unsuccessful iterations, I drew the building to scale in the form of a hand-written PostScript program, laid out all the internal partitions, and said, "Please draw this." At the next PCG meeting, Chris Culverwell (the chair) said, "and now the architect will present the current state of the plans." John stood up and said, "We're not architects, we're draftspeople!"

My idea was that, apart from the Dean and the Head of Department, and where the building shape dictated otherwise, there would only be two sizes of office. Each would be 3m wide to suit the column spacings and window layout and to maximise the number of rooms with external windows. Smaller offices would be for one person only, intended for people such as tutors with high student consultation loads, or for PhD students in writing-up phase. Larger offices would be for academics or to be shared by two PhD students. Having standard sizes was intended to avoid the need for a cascade of office moves when someone was promoted or a new person was appointed.

The windowless central areas in the DCS wing were for workshops, machine rooms, meeting rooms and project space. The third floor areas had skylights, which became a liability in the 2007 hailstorm.

CMIS went in a very different direction, with large open plan office areas in the centre and several offices extending 4.5m along the external walls.

In the plan of Level 3, on Page 95, you'll see a small space near the lift, labeled REST. This was supposed to be a sick bay, but it was pointed out during a DCS meeting that it wasn't actually long enough to enable a tall person to lie down. Trevor Vickers said, "I wouldn't worry about it. At UNSW the sick bay was only used for having sex", to which his wife Kerry Taylor replied, "Trevor, you didn't!!!".

When it came time to plan for the move, plans like the one on Page 95, initially without occupant names, were given to people in descending order of seniority with an invitation to pick their preferred office. Peter Bailey describes his amazement that, on arriving back in the department after this process had completed, he found that N332 was still free for him to claim.

Peter Bailey returns from Edinburgh to a sunny office in CS&IT

I came back to ANU in October 1994 having spent six months in Edinburgh, avoiding a Canberra winter – three months working on my PhD, and three months employed by the Edinburgh Parallel Computing Centre (EPCC). I spent lots of time on weekends with Quintin Cutts and Graham Kirby who were studying for their PhDs at St Andrews still. My role when I came back, I was still doing my PhD, but I was working 20 hours a week as a research programmer on the CAP Programme. We were still in the Crawford Building at that time, and you, Dave, showed me the plans for the new building and said, "Here's the office I've picked out for you. Would that be OK?" You were clustering the programmers around the light well next to the cake slice part of the building on Level 2.

It was a slightly odd office. The window was divided, and the office looked as though it might be quite gloomy, and I said, "Are there any other offices available?" and you showed me the plans, and we looked up and down the offices on Levels 2 and 3, and I saw this office on the north-west end of the building, on Level 3. I asked if that was available. You said that yes it was, that all the academics, in strict order of seniority, had picked out their preferred office and that this one was free.

It was a lovely office on a corner with windows on two sides and mature eucalypt trees at leaf level outside, so it felt you were almost sitting inside the canopy. It was quiet and secluded, but as far away from the tea room as you could get. [**Dave:** It's a great office for an academic trying to avoid students. It effectively has a second exit straight down the stairs on the north end of the building. Interesting that none of them took it.]

In May 2018, I proposed, with the support of the Dean, CECS (Elanor Huntington) and the Director of RSoCS (Alistair Rendell), that the CS&IT building be renamed in honour of Robin Stanton. His contributions to computing and computer science at ANU have been exceptional, and more than any other, he was responsible for the existence of that building.

The response from the ANU's Naming Committee was that the ban on naming buildings after living people which was introduced subsequent to the naming of the Brian Anderson Building in July 2016, had recently been weakened a little:

> "...in rare or exceptional cases, the University may select living honourees to balance naming. However, unfortunately that clause wouldn't apply in this instance. It is geared towards us rebalancing the naming on campus, which at the moment is majority Caucasian males."

3.5.4 Even More Space!

Soon after the CS&IT building was completed, construction started nearby on an architecturally similar building now known as the Brian Anderson Building. It was built to house RSISE and now houses some members of DCS.

On 25 March 2019, the stylish new Hanna Neumann Building was officially opened. It replaced the old Chemistry Building and Chemistry Lecture Theatres. In addition to housing members of DCS, and large teaching laboratories, the building is home to the Mathematical Sciences Institute (MSI) and the Co-Lab, a partnership between DCS, MSI and the Australian Signals Directorate (ASD). I'm not sure whether it was an intentional design element of a building housing ASD, but the decorative metal screens on the exterior of the building significantly attenuate mobile phone signals.

Alistair Rendell recounts the origin of the new Hanna Neumann Building:

> With our expansion plans starting in 2014, came the obvious need for additional physical space. To this end VC Ian Young changed the way new buildings were funded, initiating a competitive call for proposals. Under Dean John Hosking we joined with the College of Mathematical and Physical Sciences to put in a bid for a new building that would house all of mathematics and some of computer science. We asked for around $50M, were allocated $40M and told that the building had to go where the old Department of Chemistry building was and that demolition costs had to be included within the $40M. Needless to say there followed a long process of negotiation around funding and working with the architects to try and create something that would be acceptable. Initially we were heading towards a 3 story building with offices that were nine and a bit square meters. However just as this was converging towards a solution, myself, Raj Goré, and some of the mathematicians had a meeting with the Australian Signal Directorate to discuss how we might work closer together. In particular how we might assist them to train and subsequently recruit individuals skilled in computing and mathematics. Somewhat speculatively I suggested that they might want to fund an additional floor on our joint building and use it as a base for collaborative projects and student activities. One thing led to another and what was going to be a three storey building is now the five storey Hanna Neumann building.
>
> To design and configure a new building is hard when there is one party, but when there are mathematicians, computer scientists and public servants it is a special challenge. I quickly learnt that mathematicians love their chalk boards and individual offices. In contrast, Computer scientists have moved to the 21st century, using white boards and working cooperatively in shared spaces with breakout rooms. The Australian Signals Directorate on the other hand are bound to strict secrecy and unable to tell you anything! Testimony to these dilemmas is found in the seminar room on the ground floor where there are white boards hidden behind black boards (or vice versa)! Overall however I think the final outcome looks pretty good and works fairly well (but I may be a little biased!).

DCS is now very much better off for space, and in closer proximity to its related disciplines, than at any time in its earlier history.

3.6 1988: Threats of Amalgamation

The proposed Dawkins reforms of the late 1980s had substantial potential impact on resourcing for DCS, and some actual impact. Up until 01 January, 1989 there were three distinct sectors within Australian tertiary education: Universities, Colleges of Advanced Education, and TAFE colleges. Because of dramatically rising enrolments, the government was desperate to reduce cost per student and determined that this could be done by removing a sector and by creating larger institutions.

From 1988, a process initiated by John Dawkins, federal Minister of Education, Employment and Training, resulted in a sequence of amalgamations, and the end of many small teachers colleges, colleges of advanced education and institutes of technology. Vice Chancellor salaries were to an extent based on student enrolments, and some Universities, like Monash, embarked on a program of wholesale incorporation of smaller institutions. The goal was to reduce administrative overheads under the banner of creating a Unified National System of higher education.

The Dawkins reforms were quite revolutionary. A paper, *The Dawkins Reconstruction of Australian Higher Education* by Grant Harman, an academic from UNE, was presented at the 1989 Annual Meeting of the American Educational Research Association and is available online at `https://files.eric.ed.gov/fulltext/ED308754.pdf`.

The Dawkins strategy followed one playing out in the UK. Many in the university sector of both countries took enormous delight in watching the two series of the British TV satire, *A Very Peculiar*

Practice.[4] In the series, the fictional Lowland University is merged with the Police College at Hendon. No doubt inspired by the TV show, the Goulburn College of Advanced Education became a police college.

L: John Dawkins *Photo: Department of Foreign Affairs and Trade website – www.dfat.gov.au, CC BY 3.0 au, https://commons.wikimedia.org/w/index.php?curid=28173659* **R: Ian Chubb** *Photo: Lizette Kabré, University of Copenhagen, Creative Commons Attribution-Share Alike 2.0 Generic.*

Another Dawkins measure to control the burgeoning cost of higher education was the re-introduction of fees and the introduction of an income-contingent student loan scheme HECS (Higher Education Contribution Scheme), now called HELP (Higher Education Loan Program).

On 29 July 1988, all ANU staff were advised by the Vice-Chancellor, Laurie Nichol, that the Minister had advised his intention to amalgamate ANU, the CCAE, and the Canberra Institute of the Arts (CITA), and had invited ANU to respond. Changes to the size of the governing body (Council) and changes to the basis for future funding were also foreshadowed by the minister. The General Staff Newsletter of 25 November 1988 reported that it was the preference of Council that ANU retain its separate identity, with adequate funding, but that if that couldn't happen it would work to make the amalgamated institution a success. A memo of 07 April, 1989 from the Registrar, Rosalind Dubs, advised that Council had resolved to enter into a tripartite agreement between the three institutions covering the principles of amalgamation. The memo included a copy of the agreement.

This didn't happen because legislation to effect the amalgamation was rejected by the Senate. The General Staff Newsletter of 14 August, 1989 reported that ANU and CITA had already agreed to amalgamate by January, 1991, and that John Dawkins now proposed to establish CCAE as a university unless ANU wished to amalgamate with it.

A wide ranging consultation took place across campus. Peter Scardoni organised large public meetings of general staff to canvas opinions. In one of them I spoke against amalgamation and he asked me if I was prepared to bet my job on it.

Deane Terrell, Dean of Economics, led a campaign against amalgamation with the CCAE which was ultimately successful. Another prominent opponent was the historian John Molony, president of the ANU Staff Association and former catholic priest. I'm reminded that at a staff meeting organised

[4] https://en.wikipedia.org/wiki/A_Very_Peculiar_Practice. The Wikipedia article says that, in 2010, the Guardian ranked the series as the fifth best TV drama of all time.

in Melville Hall by Laurie Nichol, John marched to the front of the hall, and took the agenda in an anti-amalgamation direction.

The battle was not over yet. On 16 January 1991, Acting VC, Max Neutze, circulated copies of the report by Ian Chubb,[5] Chair Higher Education Council, National Board of Education Employment and Training (NBEET), to the new minister, Peter Baldwin. It recommended that the Institute of Advanced Studies be split off from the Faculties and be retained under Commonwealth legislation while the Faculties be merged with the CCAE, under ACT government legislation. Together with the Chubb Report, Max Neutze circulated copies of a press release he had issued, pointing out that the split-up of ANU was opposed by a large majority of staff, and that splitting up a successfully amalgamated institution seemed contrary to the intentions of the Dawkins reforms.

On 8 February 1991, ANU Council met and unanimously passed eleven resolutions opposing the Chubb proposal. One practical reason for the Institute of Advanced Studies (IAS) to oppose it was that, to a considerable extent, they relied on The Faculties for PhD students since, unlike The Faculties, IAS was not eligible to receive government funding to support PhD scholars.

Roger Scott, Principal of the CCAE at the time, describes the failed amalgamation process, and the strategy he employed in negotiations, in a paper in *Australian Universities Review*, 47(1), 2004, pp. 10 – 20, `https://files.eric.ed.gov/fulltext/EJ848182.pdf`. His paper is titled: *A personal memoir of policy failure: The failed merger of ANU and the Canberra CAE.* Several notable points:

- One of the sections is soothingly titled: *Shotguns and Chicanery*!
- CCAE pre-emptively reclassified Principal Lecturers and College Fellows as Associate Professors and Professors to avoid disadvantage to College staff post-amalgamation.
- Scott argued that, post-amalgamation, teaching loads should be harmonised between former college and former ANU staff. They were very different at the time, as ANU academics were expected to do research and CCAE academics were not.
- The whole ANU-CCAE saga played out in the context of the imposition of self government for the ACT, which went ahead despite a referendum vote against it. Decisions were delayed pending consideration by the about-to-be-created Legislative Assembly.
- A Dawkins amalgamation carrot was the funding of a large Engineering and IT building on the CCAE campus and denial of special funding to ANU for new developments. [Plans for the CCAE building had DCS names on the offices.]

Eventually, the ANU remained intact, and the Faculties-CCAE amalgamation did not occur. Instead CCAE became a university under three years of tutelage from Monash University.

Despite the concerns, I kept my job, and ANU and the University of Canberra went on to be very successful institutions, each significantly increasing their enrolments.

3.7 Equipment for Teaching

Because the relative funding model, as implemented, provided DCS with no equipment funding, the only internal way of obtaining equipment was by submitting bids to ANU's Computing Policy Committee (CPC) or to the Major Equipment (Boardman) Committee. In the early days, DCS had little success.

After Brian Molinari became a member of the CPC and observed how things worked, DCS adopted a different and more successful strategy for funding student computing facilities. Instead of a DCS bid for a DCS student computing laboratory, a larger, combined Faculties bid for general-access computer laboratories would be submitted after extensive consultation and discussion across departments.

The results of this strategy are described in the chapter on Student Computing, see Page 232.

[5] Later ANU Vice-Chancellor

3.8 External Funding for Research

In 1987, DCS was successful in obtaining a large industry-collaboration grant to develop software for a robotic arm. The KRIS project was a collaboration with Scientific Instruments Australia, represented by a former DCS student, Flash (James) Gordon. One of the commercial goals was to provide a robot which could transfer delicate, unfired porcelain sanitaryware for Caroma. Brian Molinari and Chris Johnson were project leaders and Paul Mackerras and Thea Clark worked on the project. The operating system was VxWorks, a real-time version of Unix. The vast bulk of the code was written in Modula-2.

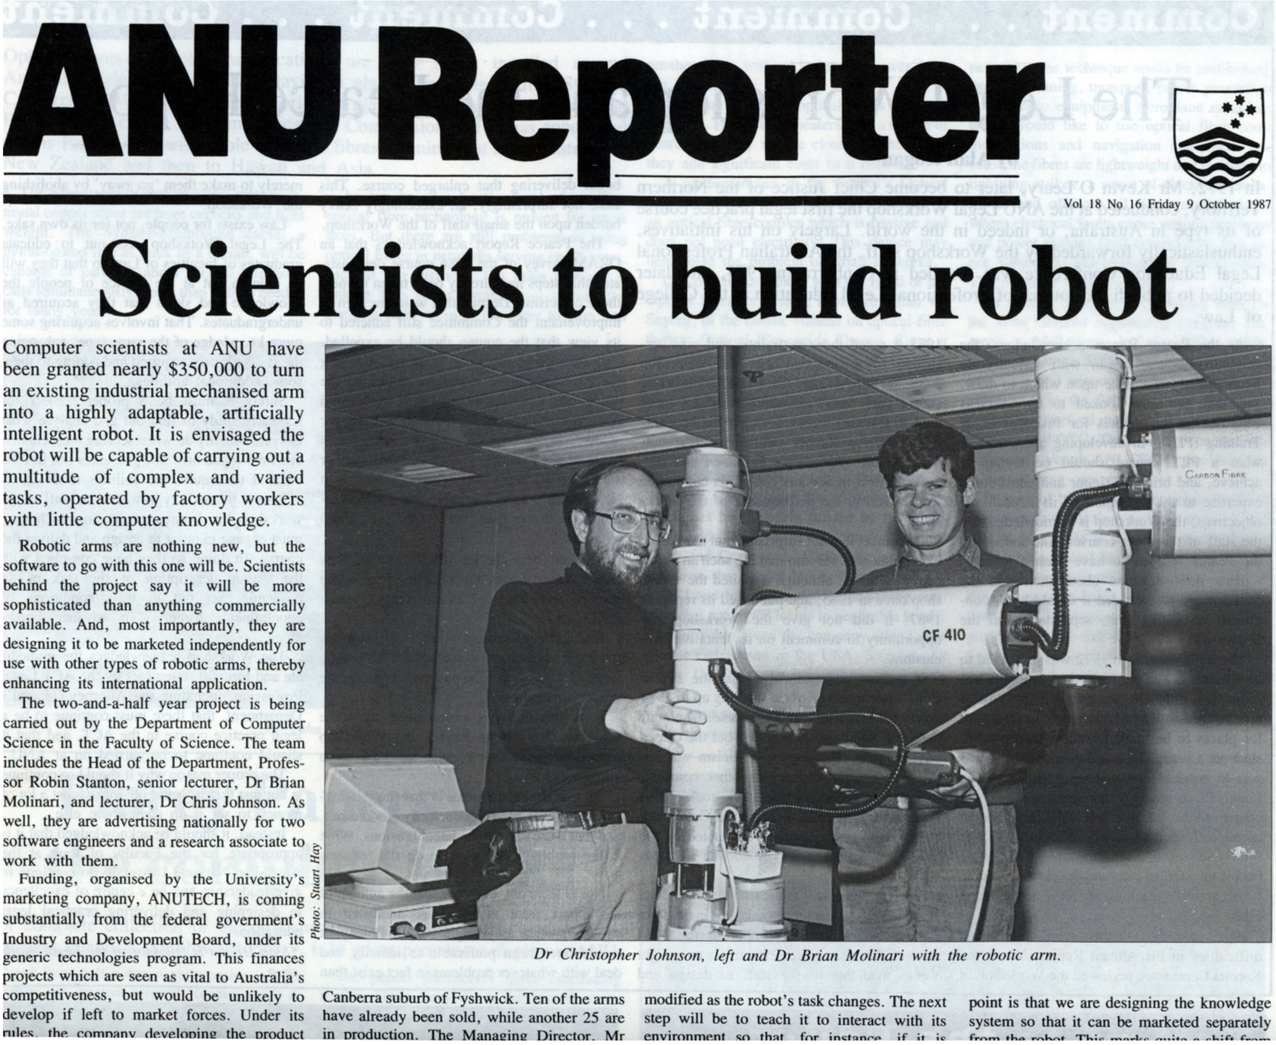
CF 410
CARBON FIBRE

ANU Reporter

Vol 18 No 16 Friday 9 October 1987

Scientists to build robot

Computer scientists at ANU have been granted nearly $350,000 to turn an existing industrial mechanised arm into a highly adaptable, artificially intelligent robot. It is envisaged the robot will be capable of carrying out a multitude of complex and varied tasks, operated by factory workers with little computer knowledge.

Robotic arms are nothing new, but the software to go with this one will be. Scientists behind the project say it will be more sophisticated than anything commercially available. And, most importantly, they are designing it to be marketed independently for use with other types of robotic arms, thereby enhancing its international application.

The two-and-a-half year project is being carried out by the Department of Computer Science in the Faculty of Science. The team includes the Head of the Department, Professor Robin Stanton, senior lecturer, Dr Brian Molinari, and lecturer, Dr Chris Johnson. As well, they are advertising nationally for two software engineers and a research associate to work with them.

Funding, organised by the University's marketing company, ANUTECH, is coming substantially from the federal government's Industry and Development Board, under its generic technologies program. This finances projects which are seen as vital to Australia's competitiveness, but would be unlikely to develop if left to market forces. Under its rules, the company developing the product

Photo: Stuart Hay

Dr Christopher Johnson, left and Dr Brian Molinari with the robotic arm.

Canberra suburb of Fyshwick. Ten of the arms have already been sold, while another 25 are in production. The Managing Director, Mr

modified as the robot's task changes. The next step will be to teach it to interact with its environment so that, for instance, if it is

point is that we are designing the knowledge system so that it can be marketed separately from the robot. This marks quite a shift from

1987: ANU Reporter article anouncing the Stanton-Molinari-Johnson grant to develop software to control a locally built robotic arm. Paul Mackerras and Thea Clark were employed on the project.

Another successful avenue for increasing resources for DCS was to seek funding for research collaborations. Apart from the KRIS project, discussed here, DCS was involved in a number of significant collaborations under the banners of CISR, ANU-Fujitsu CAP Project, and ACSys CRC. These are discussed in Chapter 6. The IDEA workshops were a further example of valuable collaboration.

3.8.1 The IDEA Workshops and Related Collaborations

Ron Morrison of the University of St Andrews has an impressive track record in athletics as well as in research. He is the current President of ScottishAthletics[6] and is a pioneer of orthogonal persistence in programming languages.

[6] https://www.scottishathletics.org.uk/staff/ron_morrison/

Ron spent several months of study leave at ANU in 1983 and returned with PhD students, such as Graham Kirby, Quintin Cutts, Fred Brown, Dave Munro, and Al Dearle, working with people such as Robin Stanton, Jeffery Yu, and Chris Johnson. He tells me: "I visited ANU on sabbatical in 1983 and returned in 87, 88, 89, 92, 94, 95, 97, 98, 99, 01 twice, 04, 06 and 08. I reckon I have spent a year of my life in Canberra and just a little less in Adelaide." [A persistent visitor!]

ANU people also attended research workshops in the wilds of Scotland. Steve Blackburn recalls attending one in Drymen and I believe there was one in Kinlochrannoch and another on the Isle of Arran. Eventually the research collaboration was expanded and formalised under the IDEA (Integrated Data Environments Australia) Project. See `https://cs.adelaide.edu.au/~idea/` for full details. Quoting from the Project Overview:

> The IDEA Project has been established by a grant from The Australian Department of Industry Science and Tourism [DIST] under its International Science and Technology Program. The principal researchers under the grant for the IDEA Project are Chris Barter, University of Adelaide, Robin Stanton, the Australian National University and John Rosenberg, Sydney University. The IDEA project aims to establish systematic interaction and collaborative research projects between Australian researchers and their European counterparts in the area of future object oriented databases and associated systems. The IDEA Project brings together a number of research groups in Australia and Europe with common interests, building on existing relationships which were previously not tightly knit, and to increase the scale of interaction, both within Australia, and between Australia and Europe.

Chris Barter says that David Smith of DIST was instrumental in getting the project off the ground. Several Americans (from UMass at Amherst and Purdue) were also involved in the IDEA workshops: Eliot Moss, who has maintained a long term collaboration with ANU, Dave Stemple; and Tony Hosking who is now Director, DCS.

Ron Morrison recalls extensive and productive collaborations initiated and promoted by ANU

In 1983 I spent six months sabbatical at ANU with my family, supported by a grant from the then Head of Department Richard Brent. I worked mostly with Robin Stanton, Brian Molinari, John Hurst and Chris Johnson. The staff introduced me to colleagues in other universities who invited me to give seminars. It was the beginning of many collaborations. At the end of the visit I was encouraged to apply for the Chair at the University of Adelaide and was offered the post. History shows I did not accept it but the foundation of collaboration and friendship with many new colleagues was well established.

I have made many of research visits to Australia to work with Chris Barter (Adelaide), John Rosenberg (Sydney/Monash etc), Robin Stanton (ANU), Chris Marlin (Flinders) and John Hurst (Monash). In 1992 I spent 3 months in Canberra working with Robin Stanton on multi-computer object stores. On the 1992 visit I brought Malcolm Atkinson (Glasgow), Al Dearle and Richard Connor to the ANU and at the request of Robin Stanton presented the EU funded Fide course on Orthogonal Persistence over two days.

On the reverse side I obtained UK EPSRC (Engineering and Physical Sciences Research Council) Senior Visiting Fellowships for John Hurst (ANU), Chris Barter (Adelaide), Chris Marlin (Adelaide), and John Rosenberg (Newcastle, NSW) to work in St Andrews for periods of six months to one year.

The idea of collaboration grew in this group, and others like Malcolm Atkinson (Glasgow), Dave Stemple and Eliot Moss (UMass), and Brian Warboys (Manchester) joined the party in visiting Australia and working there. Al Dearle and Dave Munro from St Andrews relocated to Adelaide and were in on the birth of the IDEA workshops. Perhaps the crowning glory was the joint IDEA/EU Pastel workshop held in Freemantle in 1998, WA. The EU attendees were Antonio Albano (Pisa), Malcolm Atkinson (Glasgow), Sophie Cluet (INRIA-Rocquencourt), Al Dearle (Stirling), Alfons Kemper (Passau), Jochen Liedtke (GMD), Carmelo Malta (O2 Technology), Fausto Rabitti (CNUCE/CNR), Hans-Jörg Schek (ETH), Joachim Schnidt (Hamburg), Marc Scholl (Konstanz), Dag Sjøberg (Oslo). It was the first time the EU had funded a workshop outside its geographical boundaries.

The strength of what was done in this period was the collaboration. That began at ANU and ANU was the centre of the activity.

Seven IDEA workshops were held between 1995 and 2000 – in Coogee, Melbourne, Lamington, Magnetic Island, Fremantle, Rutherglen (Victoria, not Scotland) and Victor Harbour. I'm told that the Rutherglen workshop featured a magnificent banquet and a spectacular electrical storm, but rather less magnificent accommodation. Steve Blackburn remembers the electrical storm being matched by a feisty debate with Ron Morrison over his paper.

The IDEA workshops were very successful in promoting collaboration. They supported a lot of PhDs and were highly valued by graduate students. PhD students were given a chance of 5–10 minutes to describe their work and get feedback from other students, discussions on how to write papers and techniques on how to write grant proposals. Papers were deliberately not refereed or published.

Steve Blackburn recalls thanking Ron Morrison, while driving with Ron through the wilds of Scotland, for generous support during his PhD. Ron replied, "Just make sure you do the same for your own students."

Drymen, Scotland. ***Photo: David Hawking***

3.9 An Important Teaching Collaboration

Alistair Rendell recounts the story of an important collaboration with Shandong University at Weihai:

> One of the first initiatives in the 2014 growth plan was the application for a Chinese Ministry of Education (MoE) joint program with Shandong University, Weihai. The Dean at the time, John Hosking, and I put in an application for this at the end of 2014. It sought approval for 100 places per year at Weihai. With this came a commitment by the School to deliver material into the program at Weihai with the option for students from Weihai to come to ANU for their last two years of study and be awarded degrees from both Shandong University and ANU. This was new territory for the School and indeed for the ANU as we had the first MoE joint program at ANU (Psychology was second). Since then there have been many visits by staff to Weihai and vice versa. To my knowledge there have been 30–50 students transferring to ANU each year. I enjoyed many visits to Weihai and built a reputation for being the crazy academic who would go swimming at the beach there. DCS academics who were involved include Uwe Zimmer, Ramesh Sankaranarayana, Qing Wang, and Ben Swift.

Mengyang Zhang, a Shandong University '2 + 2' student (two years at Weihai and two at ANU) reports very favourable experiences in a CECS news item.[7]

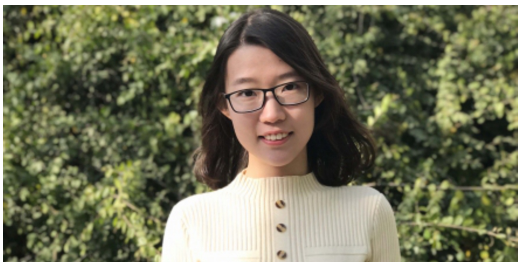

Mengyang Zhang from Shandong University, winner of the 2017 Paul Thistlewaite Scholarship. ***Photo: CECS***

[7] `https://cecs.anu.edu.au/study/meet-our-students/mengyan-zhang`

Chapter 4

DCS: A Cast of Characters

In the course of my research, I've collected many fascinating stories about the people who have contributed to DCS under its various organisational names. This chapter records those stories that haven't been told elsewhere in the book. Of course, this chapter is incomplete. I'm not aware of everyone who has contributed to DCS, I haven't been able to talk to everyone who I am aware of, and some people have declined to contribute, or not done so in time. I could have repeated information from profile pages on the web, but you can read those for yourself!

I apologise that entries are disjoint and listed in alphabetic order – I wasn't able to organise this chapter into a coherent story. I think there's enough interesting material in here to justify reading all of it, but I understand if you only read the entries of your friends, ... and enemies! (There's a small section at the end listing some people who left DCS decades ago and who have gone on to senior appointments overseas.)

2018: On the occasion of Andrew Tridgell's honorary doctorate. Back: Brian Molinari, Paul Mackerras, Malcolm Newey, John Zigman. Front: Steve Blackburn, Andrew Tridgell, Robin Stanton. *Photo: CECS Marketing.*

Peter Bailey completed his honours degree at ANU in 1991. Later, he worked part-time as a programmer on the ANU-Fujitsu CAP Project, reporting to me, while completing his PhD on parallel ML (ParaML). His supervisors were Malcolm Newey and Robin Stanton. During his PhD he worked at the Edinburgh Parallel Computing Centre, afterward at Object Technology International. He returned to DCS as a Research Scientist from 1998 to 2000, seconded to the ACSys TAR and WAR projects. Together, he, Nick Craswell and I turned our research in Information Retrieval into an enterprise search product known initially as S@NITY, then P@NOPTIC, and finally Funnelback.

In 2000–2001, Peter worked for NUIX Pty Ltd, and then became Canberra Office Manager for

Synop Pty Ltd which, for a while, was a reseller of P@NOPTIC after it was commercialised by CSIRO. In 2005, Peter joined my information retrieval research group in CSIRO. When Funnelback was spun off from CSIRO and the focus of the remaining group changed, Peter and I switched roles, and he became leader of an expanded group. In 2008 when CSIRO seemed to lose interest in researching search technology, Peter moved to Bellevue in Washington State to work on Microsoft's Bing search engine. After five years he expressed a desire to return to Canberra, and suggested that Microsoft might recruit me as part of a small Bing team in Canberra. We worked together on Bing until Peter transferred to a different part of Microsoft, shortly before I retired.

Amanda Barnard was previously Chief Research Scientist at CSIRO Data61. Her CV is impressive, filled with prestigious awards and senior appointments.[1] Her discovery of electrostatic properties of diamond nanoparticles has led to improved chemotherapy for brain cancer and to a Feynman Prize. She has taken leadership roles in the Australian supercomputing community, including at NCI and the Pawsey Supercomputing Centre.

At DCS (and the Energy Change Institute) she is excited about the opportunity to combine cutting-edge computer science methods and technologies to transform the way we approach scientific research in other domains. She hopes she will be able to influence how researchers in physics, chemistry, materials science, nanotechnology, biology and earth and climate sciences engage with technology to accelerate their impact. She also believes that creating effective relationships between the university sector and industry is key to positive change.

A Cecily Oakley interview[2] tells the story of her career up to 2010. The panel below follows up on some of the issues raised in that interview:

Amanda Barnard responds to some questions following up on the Cecily Oakley interview.

Dave: Since my focus is on computing, I'd like to know more about the computing side of your simulation work. First, do you use packages? (If so which ones?)

Amanda: Yes. Dozens of them, and many are domain specific.

Dave: If you write your own code, what do you write in?

Amanda: Yes. We have published code written in Python, C# , C++ and FORTRAN.

Dave: What computational methods do you use. E.g. finite element methods.

Amanda: Whatever the problem demands. Historically quantum chemistry, electronic structure methods, classical force field potential and mesoscale particle based simulations. Most of my work is applied machine learning so currently my time is spent with tensorflow, keras and scikit-learn.

Dave: Are there any particular challenges in the computational work.

Amanda: Mostly trying to fit "square pegs in round holes" — using methods designed for one thing and applying them to something else.

Dave: Have the recent big upgrades at NCI and Pawsey meant big jumps in what's possible? Or are they just about accommodating more users?

Amanda: A bit of both. Through the ALCGs (Australasian Leadership Computing Grants) we supported large scale projects that would not have been possible before, and the Pawsey PACER (Pawsey Supercomputing Centre for Extreme scale Readiness) scheme is similarly aligned.

Dave: What computer science training (if any) do you think computational scientists need?

Amanda: Numerical methods, C++ and Python, parallel programming, basics and software development (version control, etc) and high performance computing (including profiling, etc).

Dave: Would you now give the same answer to the last question in the interview?

> *I think you need a thick skin. I think, more and more, science is becoming very competitive. We are all competing for the same pages in journals and the same funding, so we need to learn to have a thick skin. It is also working very, very hard. I think when I was young I had the impression that an academic or scientific life was not too stressful; it's incredibly stressful. So I think also some advice would be to prepare yourself to work very, very hard and be under a lot of competitive stress.*

Amanda: Yes.

[1] https://cs.anu.edu.au/people/amanda-barnard

[2] https://www.science.org.au/learning/general-audience/history/interviews-australian-scientists/dr-amanda-barnard-computational

A more recent interview by Natarlie Kierce and Karen Phelan on the occasion of Amanda receiving an honorary Doctorate at RMIT University[3] included the following:

> For me the only thing more exciting than solving a complicated scientific problem is seeing the results put into action.

> Encouraging minority groups, including women, to join us in STEMM will better equip us to face these [physical, social, environmental, and medical] challenges. But to do this responsibly we need to ensure that a STEMM career does not expose our minorities to disadvantages, such as low wages, high rates of unemployment and a lack of job security. We need to ensure that a STEMM career is safe and sustainable, as well as exciting.

Nick Barnes is a Professor in DCS, having very recently transferred from Engineering. Prior to joining NICTA Canberra in 2003, he completed a PhD at the University of Melbourne, worked for Accenture and as a CS lecturer at the University of Melbourne, and spent 6 months at the University of Genoa, Italy. His research interests lie in computer and robotic vision, particularly bio-inspired vision, and he has worked extensively in the field of prosthetic vision.

At NICTA, Nick had the opportunity to work with the John Parker, who had 13 years experience at Cochlear, including as CTO. (John is now CEO at Saluda Medical, a company spun off from NICTA in 2010, which sells implants for controlling chronic pain.) In 2009 a bid was put together by NICTA, University of Melbourne, UNSW, the Bionics Institute and the Centre for Eye Research Australia for an ARC Special Research Initiative on Implantables, particularly the Bionic Eye. Funding of $42M was granted in 2010, and an additional $9M was eventually allocated to extend the project.

Nick described to me the appropriately rigorous device approval process which must be gone through before surgical implantation in live humans can occur. Surgeons and engineers work closely and must have a high degree of trust to ensure the safety of an implantable device. Nick led the project in Canberra, and some implantee clinical mobility trials were held here. Ethics approvals came through ANU and there was strong collaboration with the Centre for Eye Research Australia (CERA) and the Bionics Institute in Melbourne.

Nick says that a large number of prototypes were built. Testing was designed to ensure that the devices were safe, and devices were first implanted in animals and human cadavers. Even trivial changes in design required testing to start again from the beginning. Eventually, in 2012, devices were implanted in three human volunteers, each with 20 electrodes stimulating the optic nerve. Psychophysical experiments showed that subjects were only able to distinguish a very small number of signal levels on each electrode, and that the number of levels varied considerably between subjects.

Because of the very low information transmission potential (20 × 3 bits) Nick concluded that active vision processing would be essential to distill the very high information content from cameras and other sensors into the most useful signals to the electrodes. He also concluded there was no realistic prospect of restoring normal human vision with current implant technologies, and that the focus of the project should be on improving mobility. Accordingly, a head-mounted depth sensor was used to feed information into a specialised depth rendering algorithm which causes ground-based obstacles to "pop out". This proved successful.

Between 2012 and 2014, significant improvements were made and evaluated in clinical trials. In 2017 the Bionic Vision Technology (BVT) company was spun off with USD18M from Chinese investors. Four more volunteers were implanted. BVT's website[4] reports that, " The Company has recently completed initial human testing in seven patients demonstrating safety and efficacy with positive outcomes in patient mobility and device stability/durability."

Nick continues to work with BVT, but his current research direction is on fundamental computer vision, still with a biological interest.

[3] https://www.rmit.edu.au/news/all-news/2020/may/hon-doc-amanda-barnard

[4] https://bionicvis.com/about-us/

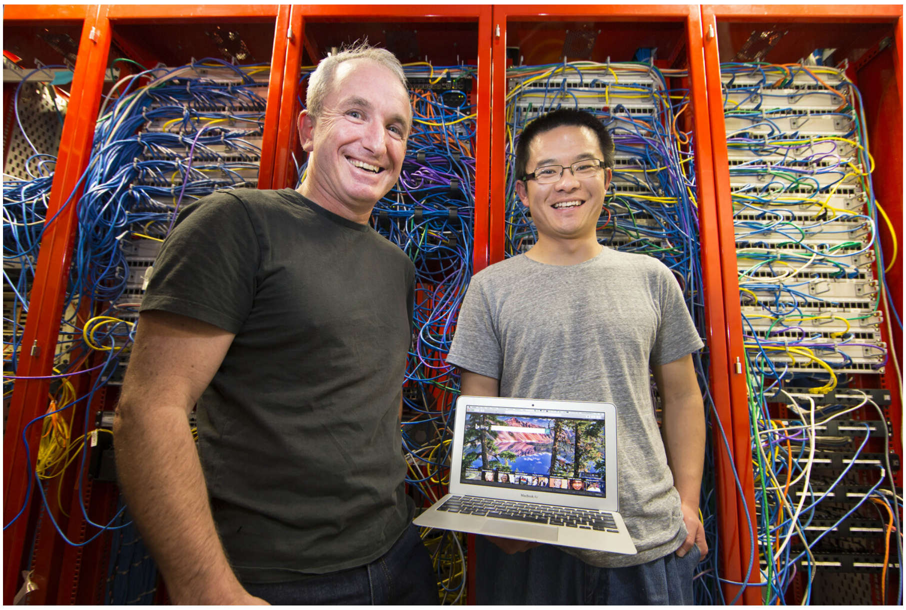

2016: Steve Blackburn with PhD student Xi Yang, who was working on a project with Kathryn McKinley (Microsoft Research (MSR)) and Steve on improving tail latency in Bing (following on from an internship at MSR). *Photo: Stuart Hay, ANU Photographic Services.*

Steve Blackburn describes his unusual career path.

I arrived at ANU in Feb 1992, following Anneke, who had a scholarship at JCSMR. Not being able to get a PhD scholarship, I got a job as a tutor, alongside Ross Parker. I started my PhD in July 1994 and finished in July 1997. My examination was plagued by the disappearance of Paul Wilson, one of my examiners (he went completely AWOL, sadly never returning to academia). It was nearly a year before my thesis was examined.

After I graduated I went to the ACSys CRC where I ran the Upside project which built an orthogonally persistent Java system. It was a huge amount of fun. There was a lot of energy and some great people. Alonso Marquez, Dave Walsh, Sam Taylor[5], and Gavin Mercer all worked on the project with me. I particularly enjoyed working with Alonzo, who was incredibly creative and energetic. He took over in 1999 when I left for UMass in Amherst, MA. [I visited Steve and Anneke in Amherst and shared the unusual experience of throwing a frisbee in the snow in front of their house.]

Darrell Williamson had offered me a position on my return, but it lapsed (to my surprise). I interviewed at UQ with Robin Stanton as a referee. Robin was wondering why on earth I was not going to ANU. I got the level C offer, then ANU matched using some money Robin scraped together that was residual money from the CRC, so I headed back to ANU, which is what both Anneke and I really wanted. One day Robin's hollow log (very) abruptly ran out, again surprising me (and apparently surprising Chris Johnson who was Head and the bearer of bad news).

By a stroke of good luck, Kathryn McKinley (who I worked with in Amherst) was visiting Intel on pretty much that very day – incidentally, the very same day that the Chemistry Building had its fire that caused parts of Civic to be evacuated. Kathryn was telling me (at some length) how Intel was trying to hire her to run a project, but she wasn't interested, but thought it would fit me perfectly – all the while I watched more and more fire trucks head to the Chemistry building and smoke get thicker and thicker. The phone call ended when we had to evacuate the building. It was a key call and very memorable for the drama outside.

I then moved to Intel, though I worked from Canberra. I ran a small team of engineers based in Russia, while my boss was in Hillsboro, Oregon. I did that for a couple of years,

[5] Now Senior Director, IoT Software Strategy at Arm, Cambridge, England

and then two things happened – a large 5-year ARC proposal with fellowship, that I'd put in right before my money ran out, was funded, and a major re-org at Intel (due to AMD's resurgence) made my job there far less promising. So I came back and was a research fellow for five years. At the end of that I transitioned into a regular teaching position within DCS, teaching COMP1110 which was and is one of the two large intro classes. I'm still doing that today. My promotion to professor was in 2012.

I think the great Linux / BSD battle captures a number of interesting aspects what we were up to, the enthusiasms and competitive nature of our friendships. I fondly remember many years of lunches in the union with a large gang of people which always finished with Brian Corrie wandering up to the union shop and coming back with a packet of two Reece's Peanut Butter Cups, which apparently kept him in touch with North America. My recollection is that Brian introduced us to ultimate frisbee while we were in the Crawford building and we played on the lawn in front of Chifley. One time Leon Smith stumbled upon us and showed us how to play properly and before too long ultimate became established in Canberra ☺

Steve Blackburn recalls a fire in the CS&IT computer room

Then there was the time that one of my computers caught fire, causing the CSIT building to be evacuated. Well, it did not quite catch fire, it just blew a capacitor which generated a large amount of smoke in the old machine room, causing the fire brigade to come. The machine was an old AMD box that I'd used for many years (at least 5, perhaps 10) for 24/7 regression testing for the Jikes RVM project. So it had been powered on and running at full speed continuously for a very long time. One of the technical staff saw what was happening and called the fire brigade, and then went to lunch. Conversely, I came back from lunch to considerable drama, and one VERY grumpy fire fighter who did not know where the source of the smoke was because, "the idiot who reported it has run off." They found it using breathing gear and an infrared camera. I was dismayed when I realised that the source of all the drama was one of my old AMD boxes. Bob Edwards kept the culprit motherboard for years.

Clive Boughton's account of 2020 modifications to eVACS electronic voting software

Our eVACS software has been in use at Elections ACT since 2001. It has been used in every election (including casual vacancies) since then.

This year, mainly due to greater security requirements, but also due to feedback from users/voters, the system has been upgraded. To vote electronically, larger touchscreens are now used. Most of the voting software has been rewritten in Ada 2012 (for security reasons) and the counting parts have been rewritten using stored procedures (again for greater security). Other things, such as randomised vote order have been improved to further help in the prevention of vote-voter connection. The entire system is deployed using server to server network connection (only on boot). So setup requires far less effort than previously. Voting servers are connected to clients via public/private key pairs, and vote transfer is achieved using fully encrypted USB devices. Blind and vision-impaired voters are still supported using the same system, but with the usual telephone-style keypads. It's a very different system, and it's been a lot of hard work pulling it all together.

BTW, all paper votes are now scanned-in. This has been the case since 2008.

Clive Boughton, like Henry Gardner and Alistair Rendell, has a background in the "hard sciences", completing a physics PhD at ANU in 1987 under the supervision of Bob Watts. Unlike Henry and Alistair, after graduation, Clive pursued a career in software engineering and project management. As an employee of C3 Pty Ltd, Pi Systems, and Ferranti Computer Systems, and through his own company *Software Improvements* he has completed a large number of software projects in government, defence, and private industry. The projects have focused on successful project completion, software quality and validation/verification of software, particularly of mission-critical and safety-critical software. Many of these projects have been outside Australia.

Clive's knowledge of the theory and practice of software engineering, for example model driven design, led to him working in DCS, first as a Visitor and then as a full-time Associate Professor, teaching Software Engineering units, designing (and having accredited) software engineering bachelors and masters courses, and supervising PhD students. His lengthy answer to my Dorothy Dix question on why professional programmers need to study software engineering rather than just computer science appears on Page 159.

Clive's company, Software Improvements, won the tender to develop the electronic voting system for Elections ACT. Their eVACS software is an example of a non-safety-critical system which must nevertheless be free of bugs.

Peter Christen is a professor in DCS, currently transitioning to Emeritus status, and will (COVID-19 permitting) soon take up a prestigious six-month Leibniz Professor fellowship at the University of Leipzig in Germany. His interests are in data mining and data matching, and he is looking forward to further collaboration with Erhard Rahm in Leipzig, and hopefully with UK colleagues such David Hand at Imperial College London, Chris Dibben at the University of Edinburgh, and Al Dearle at the University of St Andrews (a previous Dean of Science there), who visited the ANU several decades ago.

Peter completed his undergraduate degree at ETH, Zürich where he was amazed by the over-the-top methods used by Niklaus Wirth (famous for developing Pascal and Modula-2) to maintain silence in his lectures. After completing a PhD in Basel, he attended the 1998 International Conference on Supercomputing (ICS) in Melbourne, where he presented a paper in the same session as Markus Hegland, who then invited him to visit ANU. A one-year postdoc at ANU, funded by the Swiss National Science Foundation, was followed by a three-month contract at ACSys CRC (organised by Simon Hawkins and Graham Williams), leading into his appointment as a lecturer in DCS at the beginning of 2001.

Peter's lectureship was 50% funded by the ANU-Fujitsu CAP Project, and his expectation was that he would spend the first semester of his appointment as a full time researcher, and take on a teaching load in second semester. This expectation was rudely contradicted when a student asked him about the textbook for COMP2300 (Computer Systems), starting in a week's time! A dislocation had been caused when the previous lecturer, Lex Weaver, suddenly departed to join a security agency. Bob Edwards assisted with a frantic search through backups to find lecture slides and other course materials to enable Peter to teach the course on short notice.

Peter has felt the constant pressure of teaching loads in DCS in the two decades since, but has managed to establish a strong research record nonetheless. In coursework teaching he contributed with the development of the first *Data Mining* course in 2007, and then redeveloping *Programming for Scientists* together with Alistair Rendell from 2012 onwards to a very large course with over 300 students. Peter also contributed significantly to HDR supervision, taking on roles such as Graduate Convenor since 2008, and Associate Dean (HDR) from 2009 to 2011. His contribution to education at ANU was rewarded with several awards, notably the 2014 Vice-Chancellor's Award for Excellence in Supervision, and in 2013 Dean's Awards for both Teaching Excellence and for Excellence in Supervision.

Peter will finish his full-time teaching career at ANU with a bang in March 2021 by having taught (since July 2020) three versions of the course *Data Wrangling* to nearly 400 students, including to over 100 students from the new Graduate Certificate of Data Engineering. He remembers fondly a 2016 sabbatical at the University of Cambridge, and his involvement with the Isaac Newton Institute, and the Administrative Data Research UK initiative. He is looking forward to a long break from teaching.

Nick Craswell was a PhD student in DCS from 1996–2000, initially supervised by Paul Thistlewaite. I took over Paul's role when he died in early 1999. I had no prior supervisory experience but fortunately Nick was an able student and knew more than I did about the Web. Earlier, Nick attended the WWW7 conference in Brisbane in 1998 (chaired by Paul and Helen Ashman) and met two PhD students from Stanford, Larry Page and Sergei Brin. He still has their business cards (with "Google" written on them in case he forgot who they were!) This was at the stage where they said their main goal was to be a research system.

During his PhD, Nick spent an internship at the Alta Vista lab in Brisbane and another with Stephen Robertson at Microsoft Research in Cambridge. Prior to his PhD, Nick did a combined BEc/BSc (hons) degree at ANU. He says, "I was really interested in systems that involve humans, and how to make them work well for said humans. So Economics was one area, but I really got into CS more. For my PhD topic it was clear that IR (Information Retrieval) is a very important and impactful area where a system (the Web and search engines) meets humans."

In the early days of his PhD, Nick used his advanced knowledge of Excel to advantage in tuning the parameters of the PADRE retrieval system running on the Fujitsu AP1000. How many other Excel users know that Excel supports Integer Programming, and that it can be used to drive programs running on an advanced supercomputer? In his roles at Microsoft Research and Bing, Nick has continued to tune more and more IR parameters based on human interaction data.

Nick recalls that some Economics friends pulled off a graduation ceremony stunt in which one graduand handed the Chancellor a large bag labelled with a dollar sign in return for his testamur, while a carefully positioned co-conspirator took the "incriminating" photograph. He also claims that he, Jason Haines and Geoff Keating accidentally[6] won a national programming competition. This was reported in the FEIT Annual Report for 1995:

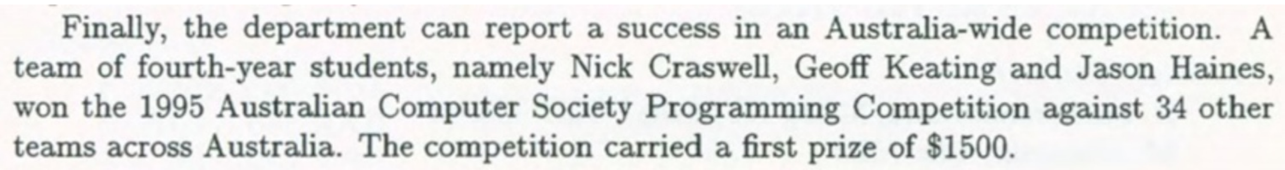

Finally, the department can report a success in an Australia-wide competition. A team of fourth-year students, namely Nick Craswell, Geoff Keating and Jason Haines, won the 1995 Australian Computer Society Programming Competition against 34 other teams across Australia. The competition carried a first prize of $1500.

Bob Edwards was recruited (by me) to the DCS Programming Group in the early 1990s – One of my best appointments! He came from UNSW, with a can-do attitude, and possessed strong skills in both hardware and software. Over the nearly thirty years since, he has built hardware gadgets, wired networks, set up computer labs, administered systems, designed furniture, written software and taught five different computer science courses. (See `http://users.cecs.anu.edu.au/~Bob.Edwards/`.)

You will find many references to Bob's contributions throughout this history. One of the more interesting achievements is building the Bunyip (see Page 218) and winning a Gordon Bell award.

To quote Mike Carden, former Masters student, "Another legend of DCS is Bob Edwards... a luminary of Open Source, a great teacher, and a builder of community."

Shayne Flint was a part-time PhD student in DCS between 1998 and 2006 and was a member of academic staff between 2007 and 2018. He previously studied engineering at RMIT and graduated in 1979, working with the RAAF until 1986 and then for defence contractor C3 Pty Ltd. At C3 he met Clive Boughton and, when C3 folded, the two became founders of the company which became Software Improvements, which did consulting work and distributed an Ada compiler. After leaving Software Improvements in 1996 Shayne formed another company which developed and sold a system which allowed Ada developers to easily access Java APIs. Shayne says:

> While working on the Ada compiler technology I continued my work on improving software engineering practice. This led to conversations with Clive, Brian Molinari and Paul Thistlewaite about continuing the work as a PhD candidate. I started my PhD in 1998. The initial topic was around software process improvement, but quickly shifted to model-driven engineering within a broad context of systems engineering, systems thinking, complexity etc. Clive was my primary supervisor, and I worked on it mostly on a part-time basis. During my PhD, my work attracted the attention of DSTO. I was subsequently employed by DSTO where I worked on my PhD (and other things) for a year or so before returning to ANU. I eventually completed the PhD in 2006.
>
> From 2003 I was employed on and off to teach specific courses, before being employed full-time between 2007 and 2018. During my time as an academic, my focus was on industrially relevant education and practice (a natural extension of my work in the military and industry). Underpinning this was the following vision:
>
> "ANU graduates will live and work in a world of rapid technological, economic, environmental and social change. My vision regarding education is to ensure that students from all disciplines are able to graduate with the knowledge, skills and confidence required to build successful careers in this uncertain, but exciting world."

[6] Nick explains that unlike some other teams, the winning ANU combination was self-organised rather than carefully selected, lacked a coach, and didn't practise much.

> My vision led to a series of initiatives that resulted in a BSEng program that produced graduates with a much broader understanding of the discipline and practice than just "programming and project management".

2011: Henry Gardner (R) with PhD student Torben Sko (C) and Tim Flannery (L, Australian of the Year 2007). Henry and Torben had a grant to develop an interactive computer graphics exhibit for Science Week. The exhibit, *Who Could You Be?*, featured a 3D Tim Flannery avatar to engage and inspire young people about careers in STEM. ***Photo: Stuart Hay, ANU Photographic Services.***

Henry Gardner started academic life as a nuclear physicist, with no particular interest in computing. However, after spending most of 1980 hitch-hiking around Europe, he completed a graduate diploma in computing from the University of Melbourne, and that sowed a seed. While working on his PhD at ANU's Department of Theoretical Physics under Bob Dewar, Henry became one of the first Australians to use a Cray supercomputer. From a teletype at ANU he dialled up the Cray at NERSC (National Energy Research Scientific Computing Center) in Berkeley, CA, working on modelling the iron core of ANU's tokamak. This mode of working was stopped after hidden telecommunications charges caused the exhaustion of the expendable research materials budget of the Department of Theoretical Physics.

Henry then worked on stellarators and used the RSPhysS VAX for modelling the ANU "Small Heliac" (SHELIA). He took up a postdoc position in Germany and participated in the design of the Wendelstein 7-X stellarator around the time of the fall of the Berlin wall. The Wendelstein 7-X is a massive project which has taken decades to bring to fruition. Henry says that funding was obtained by agreeing to move the project to Greifswald in the former East Germany, and thus become eligible for funds associated with East German reconstruction.

In 1992 a short submission, drafted initially by Henry, succeeded in attracting $8.4M for a Major National Research Facility from the Keating government for an upgraded version of the ANU heliac. Unfortunately this did not provide funding for Henry's position and he then started working with ANUSF and DCS.

ANUSF's role in supporting One Australia's America's Cup bid (see Page 207) resulted in funds coming to ANU, which were used to create a Computational Science and Engineering Program at ANU. Henry and Ben Evans set up a lab, comprising 20 Silicon Graphics workstations,[7] in the CS&IT building. Henry says that, having got the course organised, he didn't know what should be taught!

Viewing the 12.4 metre high, Australian-constructed Jeff Koons Puppy (supporting 55 tonnes of soil and plants) on a visit to the Museum of Contemporary Art in Sydney provided the inspiration!

[7] I used it once for a large-scale text retrieval experiment.

Students used the *Strand* finite element software package to create and visualise stainless steel puppy structures. Henry had a paper published on *the Puppy Paradigm* and won an international teaching award. Following its exhibition in Sydney, the Puppy support structure was dismantled and sold to the Guggenheim Museum in Bilbao, Spain.

A visit to the CAVE at Argonne National Laboratory in Illinois provided the inspiration for Henry's next move – into virtual reality. Henry and Rod Boswell felt that the critical element of the CAVE was the corner. Henry returned to ANU to supervise the creation of ANU's Wedge, located in the Huxley building. The three critical elements for success were: frame-accurate synchronisation of the two projectors; achieving a seamless join at the corner; and achieving correct perspective along the walls. Rod's technician, Peter Alexander, put the wedge together physically, Intergraph provided a computer which solved the synchronisation issue, and Drew Whitehouse's P-Space software solved the software perspective problems and enabled applications to be built for the theatre.

The Wedge was very successful. Chief Minister Kate Carnell launched the first one and ANU sold Wedges to ADFA, CSIRO Discovery Centre, and the Powerhouse Museum. Running application software written by Drew Whitehouse, the Discovery Centre and Powerhouse exhibits were among their most popular. A portable version of the Wedge toured the country including being exhibited at an international physics workshop on the plasma physics topic of "magnetic islands" which, coincidentally, took place on Magnetic Island in Queensland.

Together with Chris Johnson, Darrell Williamson, and Rod Boswell and Henry applied for a teaching grant in collaboration with RMIT University, through a scheme known as the "Science Lectureships Initiative." They were successful with a $1M grant in 2000 to develop courses in both universities with visualisation and the Wedge as a major focus. Henry claims to have established the world's first graduate eScience program.

The Jeff Koons floral puppy, constructed in Sydney, now in Bilbao, Spain. *Photo: David Hawking, 2015*

Oh, and did I mention the funny ukelele songs for software engineering?[8]

Henry was acting Head of DCS in 2004 and returned from 2008 – 2013, becoming the first director of the new Research School of Computer Science. This marked the integration of DCS and the Computer Sciences Laboratory, a merger which was not without pain. Henry says that he paid full attention to academic workloads and the need for fairness. There were vociferous criticisms from some senior staff, and the future of computer science was to some extent in doubt, but student evaluations showed a significant rise in satisfaction over the 2008 – 2012 period.

[8] See Henry's homepage: `https://users.cecs.anu.edu.au/~Henry.Gardner/`.

After the dot com crash around 2000, enrolments in computer science declined dramatically. However, from the mid 2000s enrolments grew dramatically, causing considerable difficulties in finding staff to teach the courses. In 2019, 72% of CS enrolments were international students.

Tom Gedeon came to ANU in February 2003, on long service leave from Murdoch University, as a NICTA visitor (possibly the first visitor to the NICTA Canberra Lab). His wife was returning to a post in the Attorney-General's Department. They arrived in Canberra immediately after the 2003 Canberra bushfires in which around 500 houses were destroyed. Unsurprisingly, they found accommodation hard to find.

Tom spent the first six years of his life in Hungary, then with his family followed his father's United Nations Development Program postings to Nigeria (until war broke out), Ivory Coast (very briefly) and Guyana (for about five years), before migrating to Australia and completing high school in Perth. He studied medicine for three years and one week, but converted to computer science, appreciating the willingness of a Faculty Secretary to flexibly interpret the rules.

In October 2003, Tom took up the position of Professor and Head of DCS, but soon afterward, in August 2004, was asked to take over the position of Interim Dean of FEIT. The Deanship was vacant due to the departure of John Baird. John Richards had been appointed as his replacement but was not available until his term as Deputy Vice-Chancellor ended. This was during the period when the Vice-Chancellor Ian Chubb was trying to fully integrate the Institute of Advanced Studies and The Faculties. FEIT and RSISE were initially combined as the Institute of Engineering and Computer Science, eventually becoming CECS. When this was first in the offing some thought that CECS would be pronounced SEX rather than KEKS but the latter has largely prevailed, despite some hold-outs.

Once John Richards took up his post as Dean, Tom became Deputy Dean, focusing on teaching matters. He stayed in this role until early 2005, and returned to DCS as an academic member of staff, having earlier expressed his preference not to continue in an administrative role.

In 2008, Tom created a Human Centred Computing (HCC) group within DCS. Also in 2008 he was a consultant to the Australian Bureau of Statistics (ABS) on the development of Field-Of-Research (FOR) codes. Over the years he has taught courses in Web Development and Design and in neural nets. While Head of DCS, he helped start a new degree, BA (New Media Arts) which required two computer science courses, one of which was Art and Interaction in New Media.

From 2010 to 2012, Tom was President of CORE, the Computing Research and Education Association of Australasia. One of the issues CORE faced was to persuade the RQF (Research Quality Framework, later ERA, Excellence in Research for Australia) to accept conference papers as valid outputs when rating computer science departments. By the 2015 ERA exercise computer science ratings were citation based, and conference papers were included. Tom organised the 2016 Australian Computer Science Week at ANU.

Stephen Gould came to ANU in November 2010, soon after completing his PhD. He is a professor and has been awarded a 2020 ARC Future Fellowship. His background is in electrical engineering and his interests lie in computer vision and machine learning. He has both a Masters and a PhD from Stanford University and worked at Amazon in Seattle, on Amazon's automated convenience stores (Amazon Go) and on Amazon Web Services (AWS). He's very keen to see research translated into practice and taught the *Craft of Computing* course. Now he teaches *Advanced Machine Learning*.

Prior to ANU, Stephen has worked in several interesting start-ups:

LGC Wireless: (In about 1998.) The company made mini base station repeaters for the cell phone network. These were capable of improving mobile reception in places like hospitals (thus reducing the power output from mobile devices), and train stations. The most enthusiastic customers were actually mobile phone shops, where a repeater in the shop showed off all the handsets in the best possible light. Stephen worked on developing the firmware.

Polartechnics: A medical devices company which developed a screening device for cervical cancer to compete with Pap smears. Rather than removing tissue for post analysis in a laboratory, the Polartechnics probe scanned the tissues and used machine learning to detect pre-cancerous conditions. Launch of the device was substantially delayed by a change in business model, and by the time it was launched the accuracy of Pap smears had improved, leaving the device no longer competitive.

Dilithium Networks: High-performance on-the-fly transcoding of different video and audio codecs. In the course of the work, it was discovered that most vendors did not conform to the ITU (International Telecommunications Union) standards for codecs. It turned out that a compliance tester was the company's most valuable product.

Sensory Networks: Stephen was a co-founder of this network security start-up which developed an FPGA-based gadget like a `grep` chip for network traffic – it could do regular expression matching on high-rate network traffic flows. The company raised $20M in VC funding, peaked at around 70 employees, and was sold to Intel in 2013 for $30M.

Stephen and I had a good chat about the sort of CS training which might prepare a student for work in a start-up company, or even encourage them to found one. We agreed that there was no fixed formula or recipe for a successful start-up, and that many of the skills required are not part of a CS curriculum. However, exposing students to stories of start-ups is likely to be beneficial – there's a real excitement (along with countless setbacks) in working in the start-up world. ANU's *TechLauncher* program[9] and the ACT Government's *Canberra Innovation Network*[10] are valuable initiatives.

Stephen currently leads the ANU node of ACRV, the Australian Centre for Robotic Vision, having taken over from Rob Mahony, but it's coming to an end shortly. One project worked on the integration of robotic vision with natural language, E.g. "Robot, please go to the kitchen and make me a coffee." Another project partnered with a US company building a device to automatically assess embryo viability in IVF. ANU involvement was to track the embryos as they developed in the incubation chamber to identify individual cells and detect cell division events.

Within CECS, computer vision has been split between the research schools of Engineering and Computer Science, but the research has recently coalesced under the new School of Computing, aligning closely with machine learning, now a very important part of computer vision.

From 2021, Stephen will be working on Deep Declarative Networks, an innovative integration of classical models into deep learning models, supported by his recent Future Fellowship.

Richard Hartley is an Emeritus Distinguished Professor in DCS. He grew up in Perth and won a French Government scholarship to spend year 12 at a French language boarding school in New Caledonia. He has ever since maintained a keen interest in languages, and can speak French and German well, can read but not speak Danish, and learned to tell his children to go to bed in Czech, his wife's native tongue.

Toward the end of the year in Noumea, he decided he should see more of New Caledonia. Accordingly, he went down to the harbour and hitched rides on small boats. The first trip was to the Loyalty Islands. The next time, he asked for a ride on a boat going to collect guano from Walpole Island. They dropped him on the Isle of Pines (Kuto) and collected him on the way back, ten days later. He obviously struck up good relations with the boat captain, who later visited Richard's parents in Perth, and subsequently proposed a business venture from Venezuela which involved shipment of crocodiles, whales' teeth and slouch hats.

Richard studied mathematics at ANU between 1967 and 1970. While there he co-wrote his first, and least-cited publication: *The Scuna/Uncs Songbook*, a songbook of the ANU and UNSW Choral Societies. During vacations in Perth he gained access to the University of Western Australia's PDP-6 computer[11] and worked on plotting graphs on the CRT (cathode ray tube) screen. He was impressed that you could interact with the plots using a light pen.

After he completed a PhD in knot theory at the University of Toronto, he made use of a walk-up computing facility at the University to develop and run programs relating to knot invariants. You prepared a deck of cards for a program in PL/C, fed it into a publicly accessible card-reader and

[9] `https://cs.anu.edu.au/TechLauncher/`

[10] `https://www.cmtedd.act.gov.au/economic-development/innovation-industry-investment/canberra-innovation-network`

[11] I assume the same one used by both Ray and Col Jarvis.

collected your output almost immediately. The PL/C system imposed a three second overall time limit on compilation plus execution. Richard was very impressed at how much could be done in such a short time. Wikipedia tells me that PL/C was a subset of Pl/1 developed at Cornell University. "The PL/C compiler had the unusual capability of never failing to compile any program, through the use of extensive automatic correction of many syntax errors and by converting any remaining syntax errors to output statements."

[Richard Brent wonders whether PL/C might be related to the apocryphal system that, when given the null program, will "correct" it to give a useful working program.]

After post-doc appointments in Frankfurt and Columbia University, Richard Hartley took a 3-year post (1979 – 1981) at the University of Melbourne, continuing to work on knot theory, writing programs in Pascal and making use of Richard Brent's multi-precision arithmetic package to compute knot invariants.

Next he worked at the University of Missouri, in St Louis where he immersed himself in computing. Intel had donated a rather unusual computer to the university which came with a faulty Ada cross-compiler written in Pascal. The computer was an iAPX 432, described as a *micromainframe*, and designed to be programmed entirely in high-level languages. The story of this machine[12] is little remembered but very interesting.

Richard set about making the compiler work, converting it to operate in ASCII rather than EBCDIC, and cursing the fact that all the variable names were entirely meaningless. Eventually the Ada compiler compiled but complained about uninitialised variables. He discovered that the Pascal compiler had chosen to initialise memory to `0xEF`. He didn't have the source for the Pascal compiler, but edited the binary to replace `0xEF` with `0x00`. Although he left St Louis before completing this unusual software engineering task, he came away with the conclusion that most seemingly intractable computing problems can eventually be solved.

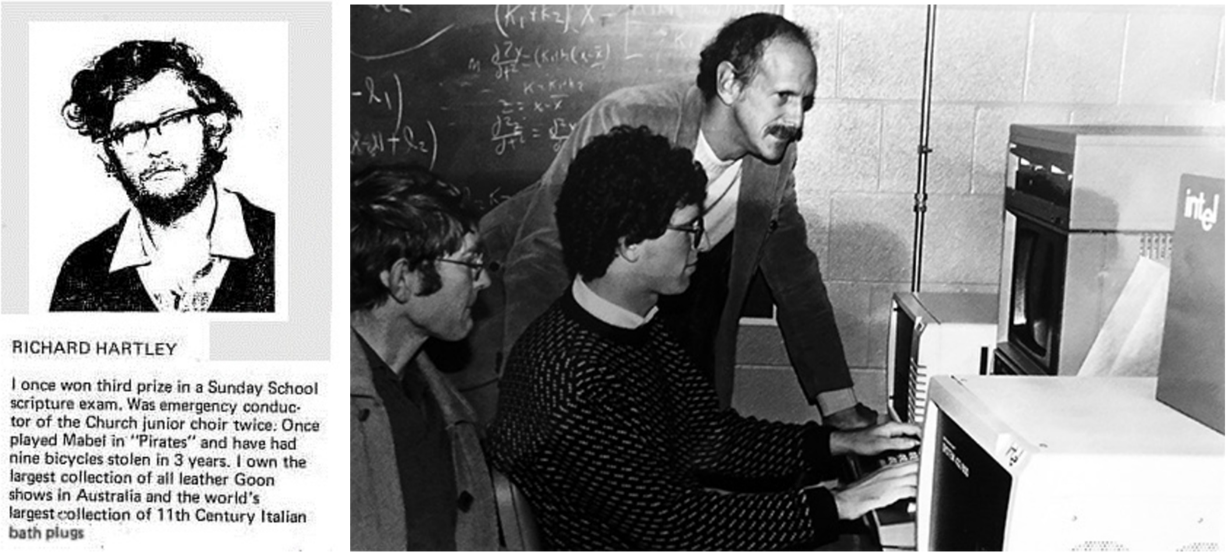

RICHARD HARTLEY

I once won third prize in a Sunday School scripture exam. Was emergency conductor of the Church junior choir twice. Once played Mabel in "Pirates" and have had nine bicycles stolen in 3 years. I own the largest collection of all leather Goon shows in Australia and the world's largest collection of 11th Century Italian bath plugs

Richard Hartley: L: Campaigning to be SRC president. R: Overseeing a Univesity of Missouri maths student operating the new Intel iAPX 432 machine. Richard is on the left. ***Photo: courtesy of Richard***

Next stop was Stanford University for a Masters in computing. While he was there Jeffrey Ullman published a book on *Computational Aspects of VLSI* and there was a focus on VLSI design. This led to a job using CAD to design VLSI with GE (General Electric) in Schenectady, NY. This location avoided the need to maintain a long-distance relationship. At GE he worked on a "silicon compiler." Eventually, GE decided that there was no point in doing VLSI design in-house, and Richard switched to working on computer vision, a field he has remained in until the present day. In academic terms this must be judged a success, since his Google h-index is almost 80 and one of his publications has

[12] `https://en.wikipedia.org/wiki/Intel_iAPX_432`

more than 29,000 citations.[13]

Around the time of the first Gulf War, he was ask to go to GE's aerospace division to try to improve the flight simulators used to train pilots. The flight simulators were from a bygone era and consisted of three cabinets of electronics. Rather than working on upgrading and improving the hardware, the on-site priority was to implement a capability to model real terrains (e.g. San Francisco airport) rather than artificial terrain. Richard's team spent time doing photogrammetry on aerial photographs, from this building realistic terrain models and filling in with realistic textures. Although so called Digital Elevation Data (DED) files already existed, they found a very serious error in one DED file, where a hill, obvious from the aerial photographs, was non-existent in the DED files. Maybe the acronym DED had hidden meaning.

At GEC, Richard also worked on projects for the FBI, including one on fingerprint analysis.

Richard told me of a number of memorable moments [so far] during his career:

- Attending a workshop in Lhasa, Tibet, organised by Harry Shum [my grand-boss when I joined Bing] and travelling there on the newly opened high-altitude railway.
- Listening to a talk on group theory in Frankfurt, delivered in German by Peter Hilton, a British mathematician who, as a first year undergraduate, was recruited to the code-breaking centre at Bletchley Park on the basis of his knowledge of European languages and his mathematical ability. Hilton has later said that taking Italy out of the war had been a ghastly mistake since, unlike the German ones, the Italian codes were easily broken over a pint at the local pub.
- Attending a course of lectures at ANU in algebraic topology by Max Newman,[14] whose work at Bletchley Park led to the development of Colossus, the world's first programmable electronic computer.

Richard was saddened by the decline of NICTA as it came to the end of its funding. At its peak, NICTA had substantial world-ranked groups in computer vision and in machine learning. He also regrets the demise of the former ANU Institute of Advanced Studies.

Tony Hosking, Director of the latest DCS incarnation, the new (2021) School of Computing. ***Photo: CECS Marketing***

Tony Hosking joined ANU in 2015 with a secondment into NICTA's Trustworthy Systems Group, having developed his academic career to that point in the United States where he had been on the

[13] Note: As at 24 March, 2021, Google Scholar had wrongly attributed this Hartley and Zisserman book, *Multiple view geometry in computer vision*, to someone else, and it failed to appear in Richard's Scholar profile.

[14] `https://en.wikipedia.org/wiki/Max_Newman`

faculty of Purdue University since 1995 following a PhD at the University of Massachusetts. He returned to Australia, partly to satisfy a desire to return home (29 years of Arctic winters had taken their toll), but also because the ANU/NICTA opportunity meant being able to work with the leading applied software verification team in the world led by Gernot Heiser at UNSW/NICTA and to co-locate with his long-time collaborator Steve Blackburn at ANU. The goal was to continue work initiated during a sabbatical at NICTA in 2013, on specifying and implementing a low-level managed language virtual machine, and formally verifying its key components including garbage collector and compiler. He stepped up as interim Director of the ANU Research School of Computer Science in 2019 and was confirmed in the substantive role later that year.

Tony is a co-author of the bible of garbage collection: *The Garbage Collection Handbook: The Art of Automatic Memory Management*. It seems likely that his current administrative role also involves dealing with a modicum of garbage.

Tony's undergraduate degree is in applied mathematics and computer science (BMathSci) from the University of Adelaide, but he completed his first year, in 1982, at the University of Waikato. Between Waikato and Adelaide he worked with Mark Apperley on FORTRAN mapping software – computing centroids of parcels of land from mapping data – and enjoyed a related visit to the New Zealand's government mapping agency to see their impressively large table plotters. After Adelaide, he contemplated doing an Honours year at ANU and vaguely recalls meeting with a certain Richard Brent over the summer of 1984. Instead, he returned to Waikato and enrolled in a mixed coursework/thesis Masters degree under E.V. Krishnamurthy (who later joined ANU's RSISE) and Mark Apperley. He worked on a range of projects at Waikato in computer graphics but focusing in his thesis on efficient interpretation of the lambda calculus.

Tony went on to study for a PhD at the University of Massachusetts at Amherst, with Eliot Moss as his principal advisor. As a research assistant he worked on a virtual machine for the Smalltalk-80 programming environment, based on the classic "blue book" specification of Adele Goldberg and David Robson: *Smalltalk-80: The Language and its Implementation.* which specified the virtual machine and byte-code compiler for Smalltalk in the Smalltalk language itself. Tony built an efficient implementation in C, which was state-of-the-art at the time, interpreting threaded code and instruction addresses rather than bytecodes, and making full use of the registers of the underlying hardware. Smalltalk is an integrated development environment and Tony spent most of a year implementing the graphics engine (or BitBlt bit-block transfer function) efficiently in C rather than assembler. This had the rewarding outcome of seeing the Smalltalk environment (cloned from the image with its roots in Xeroc Parc) finally come up with zippy graphics and satisfying graphical editor, debugger, and code browser. Along the way he also implemented the first garbage collection toolkit framework, which served not only UMass Smalltalk but also other nascent programming languages such as Cecil (U. Washington) and GNU Modula-3 (U. Mass.).

The UMass Smalltalk implementation was the platform for his PhD thesis, *Lightweight support for fine-grained persistence on stock hardware*, which focused on the integration of object-oriented programming languages with database systems to yield *persistent* programming languages. He defended his thesis the day before Thanksgiving of 1994 before driving out of town to take up a tenure-track Assistant Professorship at Purdue University. Those eight years at UMass laid the foundation for his later work in persistence for the systems programming language Modula-3, garbage collection and the PJama project on orthogonal persistence for Java (in collaboration with Mick Jordan and team at Sun Microsystems Laboratories and Malcolm Atkinson, Quintin Cutts, Tony Printezis, Susan Spence, and others at the University of Glasgow). He also was a regular participant at the influential series of persistent object systems workshops originating with the universities of Glasgow, St Andrews, Pennsylvania, and others, whose attendees included Turing Award winner Barbara Liskov and many other famous computer science researchers. Much whisky was consumed along the way, including a highlight tasting with the Malt Whisky Society. These and related events (such as the IDEA workshops, see Page 101) led to his first professional encounters with ANU and the likes of Steve Blackburn, Robin Stanton, and interlopers from Scotland such as Ron Morrison. He turned down his first offer of employment at the ANU in 1998, but was happy to have a second bite of the cherry in 2015.

Marcus Hutter joined RSISE in 2006 and since 2019 has been a Senior Researcher at Google DeepMind in London, and an Honorary Professor in DCS. Between 2014 and 2016, he was Associate Director, Research for DCS.

Marcus was general conference chair for the International Conferences on Algorithmic Learning Theory and Discovery Science[15] co-located with the Machine Learning Summer School at ANU in 2010 (see Page 166).

His research was and still is centered around Universal Artificial Intelligence. His book, *Universal artificial intelligence: Sequential decisions based on algorithmic probability* has been cited more than 1000 times. In 2006, he anounced a 50,000 Euro prize for *Lossless Compression of Human Knowledge*, and in 2020 raised the prize money to 500,000 Euro. His homepage is at `http://hutter1.net/`.

Lynette Johns-Boast gained a BA in Modern European Languages from ANU in 1982 and a graduate diploma in Information Systems from the CCAE (now UC) in 1988. In 1989 – 1990 she worked in London for a small private company called Datastream International, which was later purchased by Dun and Bradstreet.

> We developed financial modelling software used by the banks and insurance companies mainly (for very big bickies at the time) to help them decide where to invest – we initially covered equities (on 40 stock markets around the world) and then moved on to things like options and futures etc. Long before object oriented was a thing,[16] we had implemented the concepts – everything was very modularised, especially the important mathematical calculation modules (once they were proven correct, they were largely left alone because our insurance required that we could certify their accuracy), which were largely treated as 'black box' pieces of code. All of this was done in PL/1. Results were graphed and as you'd expect with models, had a variety of different variables to help anticipate what different situations might give rise to. We used to provide those graphs to the Financial Times – I think we were the only organisation doing this sort of thing at the time. I was there for a second period 1999–2002, but as I had young children I ended up teaching senior high school students during that period as working in the City of London was not feasible. During my time here I developed a PC-based software tool for teachers to create student reports.

Lynette later founded a small company which provides bespoke software and consultancy services to the Australian Government in Canberra.

> In my early days, I worked with DFAT where I developed the software for producing the camera-ready copy of both the Consular and Diplomatic lists, as well as tracking all diplomatic and consular staff in Australia, their cars, purchases, fines etc. I also worked on the passport system and was involved in the early investigation and development of including digital photographs within the passports.

Lynette joined DCS as a lecturer in 2005, taking responsibility for or contributing to many of the software engineering units. In 2012 she received the Australian Council of Engineering Deans National Award for Engineering Education Excellence and in 2014 was awarded Senior Fellowship of the Higher Education Academy of the UK.

See Page 162 for Lynette's opinions on programming and software engineering and Page 277 for her perspective on gender issues in the College.

Hanna Kurniawati is an ANU and CS Futures Fellow Associate Professor. She came to ANU via UQ, MIT's Singapore-MIT Alliance for Research & Technology, and a PhD in robot motion planning at NUS in Singapore. She leads the Robust Decision-making and Learning (RDL) project within DCS.[17] She told me about the robot they are working on:

[15] `https://www-alg.ist.hokudai.ac.jp/~thomas/ALT10/alt10.jhtml`

[16] **Dave:** I can't resist mentioning Simula 67 from two decades before.

[17] `http://rdl.cecs.anu.edu.au/`

2020: Hanna Kurniawati's robot. *Photo: Hanna*

> The robot we have is called Kinova Movo. It is a mobile manipulator with omni-directional drive, 2 arms, a torso that can go up and down, a laser sensor, and a Kinect sensor. It is a product of a Canadian company (Kinova) whose rise to fame is a robot arm compliant enough to be used to help the disabled gain their independence back. My team is a Kinova beta tester for their mobile manipulator, Movo, and I believe the first Kinova beta tester outside of North America. By being a beta tester, we got a discount price and hardware support. The latter is very useful for my team because we focus on endowing robots with intelligence, specifically they become a way to showcase our robust planning and learning algorithms.
>
> My team and I are both users and producers of software. We use C/C++ to develop open source software that implements our planning and learning algorithms. Our software is then used to enable the robot to make deliberate decisions to achieve its task. We do use a host of other software and libraries too, such as CGAL, visualisers such as Rviz, specific collision detection libraries such as FCL, simulators such Gazebo and physics engine they support such as ODE. Occasionally, we use Robot Operating System (ROS).

Hongdong Li is a Professor in the School of Computing, and a founding Chief Investigator for the ARC Centre of Excellence for Robotic Vision.

He started a postdoctoral fellowship at ANU in 2003, choosing ANU mainly because of Richard Hartley's international reputation, and has remained here ever since. From 2008 to 2010 he moved to the NICTA Canberra Lab, located in the NICTA building on London Circuit working on the ARC-funded Australia Bionic Eye project.

The Australian Centre for Robotic Vision was established in 2014 with ARC funding for seven years. Initially there were 13 founding CIs from 4 Australian universities but it grew to about 20 CIs with approximately 100 researchers. Students were very interested in robotic vision research and considerable investment was received from industry. The balance of activities was intended to include 70% fundamental research and 30% applied research. The main theme was to build smart robotic vision systems, able to see and understand unstructured environments and to perform tasks requiring intelligent decision-making. Remarkable progress was made during the tenure of the Centre. When the original application proposal was written in 2013, the value of deep learning had not been fully appreciated. In 2014, it was necessary to adjust and modify the research agenda and plan to include more research in deep learning. The adoption of deep learning technology has proven to be effective, results in superior performance, and avoids large amounts of manual work in feature extraction. It has led to the democratisation of Artificial Intelligence and machine learning technology thanks to the unified computing architecture of deep networks.

Hongdong (along with Henry Gardner and Alistair Rendell) has collaborated with Microsoft Research, Beijing on HPC (high performance computing) and HCI (human computer interface) for medical image analysis. That research was funded as a major ARC Linkage project.

Hongdong's main research interests are 3D computer vision, augmented reality and autonomous driving. He has won two best paper awards at major international conferences on computer vision, and initiated and helped the establishment of the ANU Masters degree program on Machine Learning and Computer Vision (MLCV).[18]

At the present he is involved in three industry projects:

- A project with Medicago[19] to grow better nucleotide binding plants for human vaccine production. This Cannadian based company is also developing a clinical-stage vaccine for fighting COVID-19.
- A project to interpret AUSLAN signing using computer vision and natural language processing. His team has created the first labelled collection of labeled AUSLAN video data.
- Bushfire prevention through the ANU-Optus Bushfire Research Centre of Excellence.[20] They already have cameras installed at each of the four ACT Fire Lookout Towers. They have adopted a goal of research, aiming to detect medium size fires within a short time after ignition by using AI and video analysis technology.

Weifa Liang received a PhD degree from the Australian National University in 1998, after an ME degree from the University of Science and Technology of China in 1989, and a BSc degree from Wuhan University, China in 1984, all in computer science. He has been a Professor in DCS since 2015, but left at the beginning of 2021. His research interests include design and analysis of algorithms and protocols for wireless ad hoc and sensor networks, Software-Defined-Networking (SDN), cloud computing, information processing in wireless sensor networks, design and analysis of parallel and distributed algorithms, approximation algorithms for combinatorial optimization problems, and graph theory.

John Lloyd arrived at ANU in December 1998, to take up the position of Head, Computer Sciences Laboratory (CSL), RSISE, after 11 years as a Professor at the University of Bristol. John lists his research interests as computational logic, agents, machine learning, and declarative programming languages, and has recently become interested in stochastic filtering. His *Foundations of logic programming* book has been very influential. His advice to young researchers is clear: "Find a field you really like and go as deeply as possible. Don't keep switching fields like I did!"

John stepped down as Head of CSL at the end of 2007, after nine years, and Bob Williamson took over. Two years later CSL merged with DCS to form the Research School of Computer Science. In John's view that was an essential move, creating a large enough group to be more internationally competitive. His only insistence was that the research-only IAS contracts of CSL staff be respected so that CSL staff should not be forced into undergraduate teaching. In fact, from 1999 onwards most CSL staff did voluntarily teach in DCS. By the time of the merger, all CSL staff were doing undergraduate teaching, some as much as a full course per year. CSL staff also contributed actively to the design of new courses and programs. John and the Head of DCS throughout almost all this period, Chris

[18] https://programsandcourses.anu.edu.au/2019/program/mmlcv

[19] https://www.medicago.com/

[20] https://www.anu.edu.au/files/resource/DVC200149%20ANU-Optus%20BRC%20brochure%20v6%20%28150ppi%29.pdf

Johnson, had an excellent working relationship which ensured productive collaboration between CSL and DCS.

1974: John and Sue Lloyd punting in Cambridge. *Photo supplied by John.*

John retired in 2013, but has retained an emeritus position and is currently over 500 pages into a new (quite mathematical) book about empirical beliefs.

Paul Mackerras started an ANU PhD in the Information Sciences Group within the the Department of Engineering Physics, RSPhysS, with Terry Bossomaier as supervisor. After two years he switched to having Richard Brent and Iain MacLeod as supervisors. He submitted in January 1988 and graduated later that year, working on the KRIS project (See Page 101) until 1991 and then on the ANU-Fujitsu CAP Project until September 1999. Paul takes up the story:

> I started with Linuxcare Inc. in October 1999. Tridge had submitted his PhD in early 1999 and then looked around at what employment was available for open-source/Linux developers, which in 1999 was a lot. The way I like to tell the story is that he went around to various US companies saying "Hi, I'm Andrew Tridgell, I wrote Samba, would you like to give me a job where you pay me to do whatever I like", and he got five offers. ☺
>
> Anyway, Tridge decided to take a job with Linuxcare Inc, a San Francisco-based startup company with venture capital funding that aimed to provide business-grade support for Linux distributions so that business enterprises would feel confident to use Linux in mission-critical operations. Consequently they were hiring open-source developers, particularly ones with established contributions to Linux and related software. Tridge was at that point their only Australian employee. Tridge recommended that I apply for a job with them (on the basis of my work with Linux on PowerPC and the PPP daemon), which I did, and they offered me a job. That then led to setting up an Australian subsidiary and the Canberra office and the formation of "OzLabs".
>
> Linuxcare started well but ran into various woes including the NASDAQ crash in late 2000, and in early 2001 all but one of the Australian employees of Linuxcare resigned, and most of us went to work for IBM, who established a lab in the Canberra office for us. Organizationally we were part of the IBM Linux Technology Center, and remained so until just a few years ago, when the LTC got absorbed into the IBM Systems group firmware/OS division.
>
> So in March 2001 I started working for IBM along with several others from Linuxcare, including Tridge, Martin Schwenke, Stephen Rothwell and David Gibson (all ex-ANU).
>
> My role all along at IBM has been as a senior technical leader in the Linux Technology Centre, leading and helping to direct IBM's work on various parts of Linux; initially the

support for the PowerPC architecture and IBM's POWER servers in particular, and then since 2010 leading the work on KVM (Kernel Virtual Machine) support for IBM's POWER machines, which provides open-source virtualization (hypervisor) capabilities.

A story about Paulus: Brian Molinari

A family friend, let's call him X, gave me a call: "We need a programmer, fast, can you help?" "What's it about?" I asked. The story was complicated.

X was a director of a (then) small economics consulting company and their high-profile product was a monthly economics forecast. The forecast was computed by a collection of programs which collected raw data from a range of sites, ran an econometrics model, and generated several reports. The trouble was that the system had been developed by a director Z who had just fallen out with the other directors and had decamped with his software sources and, more importantly, with his knowledge of the system. A court order had established that the software belonged to the company, and had to be provided on request. The system knowledge belonged to Z, and he didn't have to volunteer anything (and he didn't). The company didn't know what it didn't have. And the monthly forecast deadline was a few days away.

"You have a reverse engineering problem", I said. "Beyond my skills. But I have somebody in mind, called Paul Mackerras." "Tell him we will pay serious consultancy rates", said X. I managed to talk Paulus into taking it on, because it "sounded interesting". The money didn't seem to be all that persuasive.

My friend X still relates what happened over the next couple of days, in a process that was indistinguishable from magic to those who watched. Paulus would look at the PC screen, type for a bit, say "Aha", and write down something on a piece of paper. The office manager would drive that down to director Z's house in the suburbs and hand it to Z. "Blast!" he would say, "I didn't think he would find that", but would hand over a floppy disk containing the requested component. This was driven back to the office and handed to Paulus, who would read the floppy disc and look at the PC screen and then write down another request. In the end the system was reconstructed and the deadline met. Paulus was spoken of in reverential terms for many years, and I even got a small honourable mention for having "found" Paulus.

I asked Paul about it later. He said he needed to use most skills in his repertoire: a `ROT13` code, reverse engineering of Pascal programs which provided hard-wired file names, Excel macros, and so on. Most programmers could get some bits done, but it took a master programmer to get it all done, against a tight deadline.

Eric McCreath came to DCS in 2001 from the Basser Department of Computer Science at the University of Sydney, having completed his PhD at UNSW. He has been Site Director for the ACM Programming Contest for a total of ten years. He has taught a remarkably diverse set of different topics, ranging from Computer Graphics and Artificial Intelligence to Software Engineering, Operating Systems and Machine Learning. His current research focus is on employing novel architectures, such as the cell processor, GPUs,[21] and FPGAs,[22] for computationally intensive tasks. The two domains he has been looking at are computational chemistry and computer vision.

Brendan McKay arrived in DCS in 1983, after completing a PhD at the University of Melbourne, and spending three years as an assistant professor at Vanderbilt University in Nashville, Tennessee. He became [in]famous for the simulated machine A02SUCS, used by students in A02. He has remained at ANU since then and became a professor in 1998 (now emeritus). His fields of research are combinatorics and graph theory. He has made significant contributions to graph isomorphism, and proved that Ramsey number $R(4,5) = 25$. Few of us knew what Ramsey numbers were but we celebrated anyway!

To wider audiences he is known for collaborative work on the "Bible codes", debunking the theory that patterns in the bible reveal divine messages predicting future events. Brendan used the debunked Bible decryption method to find "hidden messages" about Yitzhak Rabin and Princess Diana in other books. For further detail, please see the Wikipedia article (`https://en.wikipedia.org/wiki/Brendan_McKay`).

[21] GPU – Graphics Processing Unit.

[22] FPGA – Field Programmable Gate Array.

I have a vague feeling that I've seen a photo of Brendan at a Departmental Advance in Kioloa, dressed in a nightie. Surely that can't be right.[23]

Gavin Michael was a PhD student at ANU between 1989 and 1996. He worked on the ANU-Fujitsu CAP Project and his PhD thesis was entitled *Compiler Technology for Multi-computers*. While at ANU, he was a leader at Burton & Garran Halls and was instrumental in bringing the internet to ANU halls of residence, and stoves and fridges to Burton & Garran. Since graduation he has built a stellar career with a series of CTO and CIO roles at international companies including Lloyds TSB Group, Lloyds Banking Group, Accenture, Chase Bank and Citi.

Bruce Millar joined CSIRO's Division of Computing Research in 1970, after completing a PhD and post-doc in *Speech signal analysis, synthesis and perception* at the University of Keele (Staffordshire, UK). He soon developed links with linguists and psychologists at ANU. In 1973 he crossed Clunies-Ross Street to join the ANU Computer Centre, having persuaded Mike Osborne that the Computer Centre's consulting role should apply to a broader range of disciplines, and that he was able to provide that capability. The vacancy he applied for was created by the departure of Col Jarvis.

The Computer Centre academics had offices on the top floor of the Cockcroft Building but Bruce was given an office/lab space in G Block opposite the Menzies Library. The linguists and psychologists were keen to get access to Analogue/Digital converters and, in early 1974, Bruce was able to set up an "experimental terminal", comprising a PDP-11/40 running the real-time operating system RT-11.

Although this building is now labeled "H Block", I think part of it may have been called "G Block" in earlier times. Nearly all of the old administrative blocks have now been demolished. ***Photo: David Hawking.***

In 1974, following a Vice-Chancellor's Working Party report on University Computing, Bruce was propelled into the temporary position of Acting Manager of Computer Services, in order to strengthen the role of Computer Centre Operations Manager, Alan Harris during an intense time of setting up the initial ANU computer network. He held that position for about 15 months until the appointment of Jerry Blackeby. During this time research took a back seat to meetings and consultations with external suppliers. After Jerry arrived, Bruce was freed to return to his academic position. Now Bruce and the "experimental terminal" moved to the Leonard Huxley Building in 1976, next to the

[23] Brendan: We were at Kioloa and we were supposed to break into groups with each putting on a short play based on a fixed script. In my group were about 6 people who included Bev Johnstone and Brian Molinari. For some reason I can't remember, we needed someone to wear a nightie, but Bev refused. So Brian volunteered to do it, but when he tried it was too small and he couldn't get it on. That's when I should have hidden.

computer suite. The experimental terminal was enhanced to satisfy requirements of colleagues Dan Nessett and Garth Wolfendale and a sound proof room, used by the linguists, was also established.

In 1977, following a University review of computing, the Computer Centre academic staff were temporarily assigned to the Vice-Chancellor's Computing Research Group (VCCRG). Bruce eventually chose to join the Information Systems Group in RSPhysS which had compatible research interests. Dispersal of the VCCRG was overseen by Deputy Vice-Chancellor Ian Ross, and facilitated by a small quantity of funds. When Bruce joined RSPhysS in 1979 and negotiated with John Morphett (Laboratory Manager, RSPhysS, and later Managing Director of ANUTech) for an office, computer room and sound-proofed lab space, the chancelry funds were used to remodel the available space. John apparently said, "Welcome to RSPhysS but we need you like a hole in the head!" Bruce's new, lowly but adequate premises were in the basement of the Oliphant Building among the piping, ducting and wiring.

During this period Bruce interacted with many people around the campus. Malcolm Newey was experimenting with mathematical typesetting and made use of Bruce's equipment including a Versatec printer/plotter. Bruce worked with visiting linguist, Warren Glover, to develop a printer for high-quality Devanagari script (used in classical Sanskrit and modern Hindi) using techniques he had previously used for synthetic speech from text by rule. There was interest in the Devanagari and other non-Roman scripts by Richard Barz and Tony Diller from Asian Studies. However, in 1984, these methods for printing such scripts were surpassed by the flexible font system on the Apple Macintosh.

During the uncertain period of the VCCRG, Bruce began a collaboration with Graeme Clark (of cochlear implant ("Bionic Ear") fame). Clark had travelled internationally looking for speech scientists, and encountered Bruce's former PhD supervisor in the UK. Bruce spent a month with Clark's group once the first patient had recovered from surgery and subsequently spent several days per month there. Bruce's collaboration with the otolaryngyology group continued for nearly 20 years, involving the supervision of several research students. However it began with the design of the first patented speech processor for the multi-channel cochlear implant, partly based on work carried out during his post-doc at the University of Keele in 1969.

In 1985 Richard Brent was appointed to a chair in the Information Sciences Group alongside Bruce (speech and auditory communication) and Iain Macleod (graphical communication) and it was renamed the Computer Sciences Laboratory (CSL).

In the early 1990s, Bruce organised a successful syndicated research bid, on the subject of computer user identification based on multimodal interactions through speech, cursor movements and typing, to Harry Triguboff.[24] The bid involved collaboration with Psychology (Doug Mahar) as well as Bruce, Iain MacLeod, and E.V. Krishnamurthy from CSL. This TRUST project had a board of which Bruce was Chair, and about 12 staff. The Manager was Michael Wagner who had earlier been Bruce's first PhD student. The growth in staff necessitated a move to the top floor of the Le Couteur (Mathematical Sciences) Building (now demolished and being rebuilt.) The TRUST project created IP, which on his retirement in 2006, Bruce arranged to formally transfer to Michael Wagner who continued related work at the University of Canberra.

In 1996, the Computer Sciences Laboratory moved to the RSISE building (now the Brian Anderson Building), in which a new sound-proof room was created, and a large space for the then declining TRUST project. Bruce describes Richard Brent's departure for Oxford in February, 1998 as a calamity. Bruce became acting Head of CSL. There was uncertainty about how to re-establish and Bruce had many discussions with Markus Hegland on how to bring the lab, with its continuing disparate interests, together. He also recalls making a trip to Bristol to discuss futures with John Lloyd, who had applied for Richard's position.

From 1994-1998 Bruce chaired the bid and organising committee for the 1998 International Conference for Spoken Language Processing held in Sydney. This was successfully completed with nearly 1000 delegates, and placed Australia more securely on the map with the International Speech Communication Association (ISCA) which later honoured Bruce as an ISCA Fellow.

[24] https://en.wikipedia.org/wiki/Harry_Triguboff

A 1998 review of RSISE had recommended greater focus on recruiting and supporting higher degree students. In 1999 Brian Anderson appointed Bruce as Associate Director (Students). The role was part time but the goal was to triple the number of students, and to incorporate 100 hours of coursework provided by existing academic staff for RSISE PhD students. Bruce worked strenuously with the staff, the ANU Graduate Degrees personnel and various feeder pathways to work towards this goal via policy adjustment and internet advertising. What little time was left was taken up with PhD student supervision, international committee work, and consulting.

In 2001 Bruce participated in the National ICT Centre Of Excellence bid team and, following the success of NICTA in 2002, engaged in sometimes tense discussions with both the ANU Graduate School and the ANU executive about changing ANU PhD rules to suit NICTA. After NICTA started in 2002, Bruce became Associate Director of the Canberra NICTA lab, with a focus on education. He continued in this role until 2005 and retired in 2006.

Two other activities marked Bruce's career. He was insrumental in the foundation, and subsequently secretary and president of ASSTA (Australasian Speech Science and Technology Association) Inc., a professional association for those working in this broad multidisciplinary field. He was also successful in obtaining an ARC LIEF grant2 (renewed three times) for the creation of an Australian speech data corpus. The grants involved ANU, University of Sydney and Macquarie University.

Ross Parker worked in DCS as a tutor in the early 1990s, after a cancer diagnosis. Trevor Vickers recalls that, "He had a wisdom about him, partly I think because the DCS gig wasn't on a critical career path for him, so he didn't have to be careful about saying the right thing." Trevor and Kerry also remember him for a great fruit cake recipe, which has powered a number of DCS morning teas.

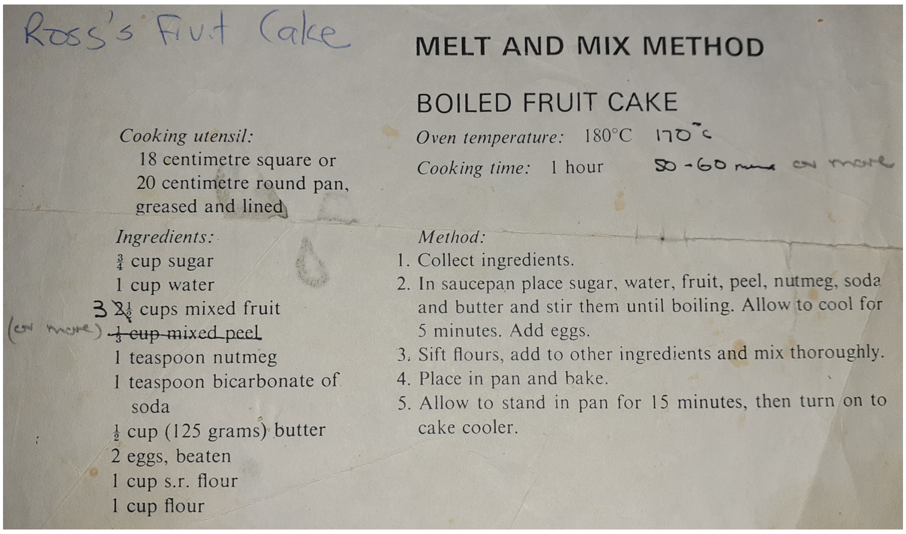

Ross's Fiut Cake

MELT AND MIX METHOD

BOILED FRUIT CAKE

Cooking utensil:
18 centimetre square or
20 centimetre round pan,
greased and lined

Oven temperature: 180°C 170°c

Cooking time: 1 hour 50-60 mins or more

Ingredients:
¾ cup sugar
1 cup water
3 ~~2½~~ cups mixed fruit
(or more) ~~⅓ cup mixed peel~~
1 teaspoon nutmeg
1 teaspoon bicarbonate of soda
½ cup (125 grams) butter
2 eggs, beaten
1 cup s.r. flour
1 cup flour

Method:
1. Collect ingredients.
2. In saucepan place sugar, water, fruit, peel, nutmeg, soda and butter and stir them until boiling. Allow to cool for 5 minutes. Add eggs.
3. Sift flours, add to other ingredients and mix thoroughly.
4. Place in pan and bake.
5. Allow to stand in pan for 15 minutes, then turn on to cake cooler.

Ross Parker's fruit cake recipe. ***Courtesy of Kerry Taylor and Trevor Vickers***

The following is from the 1996 DCS Annual Report:

> It is with deep regret that the Department records its first bereavement. Mr Ross Parker passed away on 8 May 1996, after a long bout with cancer. He held a Level A position with the Department through 1992–1995. Ross Parker came to computing after a career as an engineer in the automotive industry, and brought to the Department the common sense and pragmatism of a true engineer, as well as a deep concern for students in the learning process. His point of view, always politely and courteously argued, was often a valuable counterbalance to the received academic wisdom of the rest of the group.

Shaun Press joined the DCS programmer group in the mid-1990s and headed the group after I joined CSIRO in 1998. Nola Whitecross recalls, "Shaun was a serious chess buff and arranged serious chess events. He was always willing to come to my aid and assist with my many computer difficulties." Richard Brent met him in the Canberra Chess Club long before he joined DCS. Says Richard, "He was a good player and one of the best chess organisers in Canberra, being heavily involved in organising the annual Doeberl Cup."[25]

Alistair Rendell gained a PhD in Theoretical and Computational Chemistry from the University of Sydney and, after a number of postdoctoral positions in Sweden, the US, and the UK, came to the ANU Supercomputer Facility in 1995 as an academic consultant. In 2001 he took up an academic position in DCS, becoming Director in 2013. At ANU he was involved in a number of collaborations with IBM, Microsoft, Fujitsu and Intel and has made significant contributions to computational chemistry packages.

1985: Alistair Rendell (furthest back, very tall, two or three to the left of John Pople, Nobel Prize winner) on his first visit to ANU – a Chemistry Conference in honour of David Craig. ***Photo supplied by Alistair.***

Alistair was Director of the Research School of Computer Science from 2013 – 2018 and joined Flinders University as Vice President and Executive Dean for the College of Science and Engineering at the start of 2019. He recalls:

> I joined what was then DCS, within the Faculty of Engineering and IT in January 2001 as a senior lecturer. Chris Johnson was Head at the time. Prior to that I had been working at the ANU Supercomputer Facility. This included some interaction with DCS through the Fujitsu Area 2/3 projects. I was also interacting with Henry Gardner on delivery of some early courses focused on computational science. Indeed my position was initially funded through the Australian Partnership in Advanced Computing (APAC) to support the development of computational science at the various partner sites. At ANU, under this arrangement, DCS and the Department of Mathematics (DMaths) had committed to the creation of a Bachelor of Computational Science degree that resulted in my appointment in DCS and that of Linda Stals in DMaths.
>
> Wind forward a few years. The faculty becomes the college. DCS and CSLab become the School of Computer Science, then this transitions to the Research School of Computer Science with Henry Gardner as inaugural Director. I took over from Henry at the beginning of August 2013 and remained until the end of December 2018.
>
> We made a conscious decision to grow the School and increase our footprint on the world stage. To achieve this we needed to grow enrolments. In 2013 we had a total student load of 456. A presentation I gave in 2014 set a growth pattern that had us at 571 in 2018. In the event the actual

[25] `https://www.doeberlcup.com.au/`

enrolments for 2018 were 1269! Clearly I failed miserably in my predictive capability. Suffice it to say that the period from 2014-2018 was one of very significant growth.

Year	HDR	PGCR	UG	Total Load
2013	89	84	283	456
2018 (predicted)	11	107	361	571
2018 (actual)	82	387	800	1269

A big part of computing at ANU has been the creation of, and our association with, NICTA. This went through various phases. Early on, as Director, there was close collaboration and engagement in NICTA from staff in the School. Unfortunately funding for NICTA was not continued and led to NICTA being essentially absorbed into CSIRO.

Jochen Renz is a professor in the School of Computing. He came to CSL in early 2006, after three years at NICTA, Sydney, where he had become less enamoured of the top down organisation of research at NICTA. While visiting his group in Sydney, Sylvie Thiébaux had told him that there were some interesting positions going in CSL, closing that same evening. Very soon afterward he was interviewed by a panel including John Lloyd and John Richards, and offered a job the same day. Despite his reluctance to leave a share house on the cliffs at Dover Heights in Sydney, he accepted, joining CSL at the same time as Marcus Hutter and Alwen Tiu.

Jochen Renz in formal attire. ***Photo supplied by Jochen***

He found the environment in CSL very conducive to productivity. Michelle Moravec was a very competent administrator and she was located in the same corridor as the CSL researchers. Some of the ANU NICTA people were also there, prior to their move to the building in London Circuit. There was a great esprit de corps, and researchers often lunched together.

From 2010 to 2014, Jochen was funded through an ARC Future Fellowship which provided a grant of $50k p.a. for equipment and travel. Since Jochen is essentially a theoretician, this meant that he was able to travel extensively. One of these trips took him to an IJCAI 2011 workshop in Barcelona, which he rates as one of the best conference experiences of his career. Attendees were so engaged in the subject that they were shouting at each other!

That noisy workshop led to the Angry Birds AI competition.[26] Jochen organised a team of students of all levels from undergraduate to PhD to work on the framework for the competition and on software to participate in it. The first run of the competition was at a 2012 Australasian Artificial Intelligence conference in Sydney; since then it has run annually, mostly at IJCAI, with Jochen as Coordinator.

Playing Angry Birds is difficult. There are infinitely many possible actions and they have unpredictable outcomes. Humans are good at interacting with a physical environment but machines are not. After every competition, the winning AI agents compete with humans. The humans always win and Angry Birds remains a big challenge for AI.

Jochen always felt that Angry Birds was not the sort of subject likely to be funded by the ARC, so he never applied. He did however apply for DARPA funding in the US and was successful in gaining $A2.3M over 3.5 years. He says that working on a DARPA project is very rewarding but also very demanding. Unlike ARC grants, project reporting is required monthly and there are frequent phone conferences, at uncomfortable hours. Funding is granted project phase by project phase. Funding could be terminated at the end of each phase.

Prior to NICTA, Jochen held a Marie Curie fellowship at the Technical University in Vienna. While there he enjoyed the many cultural activities on offer, he went to the Vienna State Opera weekly, including many premieres, also theatres and classical concerts and even the famous Opera Ball. At one premiere he says he was body checked by the Austrian Chancellor, Wolfgang Schüssel. However, on another occasion, Schüssel held the door open for him, so Jochen considers they are even. "It's a very small world in Vienna. Almost everything you see on the news happens within 10km of where you are."

Jochen likes the fact that ANU allows course controllers to be as creative as they want in delivering a course. He uses this freedom in teaching a course on Research Methods and another called Studies in Advanced Computing (R&D). He trains future researchers on how to conduct research, how to present it orally to non-experts, how to write it up for publication, and how to constructively critique the work of others. He is pleased that attendance at class sessions is unusually high and delighted that some students have had their papers accepted in top conferences. He thinks that all undergraduate students wishing to pursue a research career should publish, since having a publication list is valuable when applying for PhD positions at top US universities like Stanford and MIT.

Jochen is highly motivated to generate real world impact, perhaps through a start-up company.

Ramesh Sankaranarayana is retiring from DCS after 26 years of teaching since his arrival in 1995. He completed a BE in Electrical Engineering at the University Visvesvaraya College of Engineering in Bangalore, India and worked briefly for an Indian telecom before enrolling for a Masters degree at the Indian Institute of Science, also in Bangalore. Finally, he completed a PhD in the area of graph theory at the University of Alberta in Edmonton. Arriving at Edmonton Airport in winter, his clothing was inadequate for the wintry conditions, but he was rescued by a friend who came, with a coat, to pick him up at the airport.

Several of his friends from Alberta took university positions in Australia. After graduating in 1994, he followed some of them to become a visiting fellow at Griffith University. He later visited a friend in Engineering at ANU, engaged in conversation with Brendan McKay and met other DCS members. He subsequently interviewed at ANU and commenced a fixed term teaching position in April 1995. He taught computer architecture, advanced algorithms and a course on formal languages, computability and complexity.

After three years, the course on formal languages was dropped and relevant content was dispersed into other courses. A bright student, Vijay Boyapati,[27] was determined to do the unit, despite

[26] `http://aibirds.org/`

[27] Vijay was a member of a team which built a *My Network News* system as a group project in a 1999 Software Engineering course taught by Malcolm Newey and Richard Walker. Malcolm says that this experience provided a springboard for his move to Google.

it being no longer offered. He persuaded Ramesh to let him do the unit as a special-topics course with Ramesh as the supervisor. Vijay joined Google before the Google IPO. He later funded the Boyapati Computer Science and Mathematics Scholarship at ANU.

1999: COMP3018 *My Network News* team with Vijay Boyapati at front left. *Photo: Malcolm Newey.*

In 1999, Ramesh accepted an invitation to become Associate Dean Undergraduate (Computer Science) for the Faculty of Engineering and Computer Science, taking over from Vicki Peterson. He held the post for over nine years. Initially, this involved a very heavy workload at certain times of the year, because FEIT had no faculty staff of its own, relying on part time assistance from Jan O'Connor and Mandy Metcalf from the Faculty of Science.

Over the years, Ramesh developed and taught a number of courses, including IT & E-Commerce, Computing for Engineering Simluation (which more recently became Computer Architecture and Simulation), Data Structures and Algorithms 1, Networked Information Systems, and Introduction to Data Management, Analysis and Security. His research interests are in the areas of Software Engineering, Information Retrieval, and Security.

Ramesh will retire in May 2021, after spending over 26 years at ANU. He feels that the DCS world has changed substantially over the quarter century he has been part of it. In the nineties, students used to (mostly) attend lectures, even if they flew paper planes when the lecturers back was turned! Now, the lecture content is online, student attendance is lucky to reach 30% on a consistent basis and attention spans are poorer. He found student feedback quite useful, particulary the old-style paper version. Now, with SELS (Student Evaluation of Learning System) and low student participation, the value is not as high. He's also not sure that the reduction in teaching weeks from 13 to 12 weeks was a good thing.

Despite the increased size of the School, Ramesh still feels that it has a friends-and-family feel to it.

John Slaney is an Emeritus Professor in DCS, having led the Automated Reasoning Program for many many years. He was once all-England schools champion in triple-jump, and after coming to Australia was hopeful of making the olympic team. Unfortunately, injury put paid to those plans. He completed a PhD titled, *Computers and relevant logic : a project in computing matrix model structures for propositional logics* at ANU in 1980, and taught logic in philosophy departments for several years, before returning to ANU in 1988, where he has been automating reasoning ever since. He retired in 2020. You will find more about John on Pages 165 and 174 .

Peter Strazdins has been a member of DCS since 1990 and is an associate professor. He grew up in Wollongong and left with two degrees and a taste for surfing. He completed an MSc at the University of Oxford before enrolling for a PhD at ANU in 1987. His principal supervisor was Heiko Schröder and the topic was *Control Structures for Mesh-Connected Networks*. During his PhD he also completed four fourth-year honours subjects, including AI, VLSI Design, Architectural Description Languages and Complexity. He rates Mike Robson's Complexity exam as the hardest he ever sat – a difficult NP-completeness proof.

Three days after completing his PhD in 1990 he joined DCS, 50% on the ANU-Fujitsu CAP project and 50% teaching. His first lecturing assignment was an introductory CS course to Economics and Commerce students. He found them incredibly competitive and prone to arguments over marks. Later he took on Computer Organisation (A02, later A12). This unit made use of the A02SUCS simulated machine – See the panels on Pages 59 and 151.

On the CAP Project, he worked with Richard Brent on porting the BLAS (Basic Linear Algebra Subprograms) to the AP1000. He also developed the first parallel symmetric indefinite linear system solver algorithms. This, together with other optimizations, sped up Fujitsu's commercial ACCU-FIELD electromagnetic compatibility application by a factor of 3. His research interests include high performance computing, computational science, numerical algorithms and systolic arrays.

Peter was responsible for a unit in program design and construction. He required students to build a moderately complex system in Modula-2, involving abstract data types. The assignment involved the simulation of U-boats (submarines) and E-boats (fast attack boats), with innacurate sonar. Peter Farmer developed a GUI for the assignment. There were criticisms that this exercise was too male oriented, but Peter found that in general females enjoyed it.

Peter taught in a unit on high performance scientific computing, and with Chris Johnson in a unit on parallel systems. In the early days of multicore computers, Peter ran a course on multicore computing. It used equipment donated by Sun Microsystems (a machine based on the UltraSPARC T2 (also known as Niagara T2) chip with 8 cores, each supporting 8 concurrent threads) and Intel. The Intel machine was a single-chip cloud-computer with 48 cores arranged in a 2D mesh.

In his research in high performance computing, Peter pioneered the use of look ahead in parallel matrix operations. He however advocated the alternative technique of what became known as algorithmic blocking, pioneered by Richard Brent in the CAP Program and by Bruce Hendrickson, Director of Computer Science at Sandia National Laboratories in Albuquerque, New Mexico. Peter has since collaborated extensively with Bruce and others at Sandia and Lawrence Livermore National Laboratories.

Peter has held four Linkage grants, in the HPC area, two with Alistair Rendell. Partners included Intel, Sun, Platform Computing and Fujitsu Labs Europe. Example activities included getting Gaussian going on Sun's latest shared-memory hardware, developing cycle-accurate simulators for non-uniform shared memory processors and developing service-oriented approaches to HPC.

In the collaborative project with Fujitsu Labs Europe, led by Markus Hegland of the Mathematical Sciences Institute, he developed the first parallel Sparse Grid Combination Technique algorithms, with which his PhD student of the time, Muhammad Ali, delivered a fault-tolerant version of the widely-used parallel GENE gyrokinetic application.

Hanna Suominen is the Leader of the Big Data Program in Our Health in Our Hands (OHIOH), the inaugural ANU Grand Challenge Program,[28] and Team Leader of Theory and Applications in Multimodal Pattern Analysis (TAMPA) within the Machine Learning (ML) Group of CSIRO Data61. See the panel on Page 131 for Hanna's story.

Hanna and her colleagues Inger Mewburn, Will Grant, and Lindsay Hogan are working to commercialise their research in the form of Postac®:

> The PostAc® platform is a searchable job advertisement database. It uses artificial intelligence, machine learning, pattern recognition, and data mining to find the research intensive 'needles' in the job market 'haystack'. The PostAc® algorithm that underpins the platform can 'read' millions of job advertisements and rate them via research skills intensity — the so-called 'nerdiness rating'. PostAc®[29] helps higher degree graduates and postdocs to find non-academic jobs that use their highly developed research skills.

[28] https://www.anu.edu.au/research/research-initiatives/our-health-in-our-hands

[29] https://postac.com.au/

Hanna Suominen and the Schneewittchen

I was born in the middle of a snowy winter in Loimaa, Finland. My name Hanna (I have been favoured with a child) reflects my parents' longing for the arrival of a daughter. Unlike Snow White, I grew up with two kind, loving parents, and a family tradition of supporting education and careers for daughters.

My great-grandmother told me the tale of the Aspen trees who did not know how to pause from talking to observing — without this ability one would be cursed with ever soughing leaves.

I moved to Australia in 2010, and co-founded the Document Analysis (DA) unit at ANU. In Finland, the title of *Dosentti* (Docent) is awarded to academics who achieve a high level of independent research (at least equivalent to two PhD theses) and a high standard of teaching. As part of my successful application for this title, a lecture from the DA course was assessed by staff and students to be at the highest level. These activities were in my area of expertise — developing Machine Learning (ML) and text mining to make meaning and read between the lines.

In 2016, even though it was outside my area of expertise, I was invited to revise and convene one of the most diverse, mandatory, and hence difficult units of the school. I was successful in shaping the Networked Information Systems (NIS) unit to demonstrate research-led education in ANU – I used my research in medical ML to motivate the students. Student satisfaction and outcomes in NIS improved.

However, I was fool enough to bite the odd poisoned apple of toxic feedback. In particular, the curriculum-level learning objectives on academic writing caused friction in NIS, although improved skills in academic writing are likely to lead to noticeable gains in assignments and exams. Disheartened by some negative student experiences, I felt like I was not an educator at all. The mirror on the wall told me I was not a beautiful white swan.

For Snow White and me alike, our caring community was the greatest asset that watched over us while we were experiencing our existential crises. The dwarves protected Snow White from the evil Queen and sheltered her in a glass coffin until her awakening. Likewise, my supervisors, colleagues, and students in ANU rescued me from my coffin of ugly duckling reflection.

As an outcome of this mirroring and personal growth, I focused more on education and, strongly supported by my school, gained a Masters of Leadership (Curriculum and Pedagogy) and became a Senior Fellow of the Higher Education Association (HEA).

While resolving my educational philosophy, I also realised that I have an interesting academic ancestry of PhD supervisors. (Based on a search on the supervision relationships site `https://www.genealogy.math.ndsu.nodak.edu/`.) On one hand, I am 14th-great-granddaughter of Gottfried Wilhelm Leibniz (differential and integral calculus, 1646–1716, Germany) and 20th-great-granddaughter of Nicolaus Copernicus (Sun-centric universe, 1473–1534, Poland). On the other hand, I am the only woman along this "family" tree, traced back to the 1360s, who managed to get a PhD and join the extremely competitive profession herself as an academic leader and PhD supervisor. On the bright side, CECS is aiming for a 50:50 gender balance.

I recognised that my academic ancestry gives me an opportunity to take a leadership position that models this possibility for other women in a massively male dominated field. I have a passion to lead the way and show how women can excel and become professors in science, technology, engineering, mathematics, and medicine (STEMM). I saw in myself an emerging catalyst for change who was gifted in influencing others and taking emerging leaders along. I had my Finnish *sisu* — an adamant faith in my team to focus, commit, and deliver important information and artefacts to our academic community, clients, and society.

Outreach talks, publications, and winning prizes are fundamental steps in my strategy of inspiring people to excel in STEMM, take the lessons learnt to their work, and thereby, bridge the gap between ML and its use in society. I'm proud of my real world impact too. For instance, we have shown the potential for ML in detecting symptoms imperceptible to a neurologist as simply as by saying `/aa/` into a microphone.[a] Among other resources, I have developed search engines for the Turku University Hospital and smart sports-sensors for the Australian Institute of Sports (AIS). As a co-founder and the chief ML scientist, I have launched the PostAc® web service.[b] I hope to use my ANU work to create an ML-powered workforce.

I hope you understand my choice of Snow White as a metaphor. Finally, after years of reflecting and reflecting, the mirror on the wall reveals that the true story of my life as an ANU citizen is the most precious gem of all.

[a] `https://biomedeng.jmir.org/2020/1/e13611/`

[b] `https://openresearch-repository.anu.edu.au/bitstream/1885/209176/1/PostAcTrendReport_PRINT.pdf`

Kerry Taylor. Before coming to Canberra, she worked for Oxford University Press while her husband Trevor Vickers was a postdoc in Oxford. After a period of maternity leave on arrival in Canberra in 1990, she commenced a PhD in Machine Learning at DCS in 1991, specifically *Autonomous Learning by Incremental Induction and Revision*, with Robin Stanton, Graham Williams and Claude Sammut (UNSW) as supervisors. Between 1995 and 2015, she worked in CSIRO: DIT, then CMIS, then ICT Centre, then Computational Informatics, then Digital Productivity and Services, then Digital Productivity, and very nearly Data61, all the while in the same ANU CS&IT Building where she had completed her PhD.

After a brief period at the Australian Bureau of Statistics, Kerry responded to a teaching shortage in DCS and has continued in a series of fixed-term appointments, teaching Agile Project Management, Network Information Systems, and Data Mining. She is Program Convenor for what seems to me to be a remarkable number of different degrees, and has spent considerable time promoting public service enrolments in post-graduate data science courses. Her research focus has returned to her roots in machine learning, especially over knowledge graphs.

In the pandemic year of 2020, Kerry was heavily involved in teaching for the highly-compressed, domestic-students-only, pandemic-funded Graduate Certificate of Data Engineering (GCDE).

Sylvie Thiébaux became, in 2011, the first female professor in CECS. She completed a PhD at the University of Rennes, France, and worked for a couple of years at INRIA (Institut National de Recherche en Informatique et Automatique) also in Rennes. In 1997 she came to Canberra, working for several years in the Intelligent Transport group within CSIRO Mathematical and Information Sciences. She then transferred to the Automated Reasoning Project, and soon after was seconded into the Knowledge Representation and Reasoning group in NICTA, moving to Knowledge Representation Canberra and then founding the Planning and Optimisation Group. She later became Director of the NICTA Canberra Lab. After the demise of NICTA, she became part of CECS, and led the $7.99M ARENA (Australian Renewable ENergy Agency) funded CONSORT project, planning renewable energy and storage for Bruny Island. She has also been active in the "Humanising Machine Intelligence" Grand Challenge project, an ANU-wide collaboration. In 2019, Sylvie was elected editor in chief of the Artificial Intelligence journal and in 2020 she became the first woman to be elected AAAI fellow in the southern hemisphere.

Paul Thistlewaite's mother wanted him to study medicine. After two weeks as a medical student, Paul jumped ship to philosophy, graduating with BA Hons. (1st class) from UQ in 1978. He enrolled for a PhD in logic at Latrobe University but transferred his enrolment to ANU, following his supervisor Michael McRobbie. During his PhD Paul lived in Graduate House where, in 1983, he met the love of his life and future wife Thena Kyprianou. Paul's thesis *Automated theorem-proving in non-classical logics* was completed in 1984 and is, of course, in ANU's online thesis repository. Thena recounts that Paul had been working on a particular problem but was narrowly beaten to the punch by a Canadian researcher, necessitating an extension on his PhD scholarship.

After graduation, Paul took up a post-doctoral position in DCS, funded by Rothmans (yes, the cigarette company). In 1986 he spent 3 months at Argonne National Laboratories in the US. After the post-doc, Paul joined the PISO (the Parliamentary Information Services Office), where he worked on making Hansard and other parliamentary information available via the network. He was promoted to SES level, responsible for delivery the service and managing the team, but to solve PISO internal problems was encouraged to resign and return as a consultant, apparently with a substantially larger pay packet. Someone reported this change in status to Bronwyn Bishop, who asked pointed questions in the Senate and caused both a media "kerfuffle" and distress to Paul.

After PISO, Paul formed a consulting company and carried out a number of small consulting projects for Computer Power, a prominent IT company at the time. A larger project was Paul's development of the AIDA text summarisation system. Computer Power paid him a retainer while he developed the system and contracted to pay Paul a percentage of the proceeds from sales of shrink-wrap versions. AIDA was a success and many copies were sold, but unfortunately a legal dispute ensued over the payment of Paul's percentage.

In 1990, the Australian Government established the Cooperative Research Centres program, with the first selection rounds in 1991. Paul returned to ANU, to work on the bid for the ACSys Cooperative Research Centre. When ACSys started in 1993, Paul was an ACSys Fellow, leading the PASTIME project described on Page 178. He was PhD supervisor to Nick Craswell and Raj Nagappan, and later took up a Senior Lecturer position in DCS, delivering courses on Document Technologies, Software Engineering, and Introduction to Information Systems.

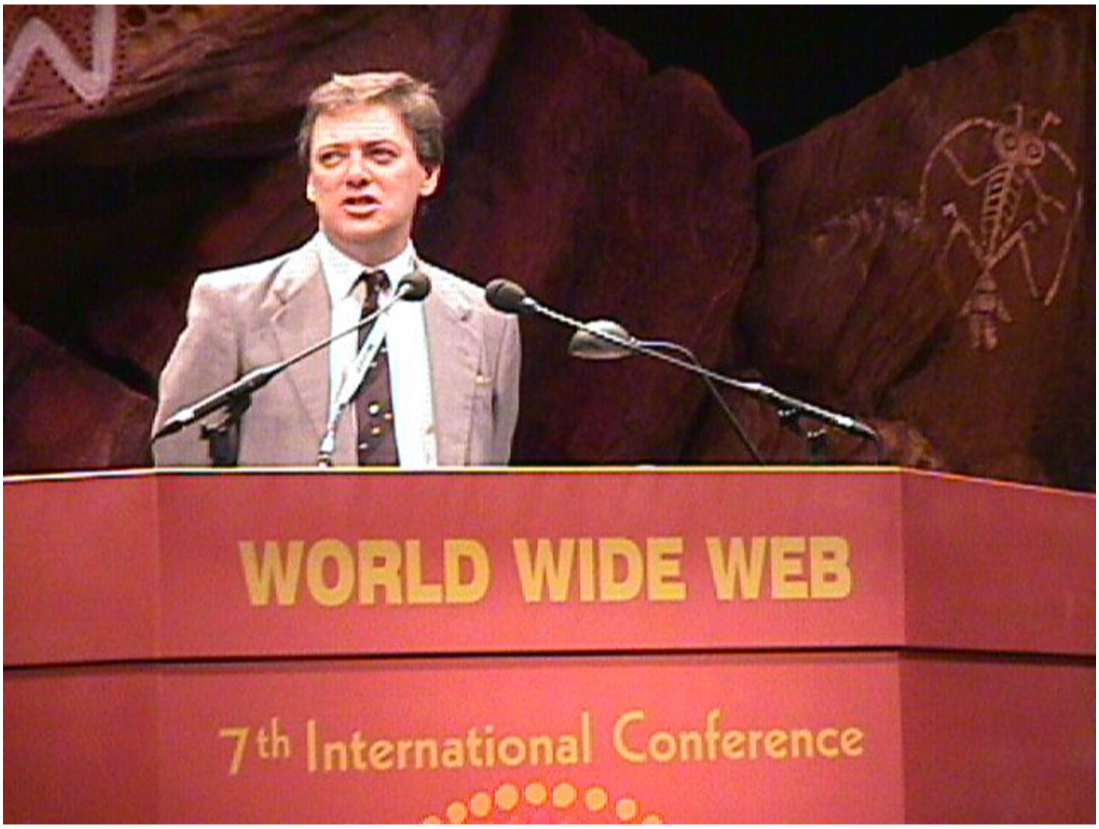

Paul Thistlewaite addressing the crowd at WWW7 in Brisbane. *Photo: WWW7 organisers.*

Paul and I participated in the Text Retrieval Conference (TREC) evaluations. I remember sharing a packet of Tim Tams with him at about 3am in the CS&IT kitchen while we plotted our next move in the campaign to "win TREC."

Paul was an instigator of the AusWeb Australian web conferences and was program co-chair of the World Wide Web conference WWW7, held in Brisbane. This conference included a breakthrough paper by Larry Page and Sergei Brin, founders of Google.

To the great sorrow of his friends and colleagues, Paul became seriously ill late in 1998, and died in February 1999, leaving his wife Thena Kyprianou and daughter Desiree. His memory is honoured by a Best Paper award at the AusWeb conference, and by the Paul Thistlewaite Memorial Scholarship at ANU. The latter was created at the initiative of Thena and the Thistlewaite family, and received many donations from the public. Then Vice-Chancellor Deane Terrell made good on an ANU undertaking to match the total of the private donations. The scholarship funds a female student to undertake an honours year – see the panel on Page 278.

Nola Whitecross, former DCS Adminstrator, remembers Paul Thistlewaite

"Paul Thistlewaite is someone I remember very fondly. He unfortunately died of cancer and his death had a profound effect on me. He was young (too young) and such a nice person. Prior to his death, the Department held an event of sorts, dress up, perhaps a Christmas party, I cannot really remember. Paul was undergoing chemotherapy at the time, and suffering the consequent conditions, loss of hair, weight gain etc. Nevertheless, Paul came to the party draped in a sheet; he was the Dali Lama. I admired his strength and his sense of humour."

I owe a great deal to Paul. It was his vision, laid out in a project proposal (AWE – Analysis of Web Evolution), to create a set of intranet management tools, including the ability not only to search the intranet but to search it as it had been at points of time in the past.[30] After Paul's death, Peter Bailey and I turned part of this vision into reality with the July 1999 launch of the S@NITY search service on the ANU intranet and websites. See the photos on Page 180. S@NITY evolved into P@NOPTIC within CSIRO, which led to the spin-off of Funnelback Pty Ltd in 2005. Funnelback was acquired by Squiz Pty Ltd in 2009. Twenty-one years after the launch of S@NITY the technology lives on in the Squiz product range, and in high-profile customer sites in Australia, the UK and the USA. Many tens of millions of dollars have been generated, and around fifty people are employed in high-tech jobs.

Paul Thomas gained a rare mid-year PhD scholarship to ANU in 2005, with me as primary supervisor. Prior to that he completed a BA(Science Subjects and Philosophy) and a MCompSci at Victoria University in Wellington NZ, then worked for the CSIRO Intelligent Transport Group in Canberra before spending three years backpacking and then temping in London.

Paul's PhD topic was in the area of personal metasearch, aiming to provide a single search interface to multiple search engines operating over the different information spaces accessible to an employee. Paul implemented a personal metasearcher and persuaded a number of ANU people to use it, learning C# and Mono in order to make a version which would run on both Windows and Linux machines. He and I developed and validated a novel way of comparing search effectiveness and Paul used it in experiments with his personal metasearcher. After graduating he worked for the CSIRO ICT Centre, which later became Data61.

He feels that it is very important that PhD students have the opportunity to meet and interact with other students and researchers. At ANU he was given many opportunities to do this through conference attendance, through lab visits, through supervisor introductions, and through a seminar and discussion series he organised for a total of ten years, spanning his PhD at ANU, and his work in the CSIRO ICT Centre. This series was called "IR [Information Retrieval] and Friends." Paul is justifiably proud of this effort which drew in participants from DCS, CSIRO, NICTA, and ANU Sociology as well as international visitors, and triggered a number of research collaborations.

Paul now works in Canberra for the Microsoft Bing search engine and was co-organiser of the CHIIR 2021 conference (Computer Human Interaction and Information Retrieval) which, absent COVID-19, would have been held in Canberra in March 2021.

Andrew Tridgell. Andrew's initial contact with computer science at ANU was via the Computer Sciences Laboratory, where he was a PhD student working on automatic speech recognition under Bruce Millar, starting in 1991. As part of that project he developed an interest in high performance computing, using the CM5 parallel computer for his speech recognition research. This brought him into contact with Richard Brent and members of DCS department.

Over the years at ANU, Andrew has taught very popular Open Source courses, contributed to the Canberra Linux Users Group, and helped countless people with IT problems. (See Page 135.) I'm not sure whether the joke was due to Andrew, but he briefly mystified people when he advertised a talk on his speech recognition work, titled *How to Wreck a Nice Beach*. (How to Recognise Speech.)

Tridge is most famous, of course, for developing heavily-used open-source software:

> **Samba**. At the end of 1991 I started working on a piece of software to help with integration of my desktop computer in CSL with the SunOS 4 file server. The primary motivation was procrastination over the christmas 1991 period, but this program eventually grew to be `Samba`, which formed a large part of my future career. [A history of Samba appears in Appendix I.]
>
> **rsync**. The focus of my PhD changed in 1993 from speech recognition to parallel computing, reflecting the work I was doing on parallel sorting, which initially came from work with speech recognition data. I changed supervisor from Bruce Millar to Richard Brent, Brendan McKay and Paul Mackerras. One of the sidelines that came out of my work on sorting was a novel delta compression algorithm which I developed with Paul Mackerras and we called rsync. I first presented this algorithm in a CS seminar at ANU on 16th October 1996 and it went on to form a chapter in my thesis, sitting rather incon-

[30] That capability would have been very useful in writing this history since, with rare exceptions, the ANU web seems only to include little information about the past.

gruously with chapters on various sorts of parallel sorting algorithms. The rsync program has gone on to be one of the better known tools for unix users.

HiDIOS filesystem. In 1995 I got involved in the ACSys CRC, affiliated with DCS. As part of that project I was the team leader for the HiDIOS filesystem (High Performance Distributed Input Output System). The key technical observation that made the first version of HiDIOS possible was that a "parallel filesystem" didn't need to support more than one file, and didn't need to support directories. This allowed for the very rapid development of a parallel filesystem for the AP1000 parallel computer to meet a one week development deadline for a meeting with Fujitsu. Later extensions to HiDIOS allowed for more than one file and even directories, but developing that initial version where the first argument to open() was ignored did make for a much simpler filesystem.

Former Masters Student Mike Carden's comments on Tridge

I was a Masters student in 2008–2009, spurred on by working full time in a public service role surrounded by government funded PhD graduates. It was a hectic time but a decade later I'm reaping the benefits.

I'm mostly responding to make sure that your book includes at least a chapter on Tridge, aka Andrew Tridgell, who set new standards for ANU DCS for a very long time. Including I'm told, a most protracted doctorate. I can relate the anecdote of him at a CLUG meeting hacking Linux to get my inkjet printer working so that he could print a photo of his puppy, or years later running the ANU's inaugural Open Source course (where I was a guinea pig student), or his roles in rsync, Samba, and remotely piloted vehicles.

I seem to remember that a motivation for developing `rsync` was to make it easier and quicker for Tridge's wife Sue to back up her files over the internet. I believe `rsync` was relied upon by Nokia in its heyday to communicate new software versions from developers in Finland to the factory in Korea.

Nola Whitecross's memory of Tridge's PhD submission

I remember the events relating to getting Andrew Tridgell to submit his PhD thesis. Time was running out and, with about five minutes to spare and a dash across campus, the deed was done. Andrew was reluctant because he was too busy doing interesting things with and for his students, consuming pizzas late into the evening and talking the talk.

Tridge's amazingly quick port of Linux to the Fujitsu AP+ parallel supercomputer is described elsewhere in this history (Pages 216 and 218, and he also developed a chess program called KnightCAP which was able to exploit the parellelism of the Fujitsu AP1000. Tridge's n-gram based speedup of multi-alternate Boyer-Moore-Gosper search (`BMG2`) is described on Page 216.

Tridge inadvertently triggered the development of the `Git` source code repository used in maintaining source code for Linux and countless other projects. To quote Tridge's Wikipedia entry:[31]

> In April 2005, Tridgell tried to produce free software (now known as SourcePuller) that interoperated with the BitKeeper source code repository. This was cited as the reason that BitMover revoked a license allowing Linux developers free use of their BitKeeper product.[3] Linus Torvalds, the creator of the Linux kernel, and Tridgell were thus involved in a public debate about the events, in which Tridgell stated that, not having bought or owned BitKeeper — and thus having never agreed to its license — he could not violate it, and was analyzing the protocol ethically, as he had done with Samba. Tridgell's involvement in the project resulted in Torvalds accusing him of playing dirty tricks with BitKeeper. Tridgell claimed his analysis started with simply telneting to a BitKeeper server and typing `help`.

Since 2011, Andrew has been involved in developing software (`ArduPilotMega`) for Unmanned Aerial Vehicles (UAVs). In 2018, ANU awarded him a DSc (Honoris Causa) and he delivered the Graduation Address. In 2020 he was awarded a Medal of the Order of Australia.

[31] `https://en.wikipedia.org/wiki/Andrew_Tridgell`

A story about Tridge: Brian Molinari

A story of the early Open Source era, c. 1997

An IT contractor (let's call him X) asked me one day: "Do you know Andrew Tridgell? He works at ANU. He's weird." I said I did indeed know Tridge, he worked with DCS, and he was not weird. But tell me.

X explained that he was a principal in an IT contract with Defence. The contract relied on an expensive software package. The software couldn't be got to work as required, and they were within an ace of having to default the contract. They then found out about a possible alternative called Samba, downloaded it and couldn't get it to work either. Gloom. At the last moment they realised the author, Tridge, was in Canberra. They gave him a call, and in a couple of hours he came on site, configured the software and had it running. The software was stable, the contract was saved, and the expensive dud software could be returned. Happiness.

"We'll pay you for Samba", said X to Andrew, indicating it was worth serious money to them. "Oh, no", explained Andrew, "It is Open Source software. Look at the licence". "Can't we pay you anyway?", said X, puzzled by his first interaction with OS software. Again, no. "Let us pay you for your time", said X. "But it only took two hours", said Tridge, "and I help people like that all the time". "Let's pretend it took two weeks", said X (nudge, nudge, wink, wink). But again, no. In the end, all they could press upon Tridge was a pizza delivery to his weekly special-interest group that met on Tuesday evenings.

See, said X, I told you he was weird.

Jochen Trumpf completed a PhD in Mathematics (control theory) in 2003 at the University of Würzburg in Bavaria and came to ANU Systems Engineering via a couple of brief post-doc positions in Israel and Indiana. John Moore, who was head of the department appointed him to a five-year fellowship and the appointment was later extended.

While an undergraduate in Würzburg, he and three colleagues formed a software company to improve travel planning for public transport in the city of Würzburg. In LISP, they built an application, complete with graphics and maps, to deliver a personal timetable for a passenger wanting to travel from any point A to any other point B. (The maximum number of connections required in Würzburg is three.) Since then the software has been made much more efficient and the scope has been expanded to cover the whole of Germany. However, Jochen is no longer involved.

Jochen is now Director of the Software Innovation Institute (SII) within CECS. It was established on 01 July, 2019 as one of the initial group of ANU innovation institutes, which were established to try out different business models for core ANU activities. SII is exploring different ways of teaching Software Engineering, motivated by employer-triggered student attrition,[32] and inspired by medical training. Jochen observes that employers of software engineers and data scientists complain that graduates of university CS courses are not job-ready. This is nothing new, but now it comes with a threat, because many large organisations, including banks as well as tech companies are creating in-house academies. Maybe the solution is to make Software Engineering training at university more like an apprenticeship or the in-hospital training given to doctors. The question is how to create the necessary flow of real software projects for students to work on.

Rather than establishing its own software company, SII sees itself as a broker of partners. In one example they have partnered with ANU's Research School of Population Health, and with other parties to build a real-time COVID alerting app for general practitioners. NHMRC funding has been awarded and co-operation has already been obtained from the Queensland state health service – others are pending. The three graduate students involved are gaining the opportunity to work on such important matters as privacy-preserving data linkage in a consequential real-world situation.

Trevor Vickers started as a lecturer in DCS in 1990. In 1991, he was given the responsibility of coordinating Robin Stanton's initiative to centralise the teaching administration duties traditionally performed by individual academics: the Cooperative Teaching Group (CTG). During his time in DCS, he created a number of important software systems for improving administrative processes within the department.

[32] Employers are sometimes happy to take students who haven't yet graduated.

> I did my undergrad and PhD at UNSW. My supervisor was Ken Robinson, who died in September 2020. Kerry and I were married just before I submitted, in late 1986. I worked at Oxford straight afterwards, for 3 years. Carroll Morgan was my boss there. He's now back at UNSW.
>
> I began at ANU DCS in January 1990. We spent the first few months renting Chris Johnson's house, while he was on study leave, until we found our current home.
>
> I've always been interested in process, and how it can be streamlined, and ANU offered lots of opportunities in that regard. For example, in 1990 students would register for individual tutorial and lab groups by writing their name on a sheet of paper taped to a wall. It was often difficult to read their writing, and some students would cross out another's name and add their own when the class was full.
>
> At the end of 1990 I began designing and implementing a system to do the registrations online. This was ATLES – an Automated Tutorial and Laboratory Enrolment System. ATLES was initially written in shell, and would regularly cause system problems because each command within the code would spawn a new process. In 1992, ANU's Centre for Development of Academic Methods (CEDAM) granted $6000 to further develop ATLES, and I assume it was then rewritten in something more efficient![33]
>
> At the same time, I began working on a system to record student marks within an individual subject (or 'unit' as they were then called). The system was called Tycho, in honour of Tycho Brahe, who was famous for detailed observation and record-keeping, and for occasionally fudging his results when required.
>
> The Cooperative Teaching Group duties, and my belief in automation of dreary tasks led to a proposal late in 1992 for funding from DEET's National Priority (Reserve) Fund for Quality Management and Assurance, which was selected by ANU to be its sole submission. We received a grant of $166,700. I remember our initial proposal was for about 2/3 of that amount, but once we were selected by ANU, we were advised to increase the budget to make it look more realistic!
>
> That funding produced DAIS – a Departmental Administrative Information System. It built on some of the concepts from Tycho, and attempted to make it simpler for academics to record and distribute information, from text book lists through to student marks. DAIS was later rewritten, and I'm sure significantly improved, to become FAIS.

Trevor's DAIS system ran on a local machine via a TCL/Tk user interface. When it came time to make it accessible through the Web, Bob Edwards recruited a student to make a web version in Java. After working on it for a while the student opined that it would have been much less cumbersome to code it in PHP. Bob took that advice and wrote the new FAIS (Faculty Administrative Information System) himself, in PHP. He also modified the Marker system to integrate with FAIS. FAIS now leaves the files in a state that the marker system can work on.

Academics I have spoken with seem to find FAIS very useful. If you're wondering why the central student admin system is unable to provide the FAIS functionality, I have been told that access is restricted and that response time may extend to minutes.

Richard Walker is a colourful character who worked on a PhD under Heinz Schmidt, then Leader of the Software Engineering Group in CSIRO DIT, while employed as a teaching academic in DCS. Richard is quite fluent in German and spent a year with the compiler group at the University of Karlsruhe. Because of the German connection, and because of his very efficient organising of the DCS cake club, he was affectionately known as the "Tortenführer" or "Kuchenmeister".

Richard was an admirer of fine typesetting and a LaTeX expert. He was also well known for his acquisitions. A pioneer in the use of Amazon, he received frequent deliveries of books. He bought the two-volume edition of the Oxford English Dictionary, including a magnifying glass permitting it to be read. He also bought a Japanese wooden flute and somehow obtained a large sofa for his office.

Since leaving DCS, he worked for a considerable time for Software Improvements and now works for the Australian Research Data Commons.

Graham Williams has had a 37-year on-and-off connection with DCS, which he recounts as follows:

> Travelling to Canberra by bus from Adelaide in 1982 I first explored ANU as a prospective PhD student, after completing undergraduate studies at the University of Adelaide. Meeting Robin Stanton, Malcolm Newey, Ray Jarvis, Terry Bossomaier, and Alan Snyder,

[33] I believe Steve Ball wrote the new version in TCL/Tk.

> and others, left a strong impression of the depth and excitement for AI at ANU. Foregoing Edinburgh, I began my PhD in the Computer Science Department in 1983. In 1985 I worked with CSIRO on a ground breaking AI system for bush fire management at Kakadu National Park. It encoded indigenous knowledge of cultural burning in an expert system I wrote in Prolog.
>
> In 1987 I implemented my PhD machine learning algorithm, based on ensembles of decision trees within a database system, and built one of the first expert systems in finance with ANZ bank in Melbourne. I returned to take up a lecturer position in DCS from 1988 to 1991 and have been involved as an adjunct since then. I joined CSIRO as Principal Research Scientist in 1991 and then in 2004 moved to the Australian Taxation Office as lead Data Miner and then the Data Science lead for the Whole of Government Data Analytics Centre of Excellence until 2015. Being a computer science adjunct allowed ongoing collaboration in research and teaching and supervising PhD students. Returning from overseas, where I was Director of Data Science with Microsoft, in 2020 I have come full circle to a new role as Professor and Chief Scientist with the Software Innovation Institute in the new School of Computing.

If this were not enough, for three years Graham was a visiting professor at the Chinese Academy of Sciences, spending one and a half months at a time in Shenzhen.

Graham recalls inspiring tutorials given by Andrew Tridgell for the first year introductory IT course he ran. He says that students loved it when Andrew pulled computers apart in front of them.

Bob Williamson joined ANU as a postdoctoral fellow in the Department of Systems Engineering, RSPhysSE, in late 1989, before becoming one of the foundation academics of the Department of Engineering, FEIT. He later became professor and head of CSL. From 2002 to 2006 Bob was director of NICTA's Canberra research laboratory and then NICTA's scientific director from 2007 to 2010, Interim CEO of NICTA in 2015 and Chief Scientist of CSIRO's Data61 from 2015-2017. Bob works in the area of machine learning; he was leader of the machine learning group at NICTA from 2010-2015, and he has been a Fellow of the Australian Academy of Science since 2012. Bob left ANU in 2020 and, from 2021, will be professor of the foundations of machine learning systems, at Eberhard Karls University of Tübingen, in Germany.

David Wolfram lectured in DCS between 1995 and 2000 and is currently a Visitor in the department in addition to his roles as founder and managing director of Ocelot Consulting. His interests lie in theoretical computer science and he chaired CATS (Computing: the Australasian Theory Symposium). He published a book entitled, *The Clausal Theory of Types*. At ANU, his office walls were adorned with the testamurs of his many degrees and fellowships.

Peter Wishart joined DCS in Jan 1980 to do a PhD under Malcolm Newey, but Malcolm was on sabbatical overseas at the time. Peter worked with Robin Stanton for a while, then did full-time tutoring. In 1981, he transitioned to the role of full time departmental programmer. Eventually he was managing a team of five programmers supporting teaching and research. He left in January 1988 to join the newly formed Software Development Centre in Canberra for NEC, which was developing imaging software for world markets.

> I remember a program called `Marker` that I designed and wrote based on a concept from Malcolm Newey. This would probably have been some time in 1981-2. Malcolm wanted to automate some parts of the process of assessing the programs written by students. So I wrote a preliminary version of a program which we tested against a number of assignments and evolved it before we decided to use it as part of the formal assessment process in a class. It was originally done to try and reduce plagiarism, which was becoming an issue in some classes. I think Marker was originally written in Simula 67, which was my preferred language on the DEC-10.
>
> Other things I was involved with at DCS:
>
> - commissioning one of first laser printers in Australia, the Canon LBP10;
> - installing and supporting the first Sun workstations;
> - setting up the first Sun user group in Australia;
> - installation and operation of the Pyramids used for teaching support;

- setup and operation of external email network through ACSNet – it was used by the whole of ANU at the time since it was the only way to get email out of ANU to Usenet and ARPANET.

Peter Wishart: The Canon LBP-10 Laser Printer

1982?: This was a new technology laser printer, one of the first in Australia, that we had to connect to our existing (old) environment of the DEC KA-10. Rob Ewin wrote the software interface which generated the Impress language to use the printer graphics and formatting capabilities. My job was to connect the printer to the KA-10. Seemed simple enough, but it wasn't! We had a "high speed" 9600 baud asynch line to connect to the KA-10. When we connected the printer, we couldn't get it to work unless we dropped the baud rate down very low, which made use of the printer problematic. We spent days with line analysers and tests to finally work out that the printer was sending its responses faster than the KA-10 could read them. The poor old KA-10 was just not able to handle "full speed" 9600 baud for even the few characters that the printer was sending as protocol acknowledgements. So, I re-purposed an old Data General Nova minicomputer which was not being used. It had no storage and no input devices. I wrote a simple machine language program which was input via the front panel switches. The program's sole purpose was to slow down the responses from the printer to the KA-10 by a few milliseconds. Debugging this simple program was messier that usual, because one of the front panel lights didn't work. The printer setup worked that way until ANU upgraded the KA-10 to a KL-10 and the 9600 baud line suddenly started working properly. This cobbled together solution typified the department's approach to making do with what we had.

Tom Worthington

In the mid-1990s there was buzz over a new technology called the World Wide Web. So I enrolled in a half day course at ANU on how to design a web page. Armed with this limited knowledge, I proposed that the Defence Department develop a web site, coordinated with other major government agencies. This was dubbed by the media as "The cabal that connected Canberra".[34]

In 1999 I was working at the Defence Department, chairing a committee writing computer policy. We were on our forty-second draft, when I decided it was time for a career change. I became an independent computer consultant. To keep up with computer developments, I wrote to all the universities in Canberra and asked who would like to have me around. Two universities had replied within a week and invited me in to discuss participation. However, ANU had already replied within five minutes of my message, simply saying an office was available and I could start as a Visiting Fellow the following week.

As promised, I arrived at ANU and found there was an office for me. I expected to be doing research, but five minutes later Brian Molinari came past, said "Hello Tom, know anything about ethics?" I explained that I had helped revise the Australian Computer Society code of ethics and he asked if I could provide some lectures. It was years before I worked out this was no coincidence: ANU needed someone to teach computer ethics and I was recruited for that.

My involvement at ANU has been much like that for 20 years, helping out at ANU with teaching, project supervision, the relationship to industry and a little research.

Lexing Xie is a professor in DCS and leads the ANU Computational Media Lab. Prior to ANU, she obtained electrical engineering degrees from Tsinghua University, Beijing and Columbia University, New York. She found the transition to New York something of a culture shock – her training in "silent English" were more than enough for reading and listening, but it seems impossible to organise thoughts on the spot and talk to people. [I have sometimes thought how difficult it would have been if I had had to learn Mandarin to complete my PhD and if the language of all major conferences and journals had been Mandarin.] Lexing says that Chinese food is better in Australia than in New York!

Her PhD spanned 2000–2005, including the dot com crash. She was in Manhattan at the time of the 9/11 attack on the World Trade Centre, luckily sufficiently far away. After graduation, she worked at IBM Research to the north of New York City for five years, carrying out a mixture of academic and applied research. An academic focus was extracting semantic concepts from video, and her team participated in the TREC Video Retrieval Evaluation (TRECVID[35] organised by the

[34] `http://www.austlii.edu.au/au/journals/CLCCommsUpd/1995/56.pdf` The cabal was essentially Tom Worthington (Defence), Tony Barry (ANU), and Kerry Webb (National Library of Australia).

[35] `https://trecvid.nist.gov/`

US National Institute of Standards (NIST). In this era before the ascendance of Deep Learning, the standard approach was to learn classifiers for concepts, and the big question was, "what concepts do we need?"

An applied project Lexing worked on at IBM was to develop a predictor for burst water mains for the local government authority in Washington DC. Her boss was presenting her prediction results to the clients when a pager went off. Yes, one of the predicted failures happened during the meeting!

Lexing later tried to understand what people do with images and videos. Her starting point was to crawl YouTube video to map the spread of visual memes among videos [ACM MM 2011]. In a few months, she amassed a dataset equivalent of everyone in IBM Research watching 5-6 hours of YouTube videos every day (behind a common corporate firewall) and received a letter from YouTube demanding that the data collection be stopped. The research lived on beyond her time in IBM though, as understanding collective behaviour and especially on YouTube has been a main theme of Lexing's lab at the ANU (see below).

When asked why she came to Australia in 2010, Lexing says, "foolhardiness!" (but it seems to have worked out okay). She saw an ANU advertisement on CRA jobs and successfully interviewed over Skype. Stephen Gould was also successful in that round. Her post-interview visit to ANU was her first visit to Australia. Her activities include both teaching and research. She has taught Statistical Machine Learning, Data Mining, Network Science, and Document Analysis. She has seen some of them grow from 40 odd to 300 students over recent years.

Lexing was also in charge of an exciting but very undersubscribed Vice-Chancellor's unit called "The Art of Computing", a first year unit designed to attract non-CS students. It made use of Snap!, a functional programming language in a browser. Students in the unit worked on remarkably advanced problems such as: solving sudoku, medical diagnosis with decision trees, biological sequence alignment, and sentiment analysis in tweets relating to hurricane Sandy. Students received a lot of support in tutorials, but largely because of the rigidity of ANU degree rules, very few students from other degrees enrolled. The first time it ran there were only nine students and two of them were children of DCS staff.

The Computational Media Lab[36] comprises around ten researchers at any given time. One of the main research themes is to formulate and solve machine learning questions for social media. Lexing claims that her team understands more about viral videos than anyone else, supported by four years worth of tweets and historical viewing statistics for millions of videos. They have found that it is possible to quantify potential popularity of videos, but the actual number of views is rather unpredictable due to the continuous influence from external events. However, the amount of time people spend on a video is predictable even before upload, based on poster's reputation, topic of the video, and a number of other factors. Lexing's team have also found interesting influences between videos and from external events. For example, the release of one Adele video, Hello, causes a rise in views for five other videos. Grammy awards, and release of videos by other artists can also cause surges in popularity.

4.1 Whatever Happened to ....?

A number of former staff members of DCS from long ago have moved on to senior appointments overseas.

Brad Broom held a joint Research Fellow / Lecturer appointment at DCS between 1984 and 1988, working on the ANU-Fujitsu CAP project. Brad has changed his research focus and is now a Professor in the Department of Bioinformatics and Computational Biology at MD Anderson Cancer Center in Houston, Texas.

Two of our former staff from the 1990s now hold chairs in universities in Hong Kong. **Qing Li** is now Chair Professor (Data Science) and Head of the Department of Computing at the Hong Kong Polytechnic University. **Jeffrey Yu** is now Professor and Chairman of the Department of Systems Engineering and Engineering Management at the Chinese University of Hong Kong.

[36] http://cm.cecs.anu.edu.au/

Mike Papazoglou taught in DCS from 1989 – 1992. His research areas include Information Sytems and Service-Oriented Computing. Mike was recruited from a senior position at the National German Research Institute of Computer Science GMD to support teaching Information Systems. ANU arranged for Malcolm Newey to interview him when he happened to be in Bonn in 1988.

Unfortunately, while he was at ANU, the programming group I headed managed to wipe out his hard disk contents during an upgrade and found that, due to a misunderstanding, the backups were empty! Despite this misfortune, he is now Chair of Computer Science at the University of Tilburg in the Netherlands, and Scientific Director of a European "Network of Excellence".

A poem written on the occasion of Richard Brent's 1998 departure for Oxford

Richard Brent is young and wise
And knows good ways to factorize.
Heaven help those banking eejits
Who put their faith in 90 digits.
In his hands such a small dec-i-mal
Is soon reduced to something primal.

Nothing disturbs our Richard's slumbers
But dreams of faster random numbers.
And in the daylight, a recurring trance,
"Randomness, don't leave to chance!"
Richard, what are you grinning for?
"I think I've saved a load and store."
Scientists, avoid disaster,
Invent your answers even faster.

Tasting the pleasures of Oxford's Hoare house
(Having fled the Acton poorhouse),
What memories you'll have of days of yore,
The smelly reek of burnt-out core,
And Univacs in elephant stables
And arsonists at work on cables!
Your goal complemented Zorba/Michael's
– To find a use for all the cycles.

At ANU we will mourn your going,
When at Henley we see you rowing.
But one big factor quiets our moan
– At least you'll leave our SPARCs alone!
No more screams of compute-starved scholars,
No more bids for computing dollars.
Gingold plans to sell some nodes
So ANU can pave its roads.

Chapter 5

DCS: Teaching

In writing about the history of ANU's computer science teaching, it would be foolish to try to document all of the different units taught in each of the last fifty years. Apart from the fact that the necessary factual information is sketchy, doing so would result in an unspeakably boring history. To misquote the so-called Hoare's Law, *Inside every complex history there is a simple one trying to get out.*

Instead, I focus on the changing forces which have shaped the curriculum, the ebb and flow in the importance of various CS topics, and the vigorous debates within the department about what students should learn. It is the case that the world of computing is continually evolving. New concepts and paradigms arise while the need for students to learn some older ones declines.

Inspiration for the evolving DCS curriculum has been drawn from the experiences and guiding principles from both Europe and the United States. One of the important sources is the ACM Curricula Recommendations (`https://www.acm.org/education/curricula-recommendations`).

My meandering story begins with how DCS initially positioned itself as an academic discipline and contrasted its course with service courses for scientists and public servants (Section 5.1). To frame the rest of the chapter, in Section 5.2 I look at the massive expansion in IT career possibilites, and in Section 5.3, how ANU graduates are likely to make use of (or otherwise benefit from) what they learn in DCS units.

Section 5.4 tells the story of what has been considered to be foundational computer science, while Section 5.5 lists other important CS topics which have been taught over the years. The relationship between Software Engineering and Programming is, and has been, the subject of considerable internal debate. A different approach to training computer professionals – the way medical doctors are trained – is outlined in the section on the Software Innovation Institute (Page 136). Software Engineering has a section to itself (Section 5.6) and some forthright views are expressed. For the record, Appendix E lists the many degrees which DCS has taught into over the years.

Three topics related to the curriculum, namely student evaluation of teaching, the Computer Students Association, and plagiarism are covered in Sections 5.10 – 5.12. Finally, I conclude (Section 5.13) with a discussion of important curriculum related changes which have occurred over the decades.

5.1 1971: Positioning Computer Science in the Academic Firmament

The teaching of undergraduate CS came about essentially as a result of student demand and was opposed by many in the university who thought that computer science was not a proper academic discipline, but rather in the category of *hamburgerology* and *washing machine science*, allegedly taught in bachelor programs at some third-rate American universities. The first CS academics at ANU were very concerned to establish it as an academic discipline and to distinguish it from pre-existing service courses which taught FORTRAN as a tool for physicists and foresters. They were also keen to distinguish the ANU CS programme from the Programmer-In-Training (PIT) course at the Canberra College of Advanced Education (CCAE, now University of Canberra). The latter was a graduate diploma course which taught public servants about the machines, programming languages and methodologies then in use in the Commonwealth Public Service.

ANU aimed to give its students a deeper understanding of underlying CS concepts and claimed that ANU graduates would write better programs, and be better equipped to deal with future changes in computer hardware, programming languages and methodologies.

Reflecting upon my 46 years of programming (or was it software engineering?) post-graduation, I realise that, in my case, the DCS aim was realised, even though I have written the majority of my code in low level languages such as M68000 assembler and C, and in hacker languages like Perl. Even in the gutter one can take inspiration from the stars.

In designing a curriculum, it is surely worth considering the range of roles our graduates are likely to assume in their future careers. In Section 5.3 I outline some role stereotypes to shed light on this. First, let's think about the massive expansion of possible roles for our graduates since 1971.

5.2 Expansion of Possible IT Careers Since 1971

In 1971, graduates of computer science who went on to be programmers, would almost certainly have been programming mainframe or mini computers, most likely for business (payroll, general ledger, etc.), military (e.g. missile guidance), or scientific or engineering applications (such as calculating orbits or designing bridges).

In 1971 a few computers were configured with real time operating systems[1] and were fitted with hardware interfaces to laboratory equipment or to process controllers.[2] However, the use of computers as controllers was only just beginning.

In 1971, computers had only recently begun to be embedded in vehicles. One example was the Apollo Guidance Computer used in the NASA moon missions from 1966.[3] Another application was flight control of military aircraft. In 1970, Garrett AiResearch completed a MOS-based chipset to compete with the electromechanical flight control computer for the US Navy's F-14 Tomcat fighter. According to Wikipedia,[4] it was 20 times smaller and much more reliable. This MOS chipset may be considered the first microprocessor, but the Intel 4004, the first true microprocessor built on a single chip, was not available until November, 1971.

In 1971, control systems for cars were very simple. There were no onboard computers. Now there may be dozens in a single car. Indeed, the 2021 Tesla Model S includes a built-in gaming system with 10 teraflops of GPU power! Computers now control rail systems, driverless trains, traffic signals and factories. Commercial aircraft operate under computer control (autopilot) for most of the flight.

In 1971, there were no mobile phones, no digital cameras, no point-of-sale (POS) terminals, no EFTPOS, no automated teller machines (ATMs), no barcode readers in supermarkets,[5], no RFID tags, no games consoles, no driverless trains, no websites.

Software libraries associated with programming languages were impoverished compared to today. There was virtually no electronic text, virtually no online images or video, and almost no user interaction data. There was very little interconnection of computers, and the bandwidth of such connections as there were was pitifully slow. There was no cloud, and no "online" services.

Fifty years later, all these things are commonplace. It's fair to say that computer science graduates in 2021 have a rather extended range of possibilities for what they may be called upon to program!

5.3 To what End Are We Now Teaching?

The population of ANU graduates who study one or more units taught by DCS goes on to a very diverse set of careers. I have no hope of describing or even enumerating all of them here. However,

[1] The RT-11 real time operating system for the PDP-11 was first implemented in 1970. https://en.wikipedia.org/wiki/RT-11

[2] As noted elsewhere, ANU acquired an IBM 1800 Data Acquisition and Control System (DACS) for Nuclear Physics in 1968.

[3] https://en.wikipedia.org/wiki/Apollo_Guidance_Computer

[4] https://en.wikipedia.org/wiki/Microprocessor

[5] Wikipedia says that the first scanning of a Universal Product Code (UPC) barcode was in 1974.

to help understand why staff in DCS have vehemently held very different opinions about what our students should learn, below is my attempt to list of stereotypical roles which require (or should require!) at least some computer science training. Of course an individual professional may play multiple roles either simultaneously or across their career.

Computer scientist: Advancing science related to computing, through experimentation, innovation or theoretical advances. The boundaries of computer science are highly debatable, but I would include algorithms, programming languages, software engineering, program verification, models of computing, image processing, machine learning, computer vision, robotics, computer architectures, databases, information systems, human-computer interaction, and automatic theorem proving.

R&D engineer (software): Inventing, evolving, designing, creating and/or evaluating prototypes of software products, packages or services. This role is typically found in tech companies, or small start-ups.

R&D engineer (software + hardware): Designing and developing controlling software for electromechanical devices, vehicles and robots. Real-time programming, often for constrained hardware resources. This role is typically found in tech companies, or small start-ups.

System and Network Administrator: Installing, configuring and updating operating systems and software packages; network expansion and management, patching, enforcing an organisation's policies, maintaining security, managing licences, recommending and procuring hardware, software and network upgrades.

Web implementor: Writing or modifying HTML, style sheets or scripts; designing web page layouts; interfacing to applications and databases.

Presales: When a vendor is selling high value hardware, software, or computer services, presales technical support is likely to be needed in preparing a realistic, low-risk bid, in making technical presentations, in answering client questions, in setting up demonstrations, and/or in running benchmarks.

Computational scientist: Computer simulation and modeling plays a critical role in many traditional sciences, such as physics, chemistry, geology, sociology, environmental sciences. Amanda Barnard (see Page 105) lists the following among the things she recommends that a computational scientist should know: Numerical methods, C++ and Python, parallel programming basics, and software development (version control, etc) and high performance computing (including profiling, etc).

Modeling engineer: A person employed to model power distribution grids, wastewater systems, or climate change may be able to use an off-the-shelf package, but will surely benefit from an understanding of computer science concepts and methods.

Security analyst: Vetting software and online services for vulnerabilities requires a special type of thinking and a high level understanding of programming languages and runtime systems. Writing secure applications is even more difficult!

Developer of software for safety critical systems: Special skills, rigorous procedures, and a strong nerve are required in this role. A high level understanding of the relevant programming language and its runtime system are required.

Commissioner of major software development projects: An employee of a government agency, or large corporation, who is responsible for issuing tenders or letting contracts for a major software project must surely benefit from an understanding of computer science concepts in determining whether a bid is likely to succeed.

Project manager of a major software development project: (government, corporate or military) Unsurprisingly, project management roles, such as project planning, staffing, decomposition into sub-projects and activities, resource allocation, timeline planning, cost estimation, client liaison, and management of contract variations are critical. It is vital that requirements are mutually understood, and it must surely help that the project manager has enough understanding of computing to be able to tell when estimates for sub-projects are realistic.

There is no possibility that a three or four year CS degree could possibly inculcate all the knowledge and skills needed for all these roles. The reality is that graduates will start at the bottom of the employment ladder, learn on the job and through subsequent external training, and rise to positions of greater responsibility. It is also true that lack of significant CS training may not be an impediment to high level IT appointments. I know of two ANU alumni who eventually became CIOs of large government agencies – one had a Computer Science Honours degree, the other had no computer science training at all! Another achieved a very senior IT position in a government agency, having passed only the introductory CS unit.

As far as I know there is no comprehensive survey of our alumni which would tell us much about their subsequent careers. Thinking about employment prospects in the local region might suggest a strong bias toward jobs in or associated with the Australian Public Service, but in recent years, a very high proportion of students are from overseas. The profile of employment for them is likely very different.

Having embarked on a rigorous statistical study with $n = 1$, I report the experiences of a representative, first-class honours, University Medal student who is Chief Scientist at Nuix, a "unicorn"[6] company which conducted a very successful IPO on the Australian Stock Exchange in November 2020. David Sitsky was a DCS student between 1990 and 1993.

> My first experience of programming was BASIC on an Atari 2600 gaming console. I then progressed to BASIC on a Tandy MC-10 and then finally a TRS-80 machine. I had the opportunity during primary school to go on a school holiday "coding camp" which used Microbee computers, again using BASIC. I attended Phillip College for years 11 and 12 where we used Turbo Pascal on Window PCs.
>
> Coming to DCS was a significant change as in first year we started with "green screens", a Unix (SunOS) shell and learning Modula-2. While the current generation can't live without instant messaging applications, we already had the `talk`, `write` and `wall` utilities for this purpose. In 2nd year and later we had access to proper X-terminals running SunOS and had access to a window manager environment.
>
> I really enjoyed the undergraduate units. In particular I have fond memories of Brad Broom teaching the first year advanced class, Chris Johnson teaching networking, Malcolm Newey teaching functional programming, Brian Molinari teaching software engineering as well as Robin Stanton on logic programming with Prolog, and finally a great operating systems unit using `minix`.
>
> The honours level units were of a very high standard and it was great to split 50% of our time on research and the other 50% on course work. My research project was investigating the use of functional languages applied to scientific programming problems. During honours year each student was granted a dedicated X-terminal for our work in a special room. Sometimes we played the multi-player X-pilots game to relax.
>
> [**Dave:** After graduation Dave worked on the ANU-Fujitsu CAP project – he describes his experiences there on Page 218.]
>
> My time ended at the CAP Research Program by the end of 1998. I had a brief stint working at Cisco Systems in embedded operating system work with their IOS product, before I had a call from Peter Bailey about a new software startup called Nuix.
>
> Apart from Peter, many other ANU people have worked at Nuix, including Nigel Snoad, Raj Nagappan, plus two others who, at the time of writing, still work there: Luke Quinane and Daniel Noll. All of these individuals have played an important part in Nuix's history.
>
> My initial goals at Nuix were to write software that was used extensively in the real world. After a bumpy start, we found a niche in being able to search over large amounts of email data in various formats (proprietary and open) for investigation purposes. As the sizes of data became larger and multi-core machines started to become common, I was able to draw on my experiences from the CAP research program in parallel computing and apply it to this problem domain (using JMS – Java Messaging Service instead of MPI). This eventually lead to what has become known as the "Nuix Engine" – a platform for processing massive amounts of "unstructured data."
>
> The early years were spent on-site with customers to understand their specific problems as well as developing and supporting the software. As Nuix grew, I found myself in positions of

[6] A company with a market capitalisation in excess of one billion US dollars.

> leadership both at the engineering level and as a part of the Nuix executive team. I still maintain a close pulse to the code and still work on it closely. My position in Nuix now is in "strategic R&D" (Chief Scientist) to work on future ideas such as making the engine "cloud native" and exploring new methods to apply over human generated text within unstructured data.

A chat with Dave Sitsky

Dave H: How well did what you learnt in your CS course prepare you for the work you ended up doing?
Dave S: All units I believe provided a solid foundation in computer science such that as new technologies / ideas came in the industry it was not difficult to adapt. Functional programming seemed esoteric at the time but now is used extensively in major programming languages. I really liked everything that was taught.
Dave H: Which units were most useful?
Dave S: I think the operating systems unit in Minix was for me the most useful since it exposed me to a large well-structured system written in C. The CAP research program reinforced this further.
Dave H: Were there any serious gaps in what you learned?
Dave S: I know things have changed significantly since I was a student, but I think more exposure to large-scale systems, particularly in 3rd year would be beneficial. Learning how to work effectively in teams with appropriate toolsets. A much bigger focus on automated testing, as well as CI/CD (Continuous Integration / Continuous Deployment) as a part of development. We use Jenkins for continuous testing, and JIRA for bug tracking. Also Confluence as our "wiki / knowledge base" and Slack for communications.
Using open-source software is very common in many jobs. Learning how to use and contribute back I think is valuable.
Finally there is the interaction with people aspect as well which is a large part of real-world work. A short unit on listening, collaboration / teamwork and personal organisation methods I think would be good. A study on different development methodologies (agile, etc) would also be good. When working in teams, doing peer code review would also be valuable.

One thing I draw from Dave Sitsky's account is that the purpose of an education is not solely to provide knowledge and skills for a career. A computer science (or software engineering) education should exercise the mind, nourish the intellect, inspire the imagination, and enable a graduate to understand and think critically about developments and opportunities beyond their specific role.

The next major section describes the teaching of what could be called the foundations of computer science.

5.4 Foundations of Computer Science

At the core of the discipline of computer science are computing machines, algorithms, programming languages, data, programs, communication, and human-computer interaction.

The practice of computer science requires the writing of programs for real machines, which in turn requires learning how to represent data, how to choose an appropriate algorithm, how to express it in a programming language, and how to get it to run on a real machine. Ensuring that the program will operate on the available machine, and that it runs in an acceptably short time may require an understanding of the architecture and instruction set of the machine in question. Furthermore, the programmer is soon forced to face the issue of program correctness, which naturally introduces the need for an unambiguous specification of required behaviour, and for effective debugging techniques, testing procedures, and (ideally) proof of correct behaviour.

Each element of programming skill stands to benefit from the theoretical base of computer science, which includes logic, computability, abstract computing machines, language grammars and their recognisers, algorithmic complexity and efficiency, and program verification. Chris Johnson outlines some DCS attitudes to the role of theory.

> DCS for a long time took the view that some formal thinking in educating students was beneficial (you will remember Ray's algorithms, Brendan, Brian, Robin, and then Malcolm and Mike

> Robson). Initially, this was just some mathematics, preferably discrete maths – even though mathematical notation was not directly applied anywhere later in the course, we collectively believed and told prospective students that the logical and semi-formal thinking was desirable as a way to have student programmers think outside the assertive giving of orders by a program. At that time some rigour in thinking was considered desirable, beyond acceptance of your program by the compiler and blundering your way to a correct program.
>
> ... COMP1013, Introduction to the Science of Computation, was a hard ask for many students and was a hurdle to completing degrees for some. Discrete maths, propositional calculus, pre- and post- conditions, loop invariants, termination, correctness.... This became more explicitly developed into a second year subject in the Bachelor of Software Engineering program, with some introduction to formal methods, predicates, and elementary program proof.

An explanation of ANU unit numbering

For the first decades of ANU CS teaching unit codes comprised one letter and two digits, where the letter denoted the year in which the unit was taught, 'A' for first year, and so on. Thus, `Computer Science B01` was a unit taught in second year.
At some point the codes changed to comprise a 4-letter code representing the subject, a digit representing the year, and three extra digits. Thus, `Computer Science B01` might become `COMP2001`.
A long-standing ANU convention is that a full-time student load comprises four units per semester for first year students, and three units in later years.

5.4.1 What to Teach First

In the 1970s, it was generally accepted that an introductory computer science unit should teach the concept of algorithms, the specifics of programming in a suitable programming language, and the process of running programs on a computer. In completing programming exercises, students were exposed to compile-time and runtime errors and to the process of debugging. In the first few years at ANU, the introductory unit also covered numerical methods. Indeed, for the first few years, with FORTRAN + numerical methods, Computer Science B01 was indistinguishable from a service course.

Over the years, what language to teach in the introductory unit has been a matter of vigorous debate. This applies not only at ANU but around the country.

As noted, DCS started with FORTRAN as the introductory language, but moved to Algol W after a few years, and then to Pascal, Modula-2, Eiffel, and Java. Finally, in recent years, Haskell has assumed the mantle. Haskell is a purely functional language and it's a very interesting choice, in stark contrast to Melbourne University which chooses C as its introductory language.[7] Clem Baker-Finch, who came to DCS from the University of Canberra, was a Haskell enthusiast. Jeremy Dawson remarks that ANU switched to Haskell as UNSW switched away from it.

Quoting Chris Johnson:

> The choice of introductory language was always driven by the idea of introducing clearer, abstracted thinking and expressive power in the introduction (e.g. structured programming, object orientation, logical assertions, functional paradigms) before students were led to learn and use industrial working languages later in the course.

One argument in favour of using a less common language is to equalise levels for students, and to thereby reduce attrition. In 1996, Steve Blackburn conducted a study of student attrition from DCS. He found that some students, particularly women, found themselves at a major disadvantage in introductory CS because other students had already mastered the introductory language prior to coming to university. Because of this, a significant number of students felt that they lacked the ability to do computer science, even though they did well in other subjects.

[7] The lecturer, Alistair Moffat, maintains student interest with *Computer Science Karaoke*. See `https://people.eng.unimelb.edu.au/ammoffat/mary-movie.mp4`

Facilities and languages used by introductory Computer Science Units, 1971–1990

Years	Language	Medium	Computer	Operating system
1971-74	FORTRAN	Cards	DG [Super]Nova	DOS
			Univac 1108	EXEC-8
1975-76	Algol W	Cards	IBM 360/50	MPFT
1977-79	Pascal(DIKU [a])	Cards	Univac	EXEC-8
	Pascal	Cards	DG Nova	RDOS
1979-80	Pascal	Cards	DEC KA-10	TOPS-10
1981-82	Pascal	PASTIE, VC404 trms	DEC KA-10	TOPS-10
1983-84	Pascal	RUFUS, Menu, VC404/VT100 trms	DEC KL-10	TOPS-10
1984-88	Pascal	VT100	DEC VAXCluster	VAX/VMS
1989-90	Modula-2	VT100	Sun 4	SunOS

Table compiled by Chris Johnson.

[a] DIKU = Datalogisk Institut, Kobenhavns Universitet

When the teaching of first year computer science commenced, the introductory units (then A01 and A02) became only a quarter of a full-time load rather than a third. Malcolm Newey recalls that the subsequent replacement of A01/A02 with four first year units (A11, A12, A13, A14) was a major step forward for the whole CS program.

After completing the first computer science unit a computer science major goes on to study some or all of the foundational topics described below.

5.4.2 Data Structures and Algorithms

Creating high quality programs to solve problems requires choosing the best combination of data structure and algorithm. Naturally the two are inter-dependent, as implied in the title of Niklaus Wirth's book (first published in 1976), *Algorithms + Data Structures = Programs*, which was used for many years as a text in DCS units. The right combination of data structure and algorithm can result in a simpler and more elegant program.

Choice of algorithm can also be critically important to achieving a program which runs efficiently. Differences between one algorithm and another can have huge econonomic ramifications. During my time working on the Bing search engine, an email came around from our indexing team: "Due to improvements in our algorithm, 11,000 servers are now available for redeployment to other uses within the organisation." That improvement would have saved tens of millions of dollars per year! Teaching of algorithmic complexity and efficiency is discussed in Section 5.4.3.

Runtime libraries associated with languges such as Java, C#, and to a lesser extent, C provide implementations of many common data structures, and of important algorithms. This language support may significantly reduce the amount of human effort required to write a program, but it's at the expense of understanding what's going on behind the scenes. Even when using built-ins, a programmer with solid knowledge of fundamental data structures and algorithms may be able to realise when the built-in is using more time or memory than it should, or why it isn't operating correctly. If performance matters, hopefully someone trained by DCS will have the courage to look under the hood and either rectify the problem or find an alternative solution.

Back in the 1990s, I was using a Unix implementation of `qsort()` to sort millions of floating point numbers. In an effort to save memory I quantised the numbers down to small integers and discovered that the sort then took most of a day instead of about a second! Yes, quantisation created massive numbers of equal values, triggering worst-case $O(N^2)$ performance from the algorithm used in that `qsort()`. A remark from Andrew Tridgell led to me implementing and using *counting sort*[8] which, by avoiding comparisons, runs in $O(N)$ time.

Units on algorithms have been taught by many people over the years, including Richard Brent,

[8] https://en.wikipedia.org/wiki/Counting_sort

Brendan McKay, Mike Robson and Weifa Liang. Chris Johnson tells me:

> By 1998 the goals of the course were identified as constructive, theoretical and pragmatic. The construction of program systems and study of data structures were separate half-courses in second year, which later combined into a third semester course, and later some into the second semester, while algorithms moved into third year with more formalism. The emphasis shifted from the study of a small library of well-worked data structures and their performance, and by 2010 to the larger pragmatic design, implementation and abstraction issues as represented by design patterns.

5.4.3 Complexity and Computability

Following on from introductory units on data structures and algorithms, are more advanced units on complexity and computability. Material includes big O notation, solving recurrences, time and space complexity, the halting problem, and formulae for the average and worst case running times of the algorithms studied. Brendan McKay recounts some of the history:

> When I came to ANU in 1983, Mike Robson was teaching a complexity unit. Not long after that, I took it over and taught it for quite a while. There were two parts, one being complexity of algorithms (NP-completeness and such) and the other was computability (the limits of computation, unsolvable problems, etc). These got separated at some point and I remember that for a long time there was a formal languages unit (also taught mostly by me) with regular languages, context-free languages, and Turing machines. That touched on computability but there wasn't much time for it.
>
> ... the main outlet for the theory was the third year algorithms unit (COMP3600) that was variously taught by me, Weifa Liang, or visitors (Jeanette McLeod and Beáta Faller were two of them). That was followed by a fourth year algorithms unit (COMP4600, Advanced Algorithms) that I also taught until Weifa took over. In the best years of that, modules were taught by random people who came along to teach in their specialty. For example, Richard Brent gave some lectures on computational number theory.

John Lloyd contributed to the teaching of COMP3600 and says that Hanna Kurniawati is teaching it now. Jeremy Dawson has also been involved in teaching this type of material.

Chris Johnson's thoughts on computability and complexity

The idea that some things might not be computable – even if it is very hard for a practising programmer to identify them – is valuable. The idea of complexity related to computations – the elementary idea that a method of $O(N^2)$ may be a worse solution than $O(NlnN)$ – must also be implanted, even though the effect may not be measurable at the scale of data used in student work (perhaps it should be.)
As for "complexity" as a concept – that is needed; but computational complexity as a practical test of what is being computed is relevant only in large scale systems ... but no longer significant in parsing in compilers, for instance ...

Chris Johnson: Theoretical computer science in popular culture

The concept of $P = NP$[a] did provide me some amusement last week while I was watching an episode of *Elementary* from 2014, a TV program modernising Sherlock Holmes – where the possibility of a proof or otherwise that $P = NP$ was part of the plot driver. Knowing the answer allegedly allowed a break into all encryption systems. Quantum computing might have been more plausible in this context but would have been even less understandable by the writers and audience, but. It's interesting in the depiction of character in the series that despite his extreme intellect Sherlock himself claimed no knowledge of the maths but relied on a semi-tame professor of computer science.

[a] See `https://en.wikipedia.org/wiki/P_versus_NP_problem`

5.4.4 Computing Machines

In the very early years of DCS, Ray Jarvis's Computer Science B02 covered theoretical topics such as finite state automata, ones and twos complement arithmetic, the construction of arithmetic units from simple logical gates, and Boolean algebra. Ray Jarvis also presented the proof that the ideal base for numbers in computing hardware was as close as possible to e, i.e. binary or ternary. He introduced the fundamental concept of computability through a study of Turing machines. The idea that anything which can be computed, can be computed by a Turing machine with infinite tape and infinite time, is very powerful. It leads to thinking about other computing machines, and machines which can be proven to be equivalent to a Turing machine. To me these ideas seem important to keep in the back of one's mind. Not everyone agrees:

Clive Boughton's thoughts on Turing Machines

I think we have to recognise that 'computing' may now not mean what it meant when it was first coined. The Turing Machine, whilst informative for earlier 'computing' models, is now largely irrelevant. Nonetheless, it would be interesting to show how and why the Turing Machine is no longer relevant. Bohr's theory of the atom falls into the same category in Physics.

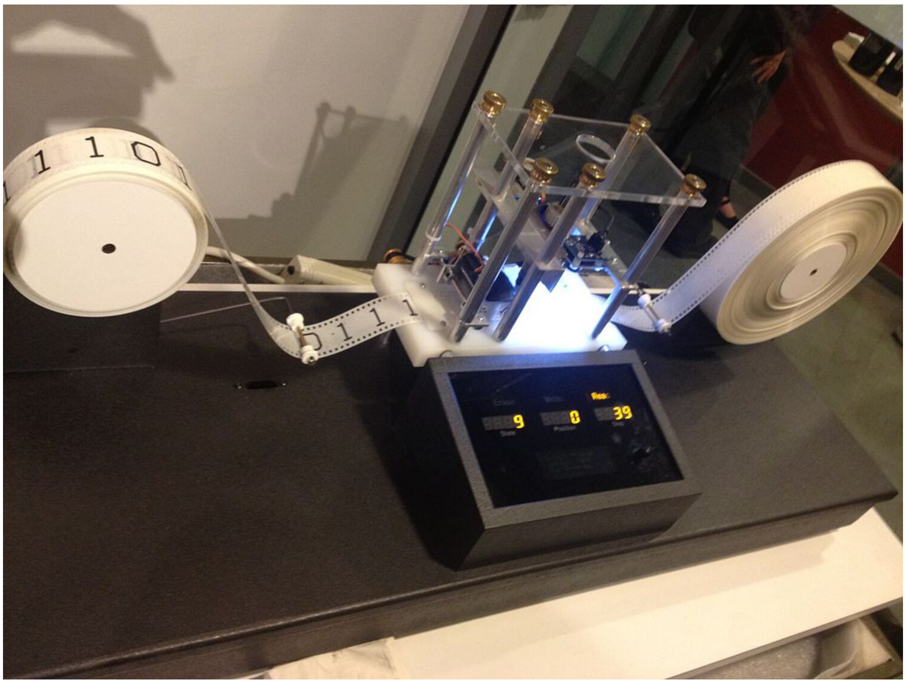

A model of a Turing machine created by Mike Davey for the Go Ask Alice exhibit at the Harvard Collection of Historical Scientific Instruments.[9] *Photo: GabrielF.* `https://en.wikipedia.org/wiki/Turing_machine#/media/File:Model_of_a_Turing_machine.jpg`. *Creative commons:* `https://creativecommons.org/licenses/by-sa/3.0/deed.en`

Chris Johnson also notes a decline in importance of abstract machines

The finite state mechine as a design notation, and exercises in finite state machines as recognisers, have lost relevance to most working programmers, as user input is often no longer entered as a stream of characters but via structured user interfaces. These ideas were originally part of the suite of first year subjects, then remained as a smaller component in the subject on computer organisation.

In the early years, an ANU student working in the public service reported back that their understanding of Boolean algebra and the minimisation of Boolean expressions enabled them to simplify a massively contorted COBOL (a programming language never taught at ANU) conditional from something like nine lines down to one. Their colleagues had great difficulty believing that the simplified version was equivalent to the complex one, even though the program was simpler, easier to understand, and produced identical results.

[9] See `http://aturingmachine.com/`.

By the mid-1980s the theory of computing and the organisation of computers could be presented as a coherent layered view, captured by Brian Molinari's 1985 textbook for these units (*Introduction to Computer Systems: a User View*, Cambridge University Press).

DCS students are of course given the opportunity to work at low levels with real and simulated machines. Mention has been made of assembler programming on the Nova and simulated Nova, of the assignment on Nova interrupt handling, and of Malcolm Newey's unit working with PDP-11s. In the 1980s Brendan McKay developed a machine simulator called A02SUCS. It, and its successors have underpinned DCS teaching of low level programming up to the present day, as described in the following panel:

The history of A02SUCS / PeANUt / rPeANUt

A02SUCS was a machine simulator and assembler written by Brendan McKay in Pascal for use in Computer Science A02. It and its successors were used for more than a quarter of a century, from the mid 1980s to 2016. Peter Strazdins says that it was a very sophisticated system – it eventually supported both module linkage (due to Brendan) and virtual memory (designed by Steve Blackburn).

Peter Farmer created the GUI interface for the Pascal version using Solaris CDE (Common Desktop Environment. He says, "In hindsight a bloated, proprietary dead end, the cross platform TCL/Tk replacement was a much better approach. Would have loved to have built that one. Funnily I must be one of the few people that didn't think A02 sucked – I loved doing that sort of programming – closer to the hardware. Would have considered Brendan's assignments a real treat!"

Drew Corrigan converted the system to Sun Modula-2, and provide a new graphical user interface. Unfortunately, Modula-2 included no definition of a standard set of libraries and Sun dropped support for Modula-2, resulting in a further port to Gardens Point Modula (GPM, from QUT).

Brendan still has an A02 assignment in Modula-2 from 1989, which required students to implement A02SUCS Model 3 by modifying A02SUCS Model 2. The changes were to add a register, replace autoincrement mode by indexed mode, add an extra instruction format with six instructions, and add corresponding assembler syntax.

Later the design of A02SUCS was modified (made worse, according to Brendan!) and renamed as PeANUt. Peter Strazdins tells me that Robin Stanton was responsible for the name change, with SPARClet another candidate.

Later a group of senior undergraduates (Mark Dickson, Brendan Humphries, and Paul Marando) rewrote PeANUt in C, using TCL/Tk and Tix. This version lasted from about 1996 to 2007.

When Peter Strazdins took over the unit again in 2008, he found that, due to a change in the Tix interface, the PeANUt window wouldn't come up. After making the necessary modifications, and correcting a few buffer over-runs, PeANUt rode again.

Interestingly, one student then ported PeANUt to Windows, and another to FPGAs!

In 2011, Eric McCreath took over the unit and wrote a new simulator called rPeANUt in Java. It emulated a RISC machine and featured a GUI, a terminal, a frame buffer with small black and white display, a simple interrupt system, and a cache simulator.

5.4.5 Programming Languages and Runtime Structures

The model of computing presented by early versions of FORTRAN essentially assigns variables to numbered memory locations when the program is compiled. Ignoring I/O, statements in the program take values out of memory locations, perform arithmetic and logical operations, and store results back in memory locations. Subroutines can be compiled separately and have distinct name spaces.[10] The link-loader ensures that variables in the main program and sub-routines do not overlap unless specified by the programmer. Parameters are passed to subroutines by transferring the memory addresses of variables.

Units taught by Robin Stanton on his arrival in 1973, presented more sophisticated languages and the runtime structures needed to support them. Chris Johnson says: "The crucial difference was that the programming language was used to express an algorithm, with its own model of storage and execution that was closer to human thinking and further away from the computing hardware."

[10] I.e. a variable X in one module refers to a different memory location than a variable X in another module.

Algol 60 is able to support recursion because, whenever a procedure (subroutine) is called, a new memory space (a frame) for it is created and pushed onto a runtime stack. When the procedure exits, the stack frame is popped. Algol 60 is a block structured language – a single program unit (main line or referenceable procedure) can be divided into blocks, and variables can be defined as local to a block – not referenceable from outside the block.

Algol 60 is a strongly typed language. Algol 60 programmers are required to declare all variables before use, specifying the type of each variable. Once declared, the type of a variable is fixed at compile time and cannot be modified. In contrast, FORTRAN allowed two variables of different types to be EQUIVALENCEd to the same memory location, or overlayed in COMMON storage.

Algol 60 formalises the semantics of passing parameters to procedures. If you pass an expression by value, the expression is evaluated before the procedure is executed and the value is assigned to the corresponding formal parameter. If the same expression is called by name, then a procedure called a 'thunk' is generated to evaluate the expression. Each time the formal parameter is referenced in the procedure, the thunk is called to evaluate a new value. This was a much slower, but more powerful, mechanism than the call-by-reference in most implementations of FORTRAN.

To quote Tony Hoare, "Here is a language [Algol 60] so far ahead of its time, that it was not only an improvement on its predecessors, but also on nearly all its successors."[11]

Recursion was an important step forward because many mathematical models are naturally expressed in recursive forms. Mind you, some students struggled to understand recursion. My office, as a tutor, was often covered in little scraps of paper, which I had used to try to explain how a runtime stack supported recursion.

Robin also presented LISP, a language in which everything is represented in the form of dynamically created lists. In 1973, it fascinated me that the minimal form of LISP originally formulated by John McCarthy (`car`, `cdr`, `cons`, `atom`, `eq` operators) can be proven equivalent to a Turing machine. Two critical items of machinery needed to support LISP (in practice) are a memory allocator and a garbage collector. We studied different forms of both memory allocators and garbage collectors. These topics have been of lasting interest given that far more recent languages such as Java rely heavily on them. Mike Robson and Richard Brent have both published research in the area of memory allocation, and Steve Blackburn and Tony Hosking have researched extensively in the area of garbage collection.

Another language and runtime system introduced by Robin was Simula 67, created by Ole Johan Dahl and Kristen Nygaard. In order to support discrete simulations, Simula 67 introduced the concept of class. Classes were data structures representing objects, packaged up with the procedures to operate on them, and executable code to enable them to operate as "agents". Class inheritance was also introduced, allowing you to define a class `truck` which inherits the attributes and operations which apply to a parent class `vehicle`. Classes became the foundation of object oriented programming, which became a very important paradigm, and was the natural way to implement many graphics applications. Steve Blackburn and Robin Stanton collaborated with researchers from St Andrews university, led by Ron Morrison, to address persistent objects – objects whose lifetimes extended beyond a single program execution.

Simula 67 also introduced the concept of coroutines. In Algol 60, once the execution of a procedure instance stops, the instance is destroyed and control returns to the caller. In Simula 67, an executable class can relinquish control using the `detach` primitive. When it does this, its frame remains in existence, and its execution may later `resume` where it left off. In 1977 Chris Johnson wrote an authoritative technical report on the forked-stack runtime structures used to support this model of computation.

Algol 68 was intended as a successor to Algol 60, but the definition of the language became very complex, and compiler implementation was difficult. Prominent European computer scientists, notably Edsger Dijkstra, Tony Hoare and Niklaus Wirth, preferred a simpler direction.

Niklaus Wirth developed a simple Algol-60-like language, incorporating records, which are like

[11] *Hints on Programming Language Design*, Dec 1973, `https://www.cs.yale.edu/flint/cs428/doc/HintsPL.pdf`

Simula 67 classes but without the procedural attributes and without the ability to act as agents. His extension was known as Algol W and, for a couple of years, it replaced FORTRAN as the language of Computer Science B01. John Hurst motivated the change by arguing that it provided a better vehicle for teaching structured programming and top-down decomposition of algorithms, as well as basic data structures.

In 1970,[12] Wirth took ALGOL W ideas still further in Pascal, a language which extended the number of primitive data types and also the mechanisms for creating compound data types. Pascal soon replaced Algol W in introductory computer science teaching. Pascal was followed by Modula and Modula-2, both designed by Wirth. Modula allowed decomposition of a program into independently compiled modules, and Modula-2 extended that to support parallelism through coroutines. It is a language designed for the writing of operating systems. Modula-2 was used in DCS teaching for many years, and was used in the KRIS robotics project in the late 1980s.

2010: A memorable guest lecture on programming languages

For the record, a brilliant lecture of the history of programming languages was delivered at the Haydon-Allen Tank in Dec 2010 by Guy Steele. While doing his PhD at MIT, Steele co-authored more than 20 papers on LISP and its implementation, and co-designed the language Scheme, which was taught for some years in DCS. He then worked on a C compiler for Tartan Laboratories, before joining Thinking Machines and developing parallel versions of both LISP and C (*LISP and C*). He co-wrote the book I have referenced most in my lifetime: *C: A Reference Manual*

Interest in programming languages continues unabated. The SAPLING group (Sydney Area Programming Languages INterest Group)[13] has been holding annual meetings since at least 2006. Its steering committee includes Steve Blackburn (ANU) and Tony Sloane (formerly of ANU).

5.4.6 Language Recognisers and Compilers

Peter Bailey's recollection of second year computer science in 1988

I did B11, B12, and B13. The lecturers were Malcolm Newey, Brian Molinari, and Chris Johnson. Chris did the Architecture unit (B13). B11 and B12 covered algorithms and data structures, and we had to learn a lot of programming languages. My memory is failing me which semester it was but my distinct recollection is that in one of them we had to learn four different programming languages in the course of the semester and implement things like parse trees and interpreters. I think the languages were LISP, Scheme, possibly Prolog (which I confess to still struggling with), and I think Modula-2. We'd done Pascal in first year. It was kind of a crazy time, but after that, I was never scared of learning a new programming language for professional reasons. Not only had to learn them but to figure out how to use them for something meaningful and get an assignment in, in very quick time. That particular set of units, those who got through them and enjoyed them was a much smaller set of people than had started.

Study of programming languages and their models of computation, plus the runtime structures needed to support them raise important questions: How are the languages defined? And how are programs in language X converted into the primitive instructions for the machine they need to run on?

Programming languages are specified using grammars. Grammars vary according to the degree of complexity of the machines needed to parse the language. Malcolm Newey and others have taught units on the types of grammar and the automata needed to recognize them.

John Lloyd tells me that COMP1600, which used to be the second year unit COMP2600, which was taught for a long time, covers a bunch of theoretical topics at the introductory level: first order logic, regular languages, finite automata, structural induction, Hoare logic.

In 1974, one of the fourth year honours units (taught by Peter Creasy) was on compiler construction. We all implemented an interpreter for an Algol-like language, of which the critical first stages

[12] I.e. just before the founding of DCS.

[13] `https://plvmq.bitbucket.io/#SAPLING/`

were a tokeniser and a recursive descent parser. Chris Johnson later taught compilers, based on his own experience working on a significant compiler project.

Computer science students sometimes doubted the value of learning how to write parsers, interpreters and compilers – When would they use this knowledge in their future careers? In fact, understanding of finite state machines and parsing turned out to be very useful in many real-world (non computer science) applications, even though in many applications command languages are replaced by user interfaces with menus and check boxes.

Chris Johnson's thoughts on the decline in importance of compiler construction

The Unix tools `lex`, `yacc` etc are no longer used much, and as far as I know, not taught. Compiler construction itself – parsing and analysis and generation – dropped out of the curriculum; in the 70s and 80s the compiler construction unit was often a form of capstone, bringing together some of these formal language processing tools and an exercise in team work, to develop a larger scale, end-to-end program that students had not encountered previously. BUT, I reckon the notion of states and transitions would still appear useful to the formal specification designer.

Understanding how the final code optimisation stages of a compiler work is also interesting and potentially valuable – how can a program be transformed into another program which is semantically equivalent but runs faster on given hardware. "In rare cases this knowledge can be applied to determining bugs in the compiler's code generation itself."[14]

Juris Reinfelds: Another view on when to optimise code

"At this stage Richard [Miller] quietly put aside the optimising pass of the C compiler for PDP-11 code, to be implemented later when time permits. To this day nobody has complained about its absence which shows that good programs do not need automatic optimization while bad programs cannot be rescued by it."

From Juris Reinfelds, *The first port of Unix*, `https://documents.uow.edu.au/content/groups/public/@web/@inf/@scsse/documents/doc/uow103747.pdf`.

One of the features which increases the difficulty of parsing a programming language is overloading. Ignoring the fact that '+' may mean either integer-addition or floating-point-addition even in FORTRAN, operator overloading was introduced in Algol 68, and has appeared in languages such as C++ and Ada, but not Java. In trivial cases such as the FORTRAN one, it is clearly very convenient. However I find it deplorable that in C++ one can define two different functions with the same name, distinguished only by the number of their arguments. Unfortunately, no-one listens to me!

Chris Johnson's view is that:

> In the hands of a skilled designer, operator overloading can provide very expressive and highly readable libraries. But the difficulties of providing meaningful error diagnostics, increased problems in debugging code, and the opportunity for confusion, have limited the use of overloading in good programming practice.

5.4.7 Multiprogramming

Solutions to certain problems are naturally expressed as parallel or pseudo-parallel algorithms rather than in sequential form. Discrete simulations are an example. Modern operating systems, even for single CPUs, must be pseudo parallel to support timesharing and to respond to both hardware and software interrupts. Section 2.13 on Page 58 has already described the launch of Peter Creasy's unit on Operating Systems, which covered issues of resource allocation, deadlock avoidance, interrupt handling, critical regions, and process scheduling. Minix and, I think, XINU were used to teach operating system principles.[15]

Later, an operating systems lab comprising four PDP-11s was set up in the CS&IT building to support this type of learning. The PDP-11 is Malcolm Newey's favourite computer and he kept two

[14] Chris Johnson's personal experience on the York Ada compiler in 1981.

[15] XINU is a self-referential acronym: XINU Is Not Unix. Such SRAs were popular a few decades ago.

of them in his office after the lab was dismantled.

To operate in the parallel world, a programmer needs to be wary of many possible pitfalls. Subtle timing bugs can lead to programs failing at random times. Such bugs are very difficult to track down. I suspect that many of the unexplained crashes of EXEC-8 on the Univac 1108 may have been timing bugs, and "hangs" may have been related to resource allocation deadlocks.

DCS units have covered concurrency in Modula-2 and in Occam. The panel below records Peter Bailey's experience with Occam and transputers.

Units in High Performance Computing (See Page 159) and a short-lived unit in multi-core computing also fall under the heading of multiprogramming.

Peter Bailey chats about programming in Occam

Dave: And you said that you had done Occam, too?

Peter: Yes, so in third year we did a whole range of units. Third year was the first time I didn't have to do any subjects other than computing, and I was extremely happy about that. Malcolm [Newey] taught a parallel programming unit. And we had this transputer, one transputer chip, which was kind of a data-parallel custom-silicon chip that had been created, probably in about 1988. We had one of those, and we had an Occam compiler. Occam was the programming language for it. We had one compiler and it took an hour to compile your program. There were no syntax-checking editors available and no modern-style development environment to warn you that things were broken. From memory it was not a typed programming language either, so if you made a single syntax error, you could wait an hour to discover that you'd missed a semicolon, or that you had made a mistake in indentation. It was a language which relied on meaningful indentation for syntax, which wasn't something that any of our other programming languages made use of. Now, of course it's become extremely popular in languages like Python and F#.

Dave: Transputers were a really interesting design, and I think Occam had a message passing primitive. Each transputer had four ethernet or comms ports and so they could be connected up into two-dimensional arrays, and your computation worked on a matrix of transputers.

Peter: Yes, it was a very interesting design.

A Transputer system from RSISE, now in storage. *Photo: David Hawking*

5.4.8 Software Bugs, Program Verification

Bugs in software are very common, and may live undetected in heavily used code, waiting for the right combination of circumstances to wreak havoc. I'm a little bit proud of my ability to debug my programs, but I fear that the reason for my prowess is extensive practice, caused by my "ability" to put bugs in my code. I'm definitely not the person to write code for surgical robots, fly-by-wire aircraft, or autonomous vehicles. It's sobering to think of all the published research results (across all disciplines) which have been obtained from runs of computer software containing bugs.

Even software assumed to be more-or-less bug-free may require maintenance over the years. Programs may need to be extended to accommodate changes in data format, changes in the hardware on which they are run, or increases in scope. Billions of dollars were spent around the globe in 1999 checking for and mitigating problems due to the "Y2K bug." Bizarrely, because of remarkable mitigation success, people afterward complained that Y2K had been a hoax!

Ian Simpson, formerly of ANU Computer Centre, reportedly pursued a lucrative career in the UK, Y2K-proofing and subsequently maintaining ancient COBOL programs many decades after they were written.

Adding functionality to very large programs can be horrendously difficult when the code is too large to be understood by a single person. Because of the complexity, new additions may introduce unforeseen undesirable effects. Modularisation is a partially successful defence against these and service oriented architectures go a step further.

One of Edsger Dijkstra's key contributions was highlighting the inadequacy of testing as the sole means of ensuring correct behaviour of software [or hardware]. "Testing can prove the presence of bugs, but never their absence." Dijkstra argued the need for proofs of program correctness. Unfortunately, proving correctness has proven to be difficult in general, and effectively impossible for programs where correct behaviour isn't able to be formally specified, as is the case with simulation and modelling.

Programming assignments in most units provide students with hands-on debugging practice but, over the decades, DCS has also taught a number of units related to verification and program semantics. Taught in the 1990s, *COMP3066 Formal semantics and Programming* included Hoare logic, predicate transformers, denotational semantics, natural semantics, and the construction of rigorous arguments about programs. Jeremy Dawson tells me that COMP2600 also touched on Hoare logic and weakest preconditions. Another unit, *COMP8130 Verification and Validation Approaches* was taught in fourth year.

In the early 2000s ANU and the University of New South Wales were involved in a joint research project in trusted computing. In 2010, Michael Norrish (now working for CSIRO Data61 but with an office in DCS) was a co-author of a very influential publication[16] which claimed. "To our knowledge, this is the first formal proof of functional correctness of a complete, general-purpose operating-system kernel." Prior to his time at ANU, Tony Hosking was also involved in formal verification of a multi-threaded garbage collector.

5.4.9 Computer Architecture and Network Programming

DCS teaching of computer architecture started with Ray Jarvis's introductory-year teaching of low level computer logic, including how to build arithmetic units from NAND gates. Later, John Hurst taught a computer architecture unit in fourth year honours from 1974 on. For his fourth year project, in the mid 1970s, Stephen Edwards built a hardware stack. Taking the opposite tack for an assignment in John's course, I designed a machine based on a queue rather than a stack.

Over the decades, a number of computer architecture units have been taught. Teachers have included Brendan McKay, Ramesh Sankaranarayana, Peter Strazdins, and Chris Johnson, while back in the 1980s, Richard Brent taught VLSI design.

Machine architectures remain a major concern for high performance computing, operating systems, and system software. However John Hurst notes a decline in importance of Computer Architecture in both teaching and academic research:

> The 1970s saw a bit of action in the Computer Architecture area, with several research groups, led by people such as Chris Wallace, Murray Allen, and Arthur Sale. Arthur Sale's work crossed the boundary between architecture and language implementation, and my work at ANU was similar. The Burroughs 1800 project was about migrating the hardware/software interface towards the programming level, culminating in outcomes such as the JustaSystem (Unix-like, written in Modula), AMPL (an Algol-like Micro Processor Language), the Software Oscilloscope, and PODL (Processor Object Description Language), all done during my time at ANU.

[16] *seL4: formal verification of an OS kernel.* `https://dl.acm.org/doi/10.1145/1629575.1629596`

> But that is on a personal level. What is of wider significance is the shift away from hardware research and computer engineering due to the advent of Very Large Scale Integration (Richard Brent taught units on this) and the consequent microprocessor revolution. The CISC/RISC debate was pretty well resolved by the time I went on study leave in 1986, and universities without access to microprocessor fabrication facilities were just left behind. There were a few pockets, such as Chris Wallace's work on capability architectures (one of the reasons I moved to Monash!), but even these lost momentum, and computer engineering research evolved into computer network engineering. We have moved on a long way in 30 years!

In our highly-interconnected, communication-dominated world, network programming becomes highly important, and this is reflected in DCS teaching. Online services relied on by billions, such as Facebook, Alibaba, Amazon, LinkedIn and major search engines build their services on top of very large numbers (millions) of distributed servers. Thousands of servers may be used just to orchestrate the workflow across the hundreds of services potentially required to deliver a results page for a single search query. Each individual service is likely delivered by large numbers of interconnected servers.

DCS units have covered protocol stacks (e.g. the standard TCP/IP model and the OSI model), the World-Wide Web, wireless networks, service oriented computing, and the Internet of Things. Teachers of these units have included Vicki Peterson, Bob Edwards, Chris Johnson, and Weifa Liang.

5.4.10 Data

Understanding data is critically important in most computing applications. For graduates entering the business or government world, databases are critical. DCS units on databases cover the critical ACID (atomicity, consistency, isolation and durability) properties of transactions, journalling, locking granularity, normalisation, and the efficiency of database computations. Of course, different types of databases have been covered, along with query languages like SQL.

Database units in DCS have been taught by Mike Papazoglou, Qing Li, Jeffrey Yu, Vicki Peterson and Paul Thistlewaite.

In recent times Data Science has become a profession in itself and a whole industry has grown up around data mining, data matching, machine learning, and privacy preserving operations on databases. These topics have been taught by Kerry Taylor, Graham Williams, and Peter Christen. .

Kerry Taylor, 2021 convenor of postgraduate and undergraduate programs in applied data analytics. ***Photo: CECS, ANU***

5.5 Other Important Topics

Here is an alphabetical list of most of the other important topics covered at some stage during the history of DCS.

Artificial Intelligence. In the mid 1970s it was said that Artificial Intelligence (AI) was doomed to eternal failure, because as soon as an AI problem was solved, it was immediately re-classified as an engineering success. Robin Stanton taught AI as an honours subject from as early as 1974. Subsequently, the subject has been taught by Graham Williams and Kerry Taylor.

Code Art. Tim Brook, who describes himself both as a lapsed mathematician and as two-thirds of a Goodie,[17] ran a subject in conjunction with Henry Gardner, called *Art and Interaction in New Media: Synthesising Visual Culture and Computer Science*. The unit was offered to students in the Arts School and in DCS. This unit has gone from strength to strength according to Henry and, in 2020, was running in ACT Colleges as an ANU "H" Unit taught by Ben Swift.[18]

Tim believes that art is about helping you see things in a way you hadn't seen them before, i.e. expanding your ways of seeing things – that's what art can contribute to the computational sciences, and to the visualisation of data.

Believing that creativity flourishes best in an environment with some constraints in place, Tim asked students to produce a major project in Java (applets) under the heading of "Secrets". He recounted two very successful projects. A mature engineer studying art for the first time, used a keyboard and screen with the message, "Type your secrets here. They will be safe with me." When a "secret" was typed it appeared on screen, appearing in a wavy line in three dimensions, wriggling like a snake and intertwining with each of the other secrets, each in their own distinct colour. The result was a changing multi-coloured pattern, in which words were visible but from which secrets could not easily be decoded.

An Art School student who had never coded before, managed to make very creative use of eye-tracking and facial recognition libraries. The finished art work consisted of a computer screen with no keyboard or mouse, displaying a 19th century family portrait. The screen was posed in an elegant Victorian frame. An eye-tracker was able to detect which person in the portrait was being looked at by the viewer, and the combined camera and face-recognition software were able to replace that person's face with the face of the viewer. Definitely a new way of thinking about the world! Not a bad first project!

Tim insists that code art requires tighter discipline than normal programming because to create art you are constantly modifying the code. To be able to do that successfully it needs to be transparent, well-structured, and well-documented. This is particularly true in *live coding*[19] as practised by Ben Swift, who now runs the unit started by Tim and Henry, for music generation.

From 2020 Ben Swift and Charles Martin, together with the ANU School of Music, have been running a "laptop ensemble" unit where students compose and improvise computer music, and graphics, in live performances using laptops.

Computational Science. According to Henry Gardner, "Computational Science can be thought of as being sexed-up HPC and stunning applications." Following the Computational Science and Engineering Education project, which Henry led from 1996, and the APAC education program, in the late 1990s and early 2000s, many computer science, engineering and mathematics students experienced units with a computational science flavour. The departments of Computer Science and Mathematics jointly sponsored a Bachelor of Computational Science from about 1999. The graduate program in "eScience", run jointly between ANU and RMIT University from 2001 to about 2007, combined computational science with software engineering, computer graphics, visualisation and human computer interaction.

Document Computing. Paul Thistlewaite introduced document technologies as a Fourth Year Honours unit in the late 1990s. The unit covered natural language parsing, metadata, document

[17] https://en.wikipedia.org/wiki/The_Goodies

[18] According to the ACT Board of Senior Secondary Studies, "H classification is given to a year 11 and 12 course which is designed and accredited by an Australian higher education provider and the ACT BSSS. Successful completion of the course can be recognised towards an undergraduate degree with that provider. Students can study a maximum of two H courses."

[19] https://en.wikipedia.org/wiki/Live_coding

summarisation, text retrieval and hypertext. Nick Craswell, Peter Bailey and I took over the unit temporarily in 1999, after Paul's untimely death. More recently, Hanna Suominen has taught a similar unit.

Graphics and Visualisation DCS has at various times taught units in computer graphics. In the early days it was necessary to talk about low level operations such as dual frame buffers, vector generation, and bit-BLTing. Now there is tremendous support from libraries and packages and most programmers don't need to worry about low levels. I suspect that this is true even for programmers developing games, where performance is critical. As mentioned below, computer graphics and visualisation was a major focus of the eScience graduate program and much of the project work that its students undertook was focussed on the Wedge virtual reality theatre.

High Performance Computing. Large scale scientific computing on advanced hardware requires much more than simple numerical methods. To achieve high performance on distributed memory machines which have a high ratio of compute power to communication latency and bandwidth, one must distribute the data carefully, and structure the algorithm to minimise communication. Similarly careful choices must be made on multi-core and/or multi-cpu shared memory machines, and vector processors. Peter Strazdins, Chris Johnson, and Alistair Rendell have taught a number of units on HPC.

Image Processing and Computer Vision: In the 1970s and early 1980s, Ray Jarvis taught this in Computer Science C03, and supervised my fourth year honours project in this area. During Ray's absence on sabbatical, C03 was taught by three CSIRO Division of Computing Research scientists: Don Langridge, Don Fraser, and Malcolm Stewart. Later, a similar unit was taught by Phil Robertson, then of CSIRO Division of Information Technology, later of Canon Research Australia and then COO of NICTA. Still later, computer vision was taught out of the Research School of Engineering until 2021 when it, once again, reverted to DCS.

Numerical Methods. Meaningful scientific computing relies on appropriate application of the principles of numerical methods to achieve results in acceptable time and to achieve desired levels of accuracy in the results. DCS units have covered these topics in some years; in others the teaching has been left to the Maths department. Brian Molinari taught the original unit in this area, Computer Science C02.

> **Brendan McKay:** Richard Brent was teaching numerical methods when I came to ANU [in 1983], then I shared it with him for a while and then I took it over (it was called C12 if I recall). It became the most popular third year unit just when the department was looking for something to chop to make room for software engineering. So we gave the whole subject to the mathematics department.

Robotics Robotics has largely been considered engineering rather than computer science, but there have been a number of exceptions. As described on Page 118, Hanna Kurniawati is currently leading a robotics project. She occasionally uses the robot in undergraduate teaching. Robots have also been used as an outreach activity and to excite the passion of secondary students for studying CS at ANU.

In the late 1980s, Paul Mackerras modified a toy electric car to fit various sensors and an onboard computer. The resulting minature robotic vehicle was called "the Moth". Peter Farmer and David Walsh subsequently worked on it. A considerable number of examples were made and distributed to ACT secondary colleges, culminating in a competition at ANU. Vicki Peterson ran this outreach program.

5.6 Software Engineering

A computing graduate whose employment requires them to work on large scale software projects, requires an understanding of the processes and tools for project management, to a much greater level if they are responsible for managing and delivering the project. Skills in managing time, staff, sub-contractors, resource procurement, documentation, downstream maintenance, and cost, and in ensuring that what is delivered meets the requirements of the client may be critical but are not part of the academic discipline of computer science. They are, however, common to other types of engineering.

Project management skills are unlikely to be sufficient by themselves. A lack of knowledge of many of the computing topics listed above would make it very difficult for the manager of a software project to sensibly break the project into sub-projects, to judge whethers estimates and progress reports from sub-projects are realistic, to ask probing questions, and to decide whether proposed hardware purchases are necessary, or adequate.

The teaching of material relating to the engineering of large scale software systems, other than that already covered in the computer science curriculum began in the early 1990s. Paul Thistlewaite, Brian Molinari and Chris Johnson were involved. Eventually, the importance of Software Engineering was elevated, and a four year Bachelor of Software Engineering degree was introduced. It achieved formal accreditation as an Engineering program.

Clive Boughton's involvement in the software engineering program began in 1995. He was later joined by Shayne Flint and Lynette Johns-Boast. All three have had a professional focus on methods, tools and processes appropriate for large scale software engineering projects, and have strong views. Shayne Flint outlines his motivations and experiences of teaching software engineering at ANU on Page 111. Here I present views on Software Engineering from Clive Boughton, Lynette Johns-Boast, Shayne Flint and Chris Johnson.

Clive Boughton

> There's a lot more to software *engineering* than programming.
>
> However, it is important to understand the underlying theory of programming just as it is important for structural engineers to understand the mathematics and physics surrounding stresses and strains, forces and moments. Unless a structural job is somewhat different to the norm and has to be 'proven' (backed up with all the appropriate calculations) then all that is necessary is to apply standardised codes (typically in the form of parameterised, pre-calculated tables). This reduces the effort of doing the engineering calculations quite significantly.
>
> The biggest issue in most engineering industries (software engineering included) is getting the requirements and planning right. Doing so takes time effort and (often) much negotiation, analysis and the forming of agreements. Part of the planning is sorting out of the appropriate architecture (often being part of the requirements) and essential design.
>
> Another part of planning is determination of the (various) processes to be used to construct what is required based on the chosen architecture and design.
>
> A further part of planning is to determine the necessary verification and validation work at the various stages of development.
>
> Of course breaking the project down into tasks with time and effort estimates is essential. Risk and risk mitigation are not trivial aspects of a structural project. They aren't for software engineers either. Project context (or where the final product is to be built and sited). Get that wrong and a project can cost a motza – way above estimates.
>
> How does all the above relate to a software engineering project? In my mind, very strongly. In terms of effort, software engineering projects can easily be as big as very significant structural projects. In fact, sometimes thousands of person years!
>
> Another problem that frequently arises on such projects is the fact that (typically) the people in charge of project management don't appreciate or understand that (like other different members of a large team) software practitioners also need guidance and leadership. There's an underlying attitude that software developers are real smart and it's not a good idea to mess with them – just let them work it all out. All of which is arrant BS. As an example, whilst I would not tell an experienced competent bricklayer how to lay bricks, I would also not tell a competent Java programmer how to write Java code – un-

less poor work was obvious. However, I would tell a bricklayer that he's laying his bricks in the wrong place or not building his brick walls to specification, if that were the case. Similarly, I would tell the Java programmer that he must not write programs that already exist and work or that his written programs don't do what is required, if that were indeed the case. Regular (and helpful) peer review can work well in programming teams, and is necessary to reduce potential friction and embarrassment.

An industry that's classified as having the worst performance record of any is nonetheless held up as the industry with the smartest people. However, it's not just the programmers that help cause the failures. The tools, and methods leave a lot to be desired. The industry is slow to take up methods and associated tools that boost productivity and quality – much slower than in other industries. This stems from the universities – mainly because of the narrowness of topics in CS.

Over the last 50 years (my time with using software techniques) I haven't seen much really effective change to enable greater reliability, stability, generation, testing and deployment of products. There are many products that claim some of these characteristics, but under the hood the old approaches of program, compile and (maybe) test, pervade. The language wars still go on, when in my mind they were essentially solved many years ago. A new language is typically not new technology – it's just the same old stuff in a different wrapper. When choosing a language for a project there's typically little analysis of what fits best. So choices are reduced to favourites – even when they don't fit.

As I used to say in my classes when I was teaching at ANU: If you want to be a programmer only focused on a particular language and its supporting tools, then join the fraternity of bricklayers who can (no doubt) lay bricks very well and professionally, but have little motivation to do something different. If you want to be an engineer, then you must learn lots of different, seemingly disparate, things including planning, analysis, architecture, design, construction, testing, processes, properties of materials etc. You don't have to be learned about ALL these aspects of engineering, but you do have to understand their relevance and purpose and respect the fact that each aspect is possibly a career in itself.

Oh! Yes! There's measurement! As a scientist, I learned that measurement is paramount to making progress in practically any domain. In general there's a symbiosis between theory and experiment. Of course, not all theories are proven (through experiment or application) before they are accepted as fact. However, with CS, experimentation seems almost anathema to many. I find this difficult to accept. I think that it's because of this lack of attention to 'measuring things' (or even trying to measure things) that software development is perhaps still more of an art than a professional practice.

Lynette Johns-Boast

Dave: Clive tells me that programming is like brick laying, an important skill but only a tiny part of an overall project. Do you agree?

Lynette: Definitely, 100%. I am frequently rude and say that so long as you have a good design (which presupposes a good and deep understanding of the problem) trained monkeys can do the coding. ☺ Unfortunately education tends to focus on the basic programming skills without helping students to understand that this is probably the least important part of what they are learning and will be doing in future. I could go on for ages about this. It's all about doing the right thing right rather than doing the wrong thing right; that testing/verification (especially 'against' requirements) only proves that you have met the requirements, but not that the requirements were what was required in the first place.

Dave: Software Engineering at ANU had a significant rise in focus from the 1990s but that focus seems to be declining. Is that true? What do you think about it?

Lynette: Yes, I think that is true of ANU and I think you could quite probably relate it directly to the staff in the Department/School/College. You may not know, but it has been canned – the School will no longer offer a BSEng. COVID accelerated its demise, but the moment its last defenders – Shayne and myself – left, it was doomed. The School is peopled by 'scientists' who

> seem not to understand that software is only a tool, not a be all and end all in its own right, and so don't understand and value the application of the craft of software development. The influence of scientists is not something unique to engineering/technology at ANU, but is a recognised worldwide phenomenon and is something that is causing increasing concern in the world of engineering education.

Lynette and I subsequently had a good chat over coffee, exploring what we each understood by "programming". She informed me of a trend in the Australian Public Service and elsewhere, where what I considered to be the single role of a programmer was separated into multiple roles such as Business Analyst, Solution Architect and Programmer. I guess that given a highly specified design, including data structures and algorithms, even a relatively unskilled person could turn it into a program in the given language.

Siew-Gim McGregor, an ANU alumna and former ANU MISD employee, who has worked for many years as a consultant to many Australian Government agencies, confirms this specialisation of roles in major projects: Enterprise Architect, Solution Architect, Business Analyst, Service Designer, Process Designer, Testing Team, Infrastructure Group, Design/Developers, ... She says that initial project meetings may involve 30 or 40 people.

Shayne Flint

> Software Engineering, like all forms of engineering, is about applying scientific and other knowledge to identify and solve problems perceived by people in the real-world.
>
> In my view, most software engineering programs fail to properly prepare graduates for engineering roles. Most programs comprise a computer science core with a few 'management' and 'project' units tacked on along with content aimed at satisfying accreditation bodies.
>
> This is not enough. It all but ignores the fact that real-world software is usually developed and operated within dynamically complex socio-technical environments, and that knowledge from a range of disciplines needs to be applied to build and operate effective software within such environments.
>
> To work effectively in complex environments, software engineering graduates need much more than deep computer science knowledge. They must also understand systems theory, systems engineering and how to deal with complexity, uncertainty, change, human behaviour and risk. They must also be capable of, and feel comfortable working in broad multi-disciplinary teams. They need to develop respect for, and embrace the contributions of other disciplines including artists, philosophers, historians, entrepreneurs, managers, sociologists, economists, ecologists and even lawyers.
>
> These capabilities and 'engineering culture' need to be developed at university and beyond. It is for this reason that units such as Discovering Engineering, Systems Engineering for Software Engineers, Unravelling Complexity and TechLauncher were introduced as core elements of the BSEng, and why deep and extensive connections with industry were developed over many years through guest lectures, joint research, mentoring, substantial real-world project work and regular public demonstrations of student work and capabilities.
>
> Unfortunately, few within the school saw any value in this view of software engineering. Despite this, the BSEng was kept alive for more than two decades by dedicated academics, very enthusiastic industry collaborators and support staff as well as past and current students.

Chris Johnson

> What we lack in software engineering is a body of theory from computer science and from engineering that is formal, tractable, and has useful applications to programs and systems. Other engineerings have theory from related sciences (physics, and chemistry particularly) and a paradigm of sufficiently detailed modelling and analysis that applies those sciences through the medium of applied maths. In computing we have a few snippets, but as a prime example the best developments of program proof have not kept up anywhere near the needs of everyday programming, and are not part of the software engineer's toolkit. Only the use of practices including extensive, planned, repeated testing have enabled us to get as far as we have. Trying to educate programmers who are intelligently constructive, and who can then be educated towards developing a critical, sceptical, testing point of view: design patterns and refactoring are fragments towards theory of

programming but we have no well-rounded paradigm (I hope that this aspect of software engineering education was my own approach but I didn't get it to work well enough in education, it was poorly regarded by gung ho programmers and only developed in a few people).

I think that the general level of industrial programming is weak, expensive, and is frequently misapplied to poorly specified systems: our abilities to bash systems together, effectively with big sledgehammers, is like 19th century engineering – not in the design of bridges, but in the production of iron and steel, large scale, world leading, somewhat random outcomes, steam boilers exploding, bridges collapsing, aircraft falling apart... the real software engineering need is still much more in the human and governance side. Our general inability to maintain and upgrade, with continuing horror stories of errors and problems in real large systems continue (Boeing 737 MAX, the recent Australian Stock Exchange system crash, our inability to design and implement an electronic health record system – and the inability to know why this is hard and how we should (or why we should not) attempt it, the awful weight of legacy code in economy-critical systems, the unquantified risks in an insecure, hostile world system... Developing new systems (greenfield programming – and greenfield system application development) is really the easy bit! and we do not teach enough about hubris.

I believe that programming is a special skill.

Programming is the stuff of the building components that the project has to put together – architect/designer vs bricklaying is a poor analogy. Turing machines and computability do not constitute a body of theory that can be applied in software development or analysis, unlike the mechanical and electrical and electronic theories of engineering that have a mathematical formulation that can be practically turned into the basis of solution designs and to validate designs, usually through abstracted modelling and calculus. Despite that, we have a need for a better base than "programming" alone – without any understanding of formal bits in the education the system designer can lack a vision of what a rigorously constructed system design looks like. Formal thinking in computing can be used as the basis of design languages, that can in turn be transformed and used as the basis for constructing software systems.

5.7 BInfTech Degrees

Roger Clarke spent 11 years as Reader in Information Systems in the Department of Commerce, Faculty of Economics from 1984 to 1995 and, together with Robin Stanton, was instrumental in the establishment of a set of degrees catering to the information needs of various future professionals in Accounting, Commerce and Economics. The four degrees were:

BInfTech (SE) – This eventually morphed into the Bachelor of Software Engineering (BSE).

BInfTech (IS) – Designed for professionals with an understanding of business, and a technical bent, but no desire for an academic career. The short-term career-target was as a systems analyst / project manager for organisational systems and packages, up to and including ERP (Enterprise Resource Planning systems, of which SAP was a well-known packaged example).

BInfTech (Comm) – With a greater business focus than the IS variant, with a focus on financial management information systems (FMIS), and perhaps destined for a CIO role.

BInfTech (Ec) – Designed to give future employees of the Merchant Banking or Finance industries an appropriate IT background.

The BInfTech degrees were controlled by DCS and the Information Systems group within the Department of Commerce. Because those departments were in different faculties, the degree rules needed to be approved by both the Economics Faculty and the Faculty of Science. The Information Systems Group was led by Ron Weber prior to Roger Clarke's arrival and by Shirley Gregor after his departure.

Roger's masters thesis was entitled: *The Implementation of Functional System Design and Development Techniques in a COBOL Environment* and his background prior to coming to ANU was as a practitioner in the information technology industry. One of his roles was as a design controller in a 150 person-year project to develop an overnight transaction processing system for the London Stock

Exchange. Post-ANU he has operated a consultancy company Xamax, named after the Neuchâtel Xamax football club from Neuchâtel, Switzerland. Coincidentally I have an honorary doctorate from the Université de Neuchâtel, and on a visit there in 2002, photographed the club's headquarters, never guessing the connection to Roger's company.

Roger continues to lecture in DCS units. He points out that the Information Systems material in his lectures needed to change just about every year since he started teaching in 1984.

5.7.1 Why the Teaching of Information Systems Evolved

In the beginning Roger was given the mission to teach accounting students the minimum they needed to know about computers. Text books for the unit were already set and he found them only marginally useful. In that era, a lot of accounting was still done on paper, and students needed to be made comfortable using computers, and had to be taught about command line interfaces, phone lines, modems and acoustic couplers. Over the years, the sophistication of computer interfaces increased, as did student familiarity with them. The advent of Macintosh laboratories, equipped with a really good accounting package, was a huge step forward.

Over the years, the role of computers in accounting and business also evolved rapidly. It was realised that organisations run on information and that there were great business advantages to organising and using information well. Initially computers improved the way current financial information was reported, but soon organisations realised that there was value in aggregating data and in exception reporting. This was the Management Information Systems era. Eventually, the value of computers extended to board rooms, where spreadsheets and discrete event simulation were used to model the effect of various possible actions. This was the era of Decision Support Systems. The case for the Port Botany container terminal was made by simulating the length of queues of ships outside Sydney Harbour, based on the known schedule of ship arrivals.

Concurrently with all this, computerised information became used to manage people, equipment and software. Computers controlling production lines in factories became integrated with the company's information systems. Information systems in one company began interacting with information systems in other companies (B2B) and individuals (B2C). Information security became a vital concern, and maintaining privacy of personal information while exploiting the value of "big data" became critical. Roger has built a strong reputation publishing and presenting on *dataveillance* and became very well known for his lectures on "Big Brother Google".

Over the decades orthodoxies in organisational IT models have changed and changed again. Monolithic information systems popular in the 1970s were debunked, and replaced with individual systems for payroll/personnel, accounts receivable, general ledger and so on. Now monolithic ERP (Enterprise Resource Planning) systems such as PeopleSoft and SAP are again in vogue. Small businesses who used to run accounting packages on their PC, now have little alternative but to keep their business information in the cloud, with SalesForce or MYOB.

Similarly, orthodoxies of software engineering have also changed. The rigid models of software lifecycles from the early 1990s were replaced with RAD (Rapid Application Development) and Agile methodologies. [I've noticed that some organisations apply Agile methodology in a very rigid way!]

As Roger says, what an accountant, business person, CFO, CIO, CEO, COO, or board member needs to know about computers has changed frequently and dramatically over the decades. Roger is an author on two papers reviewing the development of Information Systems in Australia over the decades:

- The Information Infrastructures of 1985 and of 2018: The Sociotechnical Context of Computer Law & Security. `http://www.rogerclarke.com/II/IIC18.html`
- A Retrospective on the Information Systems Discipline in Australia. `http://www.rogerclarke.com/SOS/AISHist.html`

5.8 TechLauncher

In 2004, Shayne Flint was given the opportunity to coordinate the 3rd and 4th year group project units and was able to gain support to do two things: 1) have all group project students work on projects with industry, and 2) to have 4th year students lead teams of 3rd year students.

Shayne Flint: The story of TechLauncher

By 2014, we had seen an increase in the number of students interested in creating their own startups. At the same time, government and parts of ANU were becoming more interested in the innovation ecosystem. The CBR Innovation Network[a] was established and there were signs that rapid development of a vibrant ACT innovation sector was underway. This presented a unique opportunity to take another step towards my vision by 'branding' all of our group projects as 'TechLauncher' and allowing students to work towards developing their own startups as part of their degree.

The key innovation in TechLauncher is that we impose very few constraints – any client (government, industry, university and even students themselves), any project of sufficient complexity (messiness) including start-up ideas, any tools, techniques, approaches etc., teams comprising students from different degrees and levels.

Students are assessed on their ability to deliver client value – full stop. We look at how they work with their clients and users, make and implement decisions, and adapt as required to changing needs and constraints. Assessment takes the form of a series of "real-world" inspired project reviews involving multiple stakeholders (clients, tutors, mentors, peers, unit convenor and the team themselves).

As the end of each semester, we run a public showcase that usually attracts hundreds of vistors.

We made TechLauncher as real as possible and an opportunity for students to get a great start – "TechLauncher – Launch an Idea, Launch a Startup, Launch a Career!"

[a] https://cbrin.com.au/

5.9 Summer Schools

John Slaney tells the story of the long-running series of logic summer schools, which were mainly aimed at undergraduate students from around the country:

> They started in the summer of 1992-3. The initiative originally came from Norman Foo (I think that was before he moved from Sydney Uni to UNSW) and Rod Girle, and was supported by a wad of money direct from the Vice Chancellor – I don't remember the name of the funding mechanism. Over the next couple of years, the funding gradually dried up, so we were obliged to charge fees (well, a small fee for students, just to cover the morning and afternoon teas and a bit of the excursion, and a more substantial fee for those not in full-time education). Apart from that, we had to beg for scraps of money from CSL, and later RSCS, and from time to time other sources including NICTA. We were able to bring in a few external presenters most years, usually including one or two from overseas, though of course we didn't pay them anything more than expenses. We also offered a couple of scholarships each year to outstanding students.
>
> Once, I think in 2008-9, the logic summer school was combined with the machine learning one, but that didn't bring enough advantages to compensate for the disruption it caused (e.g. moving the date from December to early February) so it wasn't repeated.
>
> The logic summer school has always been aimed at later-year undergraduate students, but has also been useful for Masters and beginning PhD students. Part of the rationale was to keep potential research students in Australia by showing them that there is a research community in the discipline here (i.e. people who are just as nerdy as we are). Another rationale was to improve the quality of PhD research and shorten completion times, by giving research students a decent grounding in logic in an intensive way. A subsidiary aim was to bring potential HDR students to the ANU more cost-effectively than could be achieved through the summer scholarship scheme. It has always been impressive that academics from other universities are prepared to come here during the summer break and teach units without getting paid – I guess it's the joy of being able to talk about your pet subject to bright and interested students, and not having to assess them!

2005 Logic Summer School: John Slaney lecturing to an attentive audience. *Photo: Qing Wang*

This [2020] has been the first summer since 1992 without a logic summer school. I have long since lost count of the participants, but virtually every logician (in computer science, maths or philosophy) educated in Australasia in the last 28 years has been through the summer school, heard about completeness theorems, modal logic, Goedel's theorems and automated reasoning here, and completed the mandatory bushwalk at Tidbinbilla – it's a rite of passage!

5.9.1 Machine Learning Summer Schools

According to an international list of machine learning summer schools,[20] ANU/NICTA organised seven such summer schools between 2002 and 2010, including the earliest two on the list. They were the brainchild of Alex Smola. ANU organisers have included Doug Aberdeen, Alex Smola, Li Cheng, Marcus Hutter, Scott Sanner, and Wray Buntine. The 2008 School was held at ANU's Kioloa campus on the NSW South Coast.

[20] `http://mlss.cc/`

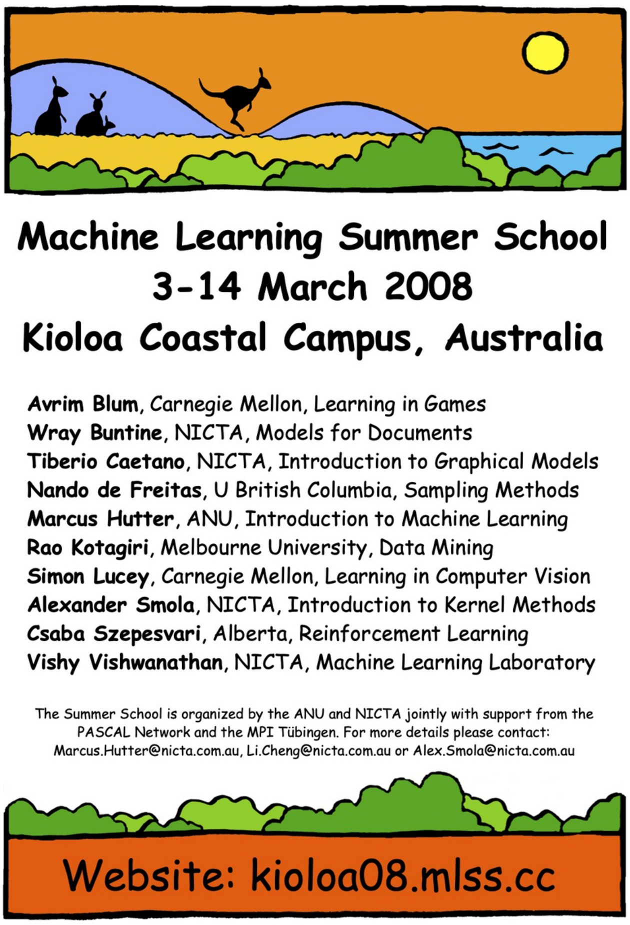

Machine Learning Summer School
3-14 March 2008
Kioloa Coastal Campus, Australia

Avrim Blum, Carnegie Mellon, Learning in Games
Wray Buntine, NICTA, Models for Documents
Tiberio Caetano, NICTA, Introduction to Graphical Models
Nando de Freitas, U British Columbia, Sampling Methods
Marcus Hutter, ANU, Introduction to Machine Learning
Rao Kotagiri, Melbourne University, Data Mining
Simon Lucey, Carnegie Mellon, Learning in Computer Vision
Alexander Smola, NICTA, Introduction to Kernel Methods
Csaba Szepesvari, Alberta, Reinforcement Learning
Vishy Vishwanathan, NICTA, Machine Learning Laboratory

The Summer School is organized by the ANU and NICTA jointly with support from the PASCAL Network and the MPI Tübingen. For more details please contact: Marcus.Hutter@nicta.com.au, Li.Cheng@nicta.com.au or Alex.Smola@nicta.com.au

Website: kioloa08.mlss.cc

Wray Buntine tells me that they were always well attended, for an Australian event, with a moderate collection of overseas students and some industry, but mostly ANZ research students. He says, "My favorite was definitely that held in Kioloa, where we got to go swimming and playing beach soccer, spent the night in beach bungalows and ate in a country canteen."

Marcus Hutter not only organised the Kioloa school, but drove the minibus, since no-one else dared to drive the 15-seater.

Scott Sanner agrees with Wray:

> I attended the MLSS's in 2008, 2009, and organised the one in 2010. I recall having about 120 registrants in 2010 with 50 international attendees. It was held on campus, nothing too notable.
>
> I think by far the best MLSS was the one held in Kioloa in 2008 organised by Marcus Hutter I think. Kioloa was packed and I believe we even had some people staying in nearby guesthouses. We had talks all day, a midday break for the beach, talks all evening, and large groups stayed up all night drinking, waking up in time for breakfast and coffee and another (hot) day of talks as we all crowded into the smaller lecture facilities before Kioloa built the more recent "shed". We had a large international contingent and a lot of overseas faculty; my favorite was Avrim Blum who had taught one of my undergrad units at CMU. Avrim was fascinated by the "T-Rex Death Stare" that the kangaroos would do when you got too close. Avrim, myself, my wife Sonia, and a bunch of others drove into Bateman's Bay one night to go to a pub... I remember it well since there were quite a group of characters there, one person in fact with a pet python who managed to get it to wrap around Avrim's shoulders. There have been lot's of MLSS's around the world, but I think for everyone who attended, the one we had at Kioloa was a uniquely Australian MLSS.

5.10 Unit evaluations

In the early 1970s, there was little or no effort to collect feedback from students about the content of units or the way in which they were taught. During the seventies, the Office for Research into Academic Methods (ORAM, later becoming CEDAM, Centre for Development of Academic Methods, now the Centre for Learning & Teaching (CLT)) was created and it became the norm that a student survey was conducted at least once per unit.

Jacinta Covington, a CS student who won the 1981 Hanna Neumann Prize for Third Year Mathematics, was very active in collecting student feedback in computer science units. She would turn up at a lecture armed with a clipboard, pencil and counter, allegedly counting the lecturer's "um"s and "ah"s. (For those seeking incentive to avoid the former, I have been informed and have had it confirmed, that in Turkish, "um" is a crude word for "vagina".)

The re-introduction of student fees (via the Higher Education Contribution Scheme, HECS) in 1989 led to a change in attitude in some quarters. If students are paying significant amounts of money to attend university, are they receiving an education worth the money? Responding to this attitude, there has been a trend to more professional course materials, recording of lectures, and institutionalisation of student feedback. One prominent member of ANU Council with a son studying at ANU is reported to have said, "There's no such thing as a bad student, only bad teachers".

To quote from the current SELT website[21]:

> The Student Experience of Learning and Teaching (SELT) is a survey run by the ANU Evaluations team and administered to the student population. It is one of the University's major tools for monitoring course development, academic development and student satisfaction.
>
> SELT surveys are run in each teaching session at ANU: Semesters 1 and 2 (standard sessions) and Summer, Autumn, Winter and Spring (non-standard sessions).
>
> There are two main survey components: the End-of-course Survey and End-of-course Teacher Survey.

In some years, the Students Association published a "Student Handbook," featuring reviews and ratings of units and lecturers.

While there is undoubted value in constructive feedback, I have spoken to some academics who have experienced considerable personal distress as a result of destructive feedback.

5.11 Computer Science Students Association

According to the ANU Students Association:

> The ANU Computer Science Students Association (CSSA) is a society for students studying (or interested in) computer science, software engineering, and related fields.
>
> We operate a common room on campus (N102 in the Computer Science & IT building, come say hi!); organise regular academic and social events, and events with industry sponsors; and a variety of other activities for the benefit of our members and the broader ANU community.
>
> Membership is open to all ANU students at only $5 for the whole year. Drop by the common room or our stall on market day to sign up

CSSA started in 1995, coinciding with the availability of a room in the new CS&IT Building, and it was relaunched in 2011.[22] Felix Friedlander is the current (2021) President.

5.12 Plagiarism

Plagiarism has been a constant issue over 50 years of CS teaching. As an undergraduate, assessment in several of my units was heavily biased toward examinations, because of the expectation that students would cheat on assignments. Some exams were worth 100% of the unit assessment. Many others required that you passed the exam as well as the overall unit.

Many methods for detecting copied computing assignments have been used in DCS over the decades, including manual inspection (which doesn't work reliably if multiple people are doing the

[21] `https://unistats.anu.edu.au/surveys/selt/`

[22] `https://anusa.com.au/life/club/cssa/`

marking). Automated tools have been used within DCS and across ANU to detect plagiarism. ANU currently has a licence for the commercial service `turnitin`.

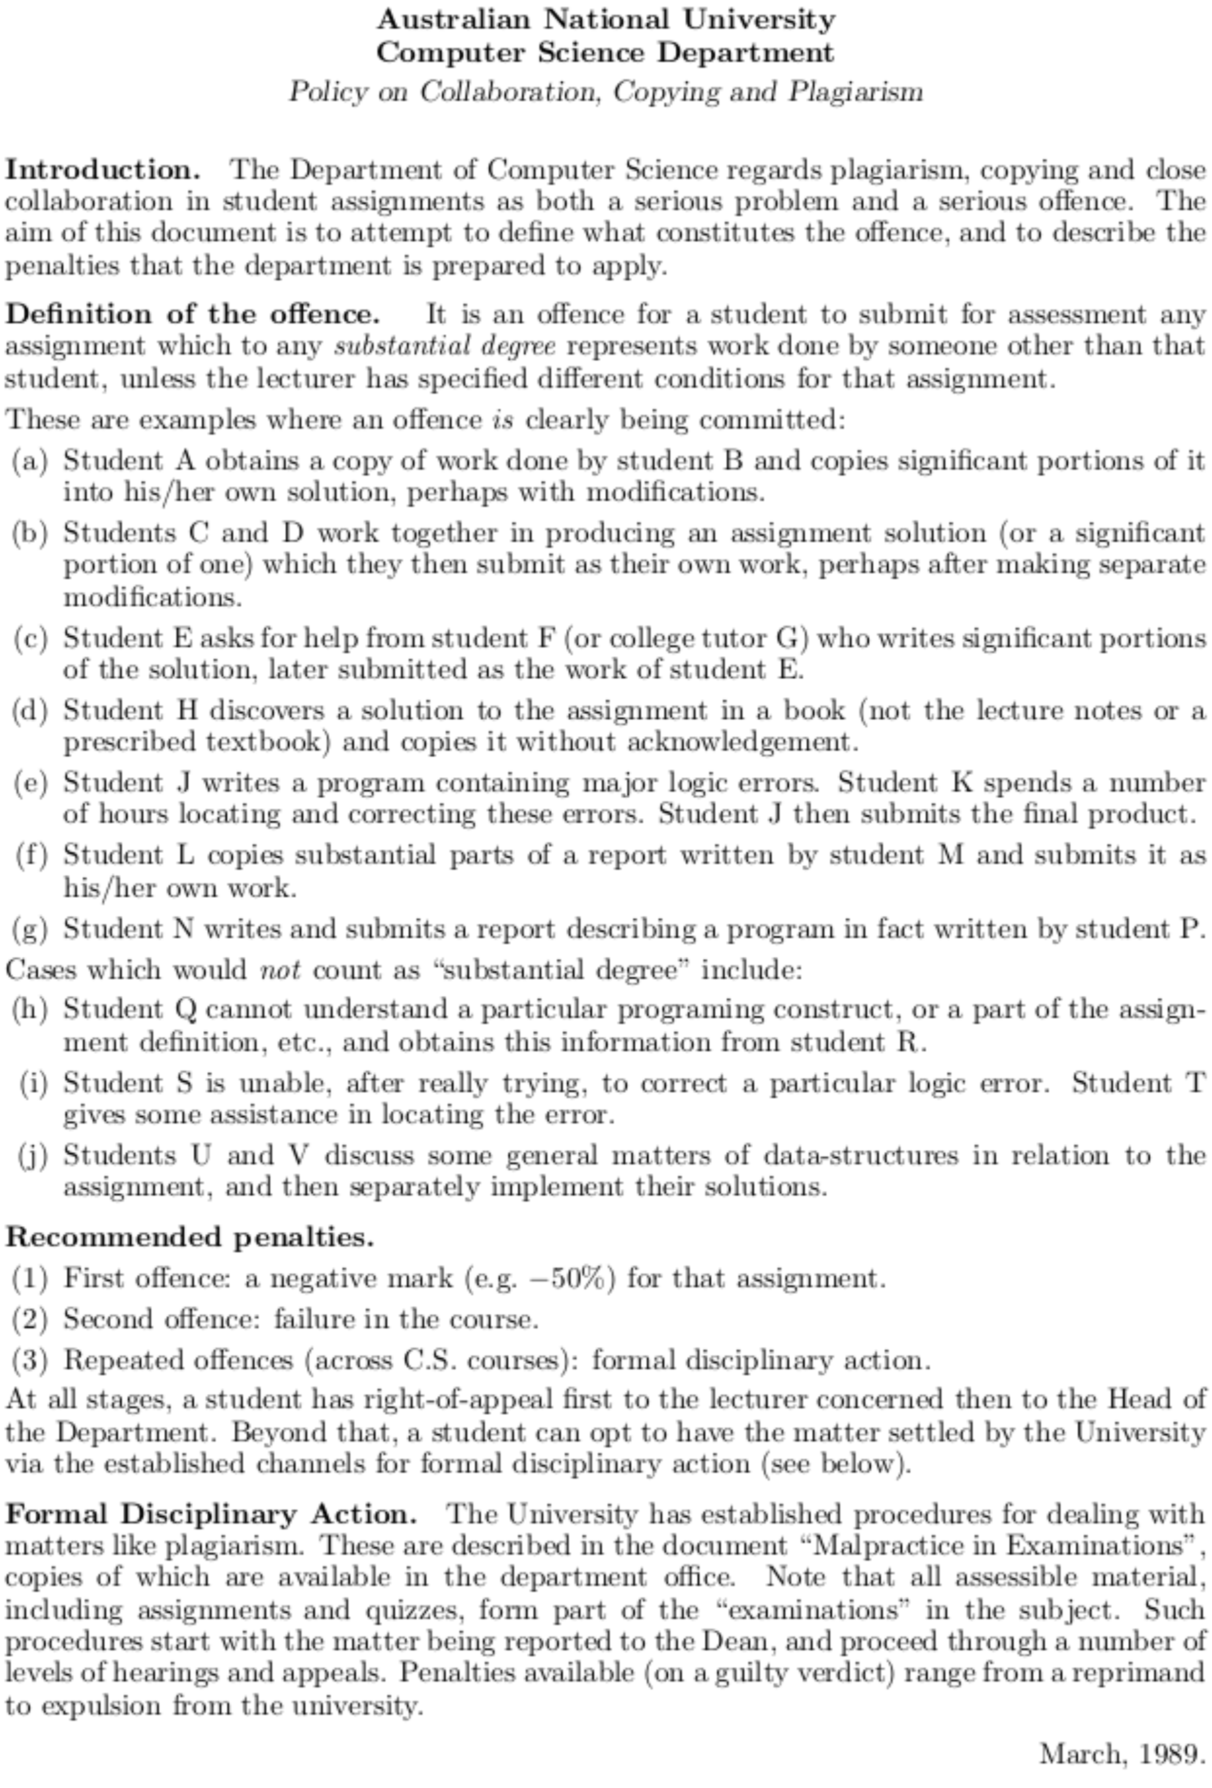

Australian National University
Computer Science Department
Policy on Collaboration, Copying and Plagiarism

Introduction. The Department of Computer Science regards plagiarism, copying and close collaboration in student assignments as both a serious problem and a serious offence. The aim of this document is to attempt to define what constitutes the offence, and to describe the penalties that the department is prepared to apply.

Definition of the offence. It is an offence for a student to submit for assessment any assignment which to any *substantial degree* represents work done by someone other than that student, unless the lecturer has specified different conditions for that assignment.

These are examples where an offence *is* clearly being committed:

(a) Student A obtains a copy of work done by student B and copies significant portions of it into his/her own solution, perhaps with modifications.

(b) Students C and D work together in producing an assignment solution (or a significant portion of one) which they then submit as their own work, perhaps after making separate modifications.

(c) Student E asks for help from student F (or college tutor G) who writes significant portions of the solution, later submitted as the work of student E.

(d) Student H discovers a solution to the assignment in a book (not the lecture notes or a prescribed textbook) and copies it without acknowledgement.

(e) Student J writes a program containing major logic errors. Student K spends a number of hours locating and correcting these errors. Student J then submits the final product.

(f) Student L copies substantial parts of a report written by student M and submits it as his/her own work.

(g) Student N writes and submits a report describing a program in fact written by student P.

Cases which would *not* count as "substantial degree" include:

(h) Student Q cannot understand a particular programing construct, or a part of the assignment definition, etc., and obtains this information from student R.

(i) Student S is unable, after really trying, to correct a particular logic error. Student T gives some assistance in locating the error.

(j) Students U and V discuss some general matters of data-structures in relation to the assignment, and then separately implement their solutions.

Recommended penalties.

(1) First offence: a negative mark (e.g. -50%) for that assignment.

(2) Second offence: failure in the course.

(3) Repeated offences (across C.S. courses): formal disciplinary action.

At all stages, a student has right-of-appeal first to the lecturer concerned then to the Head of the Department. Beyond that, a student can opt to have the matter settled by the University via the established channels for formal disciplinary action (see below).

Formal Disciplinary Action. The University has established procedures for dealing with matters like plagiarism. These are described in the document "Malpractice in Examinations", copies of which are available in the department office. Note that all assessible material, including assignments and quizzes, form part of the "examinations" in the subject. Such procedures start with the matter being reported to the Dean, and proceed through a number of levels of hearings and appeals. Penalties available (on a guilty verdict) range from a reprimand to expulsion from the university.

March, 1989.

The DCS policy on plagiarism promulgated in 1989

Brendan McKay: Is a tough policy on plagiarism workable?

There used to be a policy on plagiarism which specified negative marks for assignments done by cheating. That was my idea. [A copy of the policy appears on Page 169.] After some years, the negative marks were abolished on advice from the university lawyers, I believe.

I recall that the plagiarism policy I initiated was immediately plagiarised by other ANU departments.

An anecdote: I accused two students of writing a program together. Turns out one of them was the son of the ambassador of a foreign country (let's not report which country!). I got a phone call from the embassy inviting me to go there to discuss it, but I declined to go and instead offered to discuss it at ANU. A few days later, there was the ambassador in my office prepared to do battle for his son. He didn't know the first thing about programming, so it was a tall order to explain why it was practically impossible for two students to independently write programs almost line-by-line identical. Eventually we agreed that I would send him the two programs (anonymised) and he would ask a friend in a CS department of another university to look at them. Naturally, I never heard back. I hope the poor kid wasn't banished to that country's equivalent of Siberia.

Once in A02 a submitted program wouldn't compile. Inspection revealed that it still had the mail header from another student.

Whatever sophisticated tools may be used to detect copying, there is no guaranteed solution to the problem of students commissioning custom solutions from an assignment solution provider, or paying substitutes to sit exams. Justin Zobel presented a paper[23] at the Australasian Computing Education Conference in 2004, describing wholesale, industrial-scale cheating of these types which had been detected at RMIT University. To me, he described the anger and outrage of a student who was detected handing in a copied assignment. The student was angry because he'd paid a large sum of money for a solution guaranteed to be unique!

> **Chris Johnson**: Although the penalties for plagiarism have changed under university scrutiny, the useful definition of what practices were unacceptable was a notable advance on definitions of plagiarism that came from essay writing experience.

5.13 What's Changed Over the Decades?

Although a large body of fundamentals remains constant, it would be foolish to expect that computer science teaching in 2021 would be unchanged from 1971. This section looks at causes of change.

5.13.1 Evolving Attitudes to Service Courses

These days, now that the right of computer science to exist is scarcely in doubt, there is a view that service courses may be important. In just about every discipline, use of computers is universal, and some form of coding is common. Why shouldn't DCS put these non-CS students on a good path?

In contrast to many other universities, when ANU introduced its engineering program in 1990, the new faculty decided to use the two introductory CS units in common for computer science and engineering students, adding C as a working language further along in the engineering course.

For a period recently, DCS offered a complementary pair of first year units aimed at students from different colleges:

The Art of Computing: Teaching programming concepts via SNAP!, a visual, drag-and-drop programming language. This unit was similar to Berkeley's *Beauty and Joy of Programming*.

The Craft of Computing: Teaching how to combine existing computational tools to get things done quickly and easily. Uses Git and Python.

Alistair Rendell tells the story:

[23] *"Uni Cheats Racket": A Case Study in Plagiarism Investigation.* https://people.eng.unimelb.edu.au/jzobel/fulltext/ace04z.pdf

> When I took over as Director in 2013 we were in the midst of a discussion around the curriculum. Two challenges were how to engage, inspire and show the value of computing to students from other Colleges, and the organisation and narrative around our existing undergraduate programs. The former saw the creation of two new first year offerings entitled "The Art of Computing" and the "Craft of Computing." The latter saw a review of all our programs to align with the 2013 ACM Computing Curriculum and the development of clear and different learning outcomes for each of the Bachelor of IT, Bachelor of Advanced Computing, and Bachelor of Software Engineering.
>
> While the curriculum activities above saw the creation of two topics specifically targeting a non-computer science audience, we had also for a number of years been offering an introductory programming unit called "Programming for Scientists". This was based around Python and although originally developed for the physics students was gaining traction in other areas of science, notably biology. In the event this unit grew significantly and overtook the "Art" and "Craft" offerings, in part because of the more obvious name but also the lack of space in the degree programs for many other students (particularly those in double degrees). When I stepped down as Director at the end of 2019, Programming for Scientists was running in both semesters with enrolments in the hundreds.

5.13.2 Changes in Curriculum Due to Advances in Technology

In 1971, there was no prospect of a computer on a single chip, and electronics magazines published designs for home-build computers. It was therefore appropriate that the curriculum included elements of logic design, including the important idea that all the logical and arithmetic instructions in a computer could be constructed from NAND or NOR gates. The emphasis on the lowest level building blocks of computers has naturally declined over the decades.

In 1971, computer communications essentially consisted of local point to point links – there were no networks in the current sense of the word, no Internet, and no world wide web. Over the years, units have been introduced to cover all the issues raised by the advent of these developments, and I presume that modems no longer find a place in the curriculum.

In 1971, E.F. Codd had only just proposed a relational model for databases[24] Since then databases have become a critical part of government and business operations, and an important part of the CS curriculum.

In 1971, there were (effectively) no electronic documents and no online information. It is now possible, and useful, to teach techniques for summarising, locating, indexing and searching textual content. Now that so much information, including personal information, is online, and vast volumes of user interaction data are routinely collected, security, privacy, and ethics become important topics for the curriculum, and so do data mining and machine learning.

In 1971, graphic displays were very primitive. As they became more powerful and more common, it made sense to teach basic principles of graphical rendering on bitmapped displays. Now, graphics hardware has become so powerful, that teaching can focus on higher levels, and a new topic, that of programming GPUs, comes to the fore.

In 1971, parallel computers were a rarity. The IBM 360/50 and its successor the Univac 1108 each had only one CPU. Now every modern laptop and smartphone features CPU chips with multiple parallel cores, and advanced computers with mixed shared and distributed memories are common enough. There was little need in the 1970s for units on high performance computing (basic numerical methods sufficed) or parallel programming (apart from interrupt handlers) but they're important topics now.

In 1971, tools supporting programming activities were very primitive:

- Timesharing wasn't introduced at ANU until 1972! Turnaround on batch jobs was so slow that it was inefficient to rely on "the computer" to get rid of syntax errors.
- The first source code control system `SCCS` wasn't invented until 1972.[25]

[24] A Relational Model of Data for Large Shared Data Banks, *Communications of the ACM*, 13(6), June 1970.

[25] By Mark Rochkind in SNOBOL4 on an IBM 370.

- Integrated Development Environments (IDEs) didn't exist until the 1980s. *Think Pascal* for the Apple Macintosh was an early example. Peter Bailey worked at Object Technology International (OTI, founded in 1988), the company which developed Eclipse.
- Although WIMP (Windows, Icons, Mouse, and Pull-down menus) interfaces arguably had their origin in the Doug Englebart "Mother of all Demos" in 1969,[26], or at least at Xerox PARC in the early 1970s, ANU students had little access to interactive windowing environments until the 1980s: X11 and Apple Macintoshes began life in 1984, Windows in 1985.
- Compilers typically contained bugs. The fact that such bugs are no longer so obvious may be evidence that software development standards have improved in 50 years.
- Some operating systems may have allowed one logged in user to send a message to another, but the advent of email, conferencing software (with video and screen-sharing capability), and collaboration tools, occurred much later.
- Systems for preparing and formatting electronic documents were essentially non-existent but became available and were improved over the years. Literate programming (e.g. Tangle/Weave) was proposed by Donald Knuth in 1984.
- There were no online resources like StackOverflow or Wikipedia. Now a student can find online information on just about any topic. They can type an error message, e.g. `AWAIT/DEACT AMBIGUITY`, into a search engine like Bing or Google and expect to find an explanation and a list of likely causes.
- There were no tools for continuous testing, like Jenkins,[27] or resources to run them.
- There had been no reason to invent the CI/CD (Continuous Improvement / Continuous Deployment) model which has become widely adopted in many organisations.

The advent of powerful program development tools, and their evolution over the years makes the experience of doing CS assignments in 2021 a vastly different experience to that of 1971. The combination of all the above changes has changed what is taught, what assignments are set, and what is expected of students.

[26] https://en.wikipedia.org/wiki/The_Mother_of_All_Demos

[27] https://www.jenkins.io/

Chapter 6

Other Research Groups

A lot of research undertaken by members of DCS has been mentioned in various ways under different headings. This chapter briefly outlines research conducted in other organisational units of ANU and in collaborative organisations in which ANU was a partner.

6.1 Computer Sciences Laboratory

The Computer Sciences Laboratory (CSL) came into existence in the Department of Engineering Physics, RSPhysS in 1985, with the appointment of Richard Brent as Chair. Other academics included Iain MacLeod, Terry Bossomaier, Heiko Schröder, and Bruce Millar, who acted as head between Richard's departure in early 1998, and the arrival of John Lloyd later in the year. Terry Bossomaier, author of books on bridge playing,[1] later became a professor at Charles Sturt University, and Heiko Schroeder became professor and head of school at RMIT University.

In 1994 RSISE split off from RSPhysS and CSL became one of its founding departments. John Lloyd describes the situation in CSL during the 1998 – 2007 period of his headship:

> The Laboratory was quite small in 1998. Richard Brent had left in February to take up a chair at Oxford. Bruce Millar was Acting Head, the late E.V. Krishnamurthy was a Senior Fellow, and Markus Hegland, Steve Roberts, Michael Stuart, and Bing Bing Zhou were Research Fellows. Joe Elso and Arthur McGuffin were IT staff members and Michelle Moravec was Administrator.
>
> By April 1999, CSL had grown substantially, by dint of John's use of two available postdoc positions. He invited the Automated Reasoning Group (ARP – John Slaney, Raj Goré, and Bob Meyer) and the Machine Learning Group (MLG) from the Department of Systems Engineering (Peter Bartlett, and Jon Baxter) to merge with CSL in return for one each of the postdocs. Thus CSL was set on the path of having the main research areas of logic and machine learning. In later years, there was also an emphasis on artificial intelligence.
>
> At its peak, CSL had 12 academic staff funded by RSISE, of whom three were women. Unfortunately, two of the women left months after the peak of three was reached. One of the departing women had been appointed to a position advertised as women-only, as one of John Richards's (Director, RSISE) tactics to recruit women.
>
> From 1999 to the present day, CSL and later RSCS had modestly-sized, but internationally prominent, research groups in logic and machine learning. In addition, over the last 15 or so years, research in artificial intelligence more generally has become prominent. As well as the researchers currently at ANU, there have been a number of others who have been on the staff at CSL and have gone on to make a significant impact internationally. To mention just two: Alex Smola, who was on CSL staff between 2001 and 2004, and at NICTA from 2004 – 2008, is a research superstar with an amazing Google Scholar h-index of 130 and more than 147,000 citations (as at February 2021). He is now Distinguished Scientist/VP at Amazon Web Services. Peter Bartlett, who was in the Department of Engineering in RSISE from 1993 – 1999 and CSL from 1999 – 2003, is now a Professor at UC Berkeley and a very distinguished machine learning researcher.

[1] https://www.bridgeshop.com.au/book/authors-a-f/terry-bossomaier-books.html

After John stepped down in 2007, Bob Williamson took over as Head until CSL merged with DCS to form the Research School of Computer Science in the College of Engineering and Computer Science.

6.2 The Automated Reasoning Project

John Slaney describes the origins of the Automated Reasoning Project within the Research School of Social Sciences (RSSS).

> **Dave:** My understanding of it is that Michael McRobbie did his PhD in Philosophy, RSSS, presumably under Bob Meyer, and then they together founded ARP. Is that roughly correct?
> **John:** Roughly correct, yes. Michael was a PhD student in Philosophy RSSS at the same time as Errol Martin, Steve Giambrone and myself. Adrian Abraham was another student who arrived just before I left (in 1980) and Paul Thistlewaite a little later again. The academics in the logic group were Richard Sylvan (then Routley) and Bob Meyer, and a stream of visitors including Mike Dunn, Newton da Costa, Graham Priest, Ross Brady among others. I think Chris Mortensen was a short-term postdoc there as well.
>
> McRobbie did a postdoc in Melbourne in the early 1980s, and returned to ANU around 1983 or 1984. He and Bob put together the proposal for an Automated Reasoning Project. I still have a copy of that document somewhere. The word "logic" does not occur in it, though "reasoning" is on nearly every line! That was a 5-year project, funded (I think) by RSSS, and it ran from 1986 to 1991. I was to have joined it as a Level B researcher in 1987, but with a few delays didn't arrive until January 1988. Errol Martin was already here as another Research Fellow, and Rod Girle (more senior) came on secondment from Griffith almost immediately after me. Chris Brink came later, if I remember correctly – he was certainly here in 1990. Bob was the notional leader, but of course Michael was actually running it. Richard Sylvan and Val Plumwood[2] were around at the time, but not members of the project itself. I was actually on leave of absence from the University of Edinburgh from 1988 to 1991, but resigned from Edinburgh when ANU offered to extend my contract. There were two programmers, John Barlow from the start and Gustav Meglicki later, a research assistant Jocelyn Brown, and a number of PhD students including Igor Urbas, Jacques Riche and Miguel de Castro. Igor later had a researcher position in the group, as did Mark Grundy. Again visitors were important: in particular, Joerg Siekmann and Hans-Juergen Ohlbach spent significant time here.
>
> CISR started around 1989-1990 and grew into McRobbie's empire, so he put me in charge of the ARP from 1990 or thereabouts. At that time, Area 1 of the agreement with Fujitsu specified research collaboration on automated reasoning. Hajime Sawamura spent a year here around 1990, after which Bob Meyer spent some months in Japan and then Toshiro Minami continued the series of exchange visits by again spending a year or so here. That must have been about 1991–2. The Japanese collaboration, with Fujitsu Labs and later with ICOT, was one of the three most important collaborations for us; the others were with Argonne National Labs in Chicago, and with a European "Esprit" project called MEDLAR.
>
> In 1991, the RSSS project came to an end, but a few off us stayed around because we had contracts which were not tied to it. CISR gave the project a home, but had no funds for it beyond my salary. That was a really lean time, as Bob was seriously ill with a brain tumour and the last of the PhD students graduated, leaving just Mark Grundy and myself, with Mark's contract expiring in 1993. I got some ARC money to appoint Greg Restall as a postdoc in 1993, and a bit of soft ANU money to bring Raj Goré here in 1994. By then RSISE was starting up, so we came out of the woodwork and joined that as one of its foundation departments. We struggled along as a tiny research group in

[2] Famous for surviving an attack by a saltwater crocodile, as well as for her philosophy.

> RSISE until John Lloyd came in 1999 and re-formed CSL by bringing together the ARP, the rump of Richard Brent's research group and the Machine Learning people around Peter Bartlett and Jon Baxter. That put us in reasonable shape at last, after our seven lean years, and then from the start of 2003 we became the Logic and Computation research program of NICTA. By that time we had appointed Sylvie Thiébaux and Yannick Pencole (who split off to become the AI group) as well as Jen Davoren, Tomasz Kowalski and (earlier) Matthias Fuchs and Jeremy Dawson. NICTA brought in Michael Norrish, Andrew Slater (now with Google), Anbulagan, Peter Baumgartner, Jinbo Huang and Jia Meng.

6.3 Systems Engineering, RSISE, NICTA: Brian Anderson's story

Brian Anderson came to ANU as Professor of Systems Engineering in 1981, after positions at Stanford University and the University of Newcastle. Systems Engineering was part of the Research School of Physical Sciences (RSPhysS), itself part of the Institute of Advanced Studies (IAS). Every Research School[3] in the IAS was taxed by the University, in the sense of losing a standardised percentage of their budget. The center of the university ran a competition for new inititiaves which awarded continuing funding to a small number of new initiatives. Within a research school, individual groups could propose new initiatives which were ranked by the school, with the top one or two chosen for submission to the central competition. As a result of these two schemes, a proposal led by Systems Engineering in the first instance advanced a successful proposal resulting in the school creating the Computer Sciences Laboratory, with Richard Brent as foundation professor.

When John Carver stepped down as Director, RSPhysSE and became DVC, he asked Brian to write a proposal for the creation of a Research School of Information Science and Engineering. This was a challenge because the existing research schools feared the inevitable shrinkage of their slices of a fixed-sized pie. Brian successfully argued the benefits of having clearly visible research strengths in ICT when the Institute of Advanced Studies was next reviewed.

RSISE was established in 1994, and Brian received a letter from the university advising him of his appointment as Director, even though he hadn't applied and there had been no selection process. RSISE did well in the competitive resource allocation within the Institute of Advanced studies, and early on received funding for Telecommunications, and brought in CISR and the ARP. It came to comprise three strong departments: Telecommunications, Systems Engineering and Computer Sciences.

Since 1972, Brian had been involved with the Australian Government, through the Australian Research Grants Council (ARGC, later became the Australian Research Council, ARC), the Australian Science and Technology Council (ASTEC, later became the Prime Ministers Science, Engineering and Innovation Council, PMSEIC), the Australian Industrial Research and Development Incentives Board (AIRDIB), and the Australian Academy of Science. He had also known many of the Chief Defence Scientists. He was thus well prepared to chair a review for PMSEIC, of the state of ICT research and commercialisation in Australia. The report was entitled: *Driving the New Economy: Australia's Information and Communications Technology (ICT) Research Base*

Brian's review concluded that ICT research was at a dangerously low level, with the conclusion supported by a number of metrics such as patent rates, number of postdoctoral positions relative to other disciplines, and ratio of university R&D to industry R&D. He presented to PMSEIC. Minister Alston and DoCITA (Department of Communication, Information Technology, and the Arts) were supportive and the decision was taken to fund a centre of excellence in ICT research, using a competitive selection process. In 2001, John Howard organised an event in Redfern to promote the then Government's Backing Australia's Abilities program. Two big-ticket items were announced: Centres of Excellence in ICT and in Biotech.

Brian's view was that the ICT Centre would have to draw on existing strengths. It seemed that good ICT researchers were primarily located at ANU, UNSW and the University of Melbourne and

[3] Until 2005, the name "Research School" was applied to the highest level organisational units within the university – at the level of the present Colleges.

Brian tried to set up a consortium involving those organisations. The Victorian Government was concerned to make a more Victorian-centred organisation led by the University of Melbourne, and was unwilling to participate. The NSW and ACT governments saw great merit in the idea, and their inclusion (from the earliest stages) served to demonstrate to the Commonwealth the willingness of other parties to put up considerable money.

Eventually the NICTA consortium, comprising ANU, UNSW and the governments of the ACT and NSW, was successful in the competitive process. Bob Williamson played a major role in preparing the bid along with Brian, providing, as Brian says, "a valuable source of contrary opinions." After the success of the NICTA bid, the University of Melbourne joined as a member, with Victorian Government support, and several other universities joined as partners.

Brian agreed to be CEO for the first year of operation. After that, Mel Slater was appointed as CEO, and Brian became Chief Scientist. Brian also remained as the Minister's nominee to receive the NICTA funding. When David Skellern took over as CEO in 2005, Brian relinquished all executive roles, functioning only as leader of individual projects. [Editor's note: Hugh Durrant-Whyte subsequently took over as NICTA CEO from 2010-2014 and then Duane Zitzner from 2014-2015. Bob Williamson was Interim CEO of NICTA from August 2015 until it was wound down.]

NICTA formally commenced operations in 2002 and wound down in 2015, merging with CSIRO's ICT Centre to become Data61. Brian regrets the end of separate funding for NICTA, though he sees a significant improvement in the metrics which were used to justify its initial funding. He says that a positive feature of NICTA was a monthly processing of project proposals, allowing the organisation to be agile and responsive.

Richard Brent recalls that Brian Anderson was very keen to encourage PhD students to finish quickly. He says that Departments in RSISE were given a sum of money for each PhD student in their first three years, but nothing in the fourth year and a negative amount in years after that.

6.4 Robotics

As described on Page 118, Hanna Kurniawati is currently leading a robotics project in DCS. Hanna outlines the breadth of current robotic research at ANU:

> There's a wide range of involvement with robotics at ANU. As you might know, Rob Mahony works in control of robotics and hardware development of drones.
>
> Uwe Zimmer is also working in robotics. In fact, he runs the Robotics Lab in DCS (level 2, CSIT). There's also the HCC group (Priscilla Kan John and Tom Gedeon) who are looking into some forms of human-robot interaction.
>
> Jochen Renz's spatial reasoning work is also very much aligned with knowledge representation in robotics.Then, there's Patrik Haslum and Felipe Trevizan who did some work on applications of robotics a while back. Then, there's application of robots for drug discovery at the John Curtin School of Medical Research. Then, there's 3Ai who is looking at robot art and incorporating robots in education. If we want to see robot based on its proper definition — an agent operating in the physical world — rather than being confined into thinking of robot as of a particular form, we have even more people.
>
> For instance, Sylvie Thiébaux, Felipe Trevizan, and Charles Gretton's project with Airbus would also be an application of planning to robotics. Then, there's the TAROS project in the School of Astronomy & Astrophysics (in this sense, the telescope is the robot). In a nutshell, there's a lot of people at ANU who are looking at different aspects of robotics, from work on improving the robotic system itself (its control, its planning, its perception, etc.) to a variety of applications (in education, medicine, astronomy, etc).

6.4.1 Robotics Group, Systems Engineering → Seeing Machines

The Robotics Group within the Department of Systems Engineering was started by Alex Zelinsky in 1996. He's now the Vice-Chancellor at the University of Newcastle, but was kind enough to answer my emailed questions, and follow up with a phone chat:

Dave: Wikipedia says you were appointed as Professor of Systems Engineering at ANU in 1996, what led up to that appointment?
Alex: I came back from AIST in Japan to the University of Wollongong, Department of Computer Science in March 1995 as a Senior Lecturer, and then joined ANU in October 1996. I was recruited by John Moore (head of Dept of Systems Engineering) and Brian Anderson (director of Research School of Information Sciences and Engineering). Actually I was recruited as a Senior Fellow but soon promoted to Professor.
Dave: How big was Systems Engineering at the time?
Alex: It was quite small, about 6—8 academics. But academically very strong — in addition to Brian Anderson and John Moore, it had two other outstanding academics in Bob Bitmead and Peter Bartlett. Both Bob and Peter went on to academic leadership, Bob – University of California (San Diego) and Peter – University of California (Berkeley).
Dave: You led a robotics group in Systems Engineering, who else was in it?
Alex: I established the group — recruited quite a few academics over the 5 year period:Roy Featherstone, Uwe Zimmer, David Austin and Richard Hartley, and some great post docs.
Alex: What were you working on?
Dave: Robotics for the real world – focused on mobile robots, robotic arms, driverless cars and submersible robots. With a strong emphasis on computer vision, and active vision systems.
Dave: What robots did you have?
Alex: Some we bought, some we built.

- Mobile robots. The Yamabico — I brought two from Japan and helped build them there.
- The Nomad mobile robot — we bought it from the USA.
- The Barrett robotic arm – we bought it from the USA.
- The Toyota Land Cruiser — driverless, was built by the lab.
- The KAMBARA submersible robot was built by the lab [Dave: I chatted with Alex on the phone while standing next to the pool outside the Brian Anderson building which was used to test this robot.]
- The CeDAR active vision system was built by the lab.

Dave: What computing hardware / software did you use?
Alex: We programmed on Unix boxes in C++ and also programmed FPGA cards.
Dave: What were the major discoveries which led to the founding of Seeing Machines?
Alex: Real-time algorithms for tracking faces under all lighting conditions for all types of people -– irrespective of age, gender, facial hair, eyeglasses.
Dave: Who were the founders and investors in Seeing Machines?
Alex: Founded June 2000. Founders 4 people – Alex Zelinsky (prof), Sebastien Rougeaux (postdoc), Tim Edwards (research engineer), and Jochen Heinzmann (PhD student). Investors were: Angel investors $500k, ACT govt $300k, ANU $100k, Volvo $100k
Dave: Did you remain doing robotics work in Systems Engineering after the founding of Seeing Machines in 2000?
Alex: I worked 1-2 days a week for ANU from 2000 to 2004, when I resigned and joined CSIRO. During that time I supervised research students and held an ARC grant.

We used to work hard, but if there was a big snow fall, the whole lab of about 15 people would take off for the day to ski. It was great for team building and a good way to have a break. Similarly, in summer — we would take off to Bateman's Bay.

6.4.2 Australian Centre for Robotic Vision

ANU is a participant in the Australian Centre for Robotic Vision which was established in 2014 with ARC funding for seven years. See the entries for Stephen Gould and Hongdong Li (Pages 113 and 119) for more information.

6.4.3 The ARC Special Research Initiative on Implantables (Bionic Eye)

ANU was a participant in this initiative which received a total of $51M of funding starting in 2010. The entries for Nick Barnes and Hongdong Li (Pages 106 and 119) give more information.

6.5 The ACSys CRC

In 1990 the Hawke government established a program of Cooperative Research Centres (CRCs) to encourage research institutions to collaborate with each other and with industry. Led by Michael McRobbie, Robin Stanton, and John O'Callaghan, ANU and CSIRO, with Fujitsu, DEC/Compaq, Sun and StorageTek as industry partners, submitted bids for an "Advanced Computational Systems CRC" in each of the first three rounds. Paul Thistlewaite also worked hard on developing the bids. Third time lucky![4]

John O'Callaghan reflects on the ACSys CRC

ACSys was the result of considerable effort from Michael McRobbie, Robin Stanton and me, involving key researchers in DCS and DIT and building on the strong relationships that we had with the industry participants.

The core participants in ACSys were ANU (DCS) and CSIRO (DIT), DEC, Fujitsu, Sun MicroSystems. The main research areas for collaboration were advanced server technologies, data mining, digital media libraries and virtual environments.

At the start of ACSys, the key participants from DCS included Chris Johnson and Malcolm Newey (distributed computing), Andrew Tridgell (parallel systems), Paul Thistlewaite (multimedia) and David Hawking (digital media libraries). Later projects were led by Steve Blackburn and Paul Mackerras (advanced server technologies), Brian Corrie (virtual environments) and Markus Buchhorn (advanced research networks).

The success of ACSys was largely due to the willingness and capability of DCS and DIT to provide outstanding researchers and to conduct the research programs on the ANU campus.

ACSys was an important factor in the Federal Government providing funds for a National High-Performance Computing Centre which led to funding for the Australian Partnership for Advanced Computing (APAC) in 1999.

According to *Cooperative Research Centres (CRC) Program – CRCs over time*[5] the objective of ACSys was "To focus on the development, integration and application of advanced information technologies", and it received $12.95 million in Commonwealth funding.

Once ACSys was established, Michael McRobbie became director and Robin Stanton and John O'Callaghan were co-directors. Keith France was the initial executive officer, followed by Peter Langford. Jan Bitmead was communications director. Later Darrell Williamson took over as director.

In essence, the ACSys commercialisation model was that ANU and CSIRO researchers would work on exciting research projects and the commercial partners would watch carefully and pounce on any valuable ideas they wished to commercialise. It eventually turned out that the Australian arms of the multinational industry partners had very little ability to influence what the parent company wanted to commercialise. The reason that the companies wanted to be involved was mostly in order to satisfy the government that they were good corporate citizens, supporting Australian R&D.

6.5.1 ACSys PASTIME project

One of the first projects to get underway was PASTIME (PArliament Sound, Text, and Image Environment) led by Paul Thistlewaite, who had previously worked for PISO (the Parliamentary Information Systems Office). Steve Ball and Jason Haines were employed on the project and Nick Craswell's PhD was supervised by Paul. I was seconded as a researcher 40% of my time. This was in the very early days of the world wide web, and Paul created a stunningly impressive demonstrator, based on the newly available Mosaic web browser from the National Centre for Supercomputing Applications (NCSA) and the Conseil Européen pour la Recherche Nucléaire (CERN) httpd web server.

[4] In one unsuccessful year, it was claimed that ACSys was on the successful bids list when it went into the minister's office but not when it came out.

[5] `https://www.business.gov.au/-/media/Grants-and-programs/CRC/Cooperative-Research-Centres-over-time-PDF.ashx?sc_lang=en`

1998: Members of the PASTIME/WAR team: Jason Haines, Paul Thistlewaite, David Hawking, and Nick Craswell. ***Photo: Unknown ACSys***

Paul set up a procedure by which Hansard documents in WordPerfect format were transferred to servers at ACSys. There they were converted into HTML pages. Visitors to *pastime.anu.edu.au* could navigate to material in Hansard using a graphical hyperlinked calendar showing sitting days for the Senate and the House of Representatives, or they could perform a full text search. When a visitor clicked on a link to a section of Hansard, the page was generated from the basic HTML using a `flex` program to dynamically recognize references to members of parliament and references to legislation. In the former case the target of the link was to the member's biography, and in the latter to the text of the legislation. If you clicked on a timecoded link in the Hansard page, you could see the video of the session in parliament, starting at the relevant point, "I withdraw Mr Speaker. The honourable member DOESN'T have the brains of a sheep."

After operating the system as a demonstrator for several years, PISO decided to go to tender for a supported version of the system. Fujitsu Australia obtained the right to commercialise PASTIME and responded to the tender. Surprisingly, PISO opted not to continue with PASTIME, and contracted with another supplier. We continued to run `pastime.anu.edu.au` for at least a year to cover the transition to the new system. Even after we turned it off, complaints and queries kept coming in!

6.5.2 ACSys TAR and WAR projects

Text And Related (TAR), and Web And Related (WAR) projects were successors to PASTIME. The initial PASTIME demonstrator made use of FreeWAIS as the text retrieval engine, but Paul, Nick and I were very interested in text retrieval. My retrieval system PADRE (Parallel Document Retrieval Engine) initially ran only on the AP1000 and was thus not suitable for PASTIME. Over time PADRE morphed into a system which would run on networks of workstations, with a focus on maximising the performance from a single node. Eventually, with developments in hardware, and code improvements, PADRE could handle large volumes of text on a single workstation, orders of magnitude faster than FreeWAIS.

The team developed expertise in evaluating the quality of search engines, including web search engines. We created new test collections, including VLC (Very Large Collection) and VLC2, which were used in the TREC conference,[6] and we employed casual relevance assessors.

We earned $1000 or so through a consultancy with DEC/Compaq to evaluate the effectiveness of the Alta Vista retrieval software, and also proposed an evaluation of the Alta Vista web search engine. DEC/Compaq in the US said, "that's a good idea," and arranged for their own staff to do it!

But we went ahead and ran independent evaluations of the quality of web search engines, including embryonic Google, and presented the results at the 2000 Infonortics Search Engines Meeting in

[6] `https://trec.nist.gov`

Boston. This caused quite a furore!

In 1998, I completed a PhD by Published Work, and joined CSIRO Mathematics and Information Sciences, but I maintained a 40% involvement ACSys. Responding to pressure to commercialise from both CSIRO and ACSys, Peter Bailey (employed by ANU as a contribution to ACSys), Tim Potter (employed by ACSys), and I worked flat out to turn PADRE into an enterprise search product. We initially called it S@NITY and launched it on the `anu.edu.au` domain (with both external and internal versions of the searches) in July 1999.

29 July 1999: Launch of the S@NITY search service for ANU. Vice Chancellor Deane Terrell and ACSys Director Darrell Williamson donning S@NITY t-shirts after the opening speeches. *Photo: ACSys*

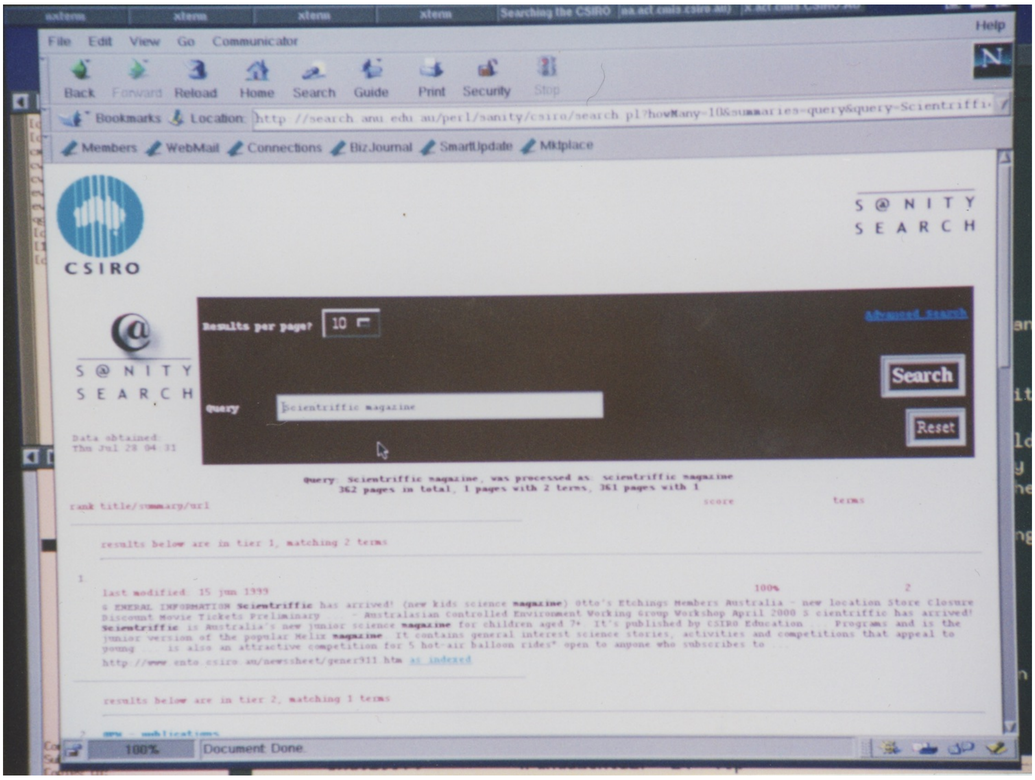
xterm
xterm
xterm
Searching the CSIRO
File Edit View Go Communicator
Help
Back Forward Reload Home Search Guide Print Security Stop
Bookmarks Location:
Members WebMail Connections BizJournal SmartUpdate Mktplace
CSIRO
S@NITY SEARCH
S@NITY SEARCH
Advanced Search
Search
Reset
100%
Document Done

http://search.anu.edu.au/perl/sanity/csiro/search.pl?howMany=10&summaries=query&query=Scientriffi

Results per page? 10

Query Scientriffic magazine

Data obtained:
Thu Jul 28 04:31

Query: Scientriffic magazine, was processed as: scientriffic magazine
362 pages in total, 1 pages with 2 terms, 361 pages with 1

rank title/summary/url score terms

results below are in tier 1, matching 2 terms

1
last modified: 15 jun 1999 100% 2
G ENERAL INFORMATION **Scientriffic** has arrived! (new kids science **magazine**) Otto's Etchings Members Australia - new location Store Closure Discount Movie Tickets Preliminary ... - Australasian Controlled Environment Working Group Workshop April 2000 S cientriffic has arrived! **Scientriffic** is Australia's new junior science **magazine** for children aged 7+. It's published by CSIRO Education ... Programs and is the junior version of the popular Helix **magazine**. It contains general interest science stories, activities and competitions that appeal to young ... is also an attractive competition for 5 hot-air balloon rides* open to anyone who subscribes to ...
http://www.ento.csiro.au/newssheet/gener911.htm as indexed

results below are in tier 2, matching 1 terms

29 July 1999: The S@NITY search interface. *Photo: ACSys*

29 July 1999: Guests at the S@NITY launch. Roger Clarke, Chris Johnson. *Photo: ACSys*

29 July 1999: Guests at the S@NITY launch. Tony Boston and Monica Berko, National Library of Australia. *Photo: ACSys*

6.5.3 Other ACSys Projects

The below summary of other ACSys projects was compiled by Jan Bitmead from an incomplete collection of annual reports. Note that individual *projects* were grouped into *programs*.

1996-97 Annual Report

Advanced Server Technologies – Safety, scalability and persistence of distributed object systems – Research students: Peter Bailey, Stephen Fenwick, Stephen Roderick, Steve Blackburn, David Miron, Brian Murphy, Sam Taylor, Kevin Moore, Stephen Barrass, Bill Keating, Sergey Bakin, Craig Eldershaw, Nick Craswell, Matthew Hutchins, Rochelle O'Hagan, Scott Milton, Roongtawanreongsri Sompat, Martin Schwenke, James Hercus, James Heath, Katrina Kerry, Jesudas Mathew, Craig Patten, Andrew Silas, Alison Hume, Brendan Humphries, Thomas Larner, Kent Skinner.

HeROD	Chris Johnson, Stephen Fenwick
JAWS	Malcolm Newey, Jim Grundy
MultiComputer Persistent Java	Chris Johnson
Parallel IO for Multicomputer Systems	Andrew Tridgell

Data Mining – Computational and Statistical Methods for Data Mining Applications – Markus Hegland, Steve Roberts, S Bakin, Craig Eldershaw, I McIntosh, D Miron, M Ng, K Skinner (honours Student), B Turlach, P Hall.

Data Mining in the Large	Peter Milne

Digital Media Libraries – –

TAR Project	David Hawking
On-Line Data Archives Program	John Lillyman
Virtual Environment Program	Duncan Stevenson (Rochelle O'Hagan – grad student)
NewMAPP Project	Peter Lamb

1997-98 Annual Report

New Students: William Clarke, Rajehndra Nagappan, Sam Taylor, Linda Wallace, John Zigman, Duncan Grove, Katrina Kerry, Dorian Colvin, Zhen He, Gareth Hughes, Luke Kirby, Peter Sienkowski.

Advanced Server Technologies – Chris Johnson, Robin Stanton

JAWS	Malcolm Newey
MCPJ	Chris Johnson
PIPPON	Paul Mackerras

Data Mining – Peter Milne, Markus Hegland

PARALGS	Markus Hegland
DMITL	Peter Milne

Digital Media Libraries – Graham Reynolds, Paul Thistlewaite

TAR	David Hawking
CACTUS	Jordi Robert-Ribes

On-Line Data Archives – Ken Harwick, Francis Vaughan

Virtual Environments – Kevin Smith, Paul Mackerras

3DINFO	Duncan Stevenson
COVER	Kevin Smith
HWS	Chris Gunn

NEWMAPP– Peter Lamb

ACORN – Markus Buchhorn, Tim Potter

1998-99 Annual Report

New Grad Students – Roland Goecke, Reehaz Soobhany, Alex Krumm-Heller, Habib Heydarian, Bernard Duggan, Matt Gray, Stephen Nees, Justin Waddell, Peter Buckingham, Nick Potter

Advanced Server Technologies – Chris Johnson, Robin Stanton

UPSIDE	Steve Blackburn
PIPPON	Paul Mackerras
JAWS	Malcolm Newey

Data Mining – Peter Milne, Markus Hegland

DMITL	Peter Milne
ParAlg	Markus Hegland

Digital Media Libraries – Graham Reynolds, Paul Thistlewaite

TAR/WAR	David Hawking
CACTUS II	Jordi Robert-Ribes

Virtual Environments – Kevin Smith, Paul Mackerras

COVE	Brian Corrie
HWS	Chris Gunn
HWMA	Duncan Stevenson
NEWMAPP	Peter Lamb
ACORN	Markus Buchhorn

6.6 Research Data Networks (RDN) CRC

ANU was a participant in the RDN CRC which ran 1993 – 1999. According to *Cooperative Research Centres (CRC) Program – CRCs over time*,[7] its objective was "To focus on high speed telecommunications utilising Telstra's experimental broadband network", and it received $13 million in funding. The ACSys and RDN CRCs were anounced on 14 Dec 1992 in a joint statement from the Prime Minister (Paul Keating) and Minister Assisting.[8]

[7] https://www.business.gov.au/-/media/Grants-and-programs/CRC/Cooperative-Research-Centres-over-time-PDF.ashx?sc_lang=en

[8] https://pmtranscripts.pmc.gov.au/release/transcript-8773

1992 Prime ministerial anouncement of the RDN CRC

The Research Data Network (RDN) CRC will conduct research and education programs in areas relating to communications network technologies, applications and services. It will also support network infratructure development, including the upgrade of the Australian Academic and Research Network, AARNet. The CRC will be established as a network of nodes, some of which are associated with other CRCs, with a Management Committee to foster cooperation between the different groups, to oversee the funding and reporting requirements of the CRC, and to manage access and connections to experimental broadband network capacity.

Participants

The participants in the research programs include: the CRC for Advanced Computational Systems; the CRC for Distributed Systems Technology; CSIRO Macquarie Joint Research Centre for Advanced Systems Engineering; the Department of Parliamentary Reporting Staff; the South Australian Centre for Parallel Computing; Microsoft Institute; Thinking Machines Corporation; Australian Supercomputing Technology; Softway; Film Australia; CSIRO Division of Information Technology; Fujitsu Australia; AOTC; Siemens; and the Australian Computing and Communications Institute (ACCI). (Not all participants are involved in every research program). In addition, the Australian Vice Chancellors' Committee and AOTC are expected to participate in the Management Committee in connection with their role as network service providers.

Research Activities

There are four research programs:

The **Distributed Interactive Multimedia Information Services** program and the **Distributed High Performance Computing** program will be established as Demonstration Projects of the CRC for Advanced Computational Systems. The Multimedia program will focus on systems architecture, tools and techniques for interactive multimedia access and navigation. This is an area which is of major importance for future broadband networks. The High Performance Computing program will integrate the high performance computing resources of Australia through broadband connections (as seen in the embryonic national cooperative supercomputer facility), and the management of data on the distributed computer network. The Resource Discovery program will be established as an additional program of the CRC for Distributed Systems Technology. The program aims to develop resource discovery prototypes that enable users to access the vast amounts of information scattered in repositories throughout Australia and the rest of the world.

The **Network Performance** program will involve research and development into the utilisation and performance of high speed telecommunications network technologies such as frame relay, DQDB and AI. This will assist in the design of the future AARNet. These four programs will enable an integrated approach to the solving of some important problems, as networks are faced with increasing demands in relation to capacity and services.

Network Infrastructure Development: Funding will be provided for an immediate upgrade of the existing research data network, AARNet, and to assist meeting the projected cost of providing upgraded services over the next few years. Funding will also be provided to the RDN CRC to fund access to broadband facilities such as the proposed AOTC Experimental Broadband Network, at some 10 sites. This reflects the fact that each of the recommended proposals will investigate technologies and applications which will require access to broadband testbed facilities beyond those available at present. This will make a significant contribution to the development of the next generation of network applications and services.

6.7 1993: Establishment of FEIT

In 1988 the Department of Education Employment and Training (DEET) commisioned a review of the Discipline of Engineering, conducted by Bruce Williams, former Vice-Chancellor of the University of Sydney. Taking account of its recommendations, ANU established a teaching Department of Engineering whose focus was on cross-disciplinary systems engineering rather than on traditional electrical/mechanical/civil streams. This move came in the midst of the discussions about possible ANU amalgamation with the CCAE, and was opposed by ANU graduate Mary O'Kane who ran the Engineering program at the CCAE. (Mary had been a student of Iain MacLeod, and went on to be Vice-Chancellor at the University of Adelaide, and from 2008 – 2018 was Chief Scientist and Engineer for the state of New South Wales.) Darrell Williamson was the first Head of ANU's Department of

Engineering. Early staff included Iven Mareels, Michael Green, and Bob Williamson. The first intake of students was in 1990.

In 1993, a new faculty was formed, the Faculty of Engineering and Information Technology (FEIT), comprising the departments of Engineering and Computer Science. It was formed in haste because the first intake to the four-year program was due to graduate at the end of 1993 with a Bachelor of Engineering degree. Robin Stanton was the inaugural Dean.

6.8 The School of Cybernetics

In September 2017, Genevieve Bell was appointed to head the 3A Institute. At the end of 2020, the Institute was broadened to become the School of Cybernetics, with Genevieve remaining as head. Interestingly, Genevieve spent time in the Coombs Building as a girl, while her mother worked on a PhD in Anthropology. She recalls being taken to afternoon tea, by a kind man called Nugget, not realising until much later that he was the person after whom the building was named.

Students and execs of 3Ai with some well-known people. You may recognize Brian Schmidt (ANU VC), Elanor Huntington (Dean, CECS), Satya Nadella (CEO, Microsoft), and Genevieve Bell *Photo supplied by 3Ai*

Nov 2020: Information provided by the 3A Institute

Established in 2017, the 3A Institute (3Ai) sits in the College of Engineering and Computer Science. We are a flagship initiative of the College to reimagine a new type of engineering and computing, one that is custom built and fit for the middle of the 21st century. The pulse in the 3Ai logo represents a Wi-Fi signal combined with a human ECG signal – it epitomises our mission of keeping humanity in technology.

3Ai is working on building a new branch of engineering to effectively and ethically manage the impact of artificial intelligence (AI) on humanity through better design and management of technology. The work of the institute centres around cyber-physical systems – next generation technologies characterised by edge computing, sensing technologies, large datasets and connectivity and increasingly sophisticated learning algorithms. The institute explores not just the technology, but also the social, environmental, cultural, political and regulatory contexts these cyber-physical systems are being built within. Activities of the institute include education programs focused on post-graduate studies and external training, research focused on exploring cyber-physical systems in industry settings, and novel theoretical work to support better decision-making in design, management and decommissioning of these systems. This work is aimed at equipping people for the challenges of the 21st century, so that we can all contribute to building a safe, sustainable and responsible world.

“Twenty years in Silicon Valley has left me with the distinct sense that we need to keep reasserting the importance of people, and the diversity of our lived experiences, into our conversations about technology and the future.” – Genevieve Bell.

When Genevieve returned from Silicon Valley to create the 3A Institute she put out the call for people to join her in building something new. 3Ai is bringing together diverse voices from diverse backgrounds to shape the future of engineering, opening the conversation to ensure our work will be genuinely inclusive, representation and transdisciplinary. Over half of our team at 3Ai are women, and we are committed to ensuring gender equality in all education programs at the institute. Not only that, but we strive to ensure that all course materials reflect a variety of voices and perspectives, to in turn reflect our diverse cohort of students and alumni. They come from a range of countries, including Nigeria, the United States, Nepal, Mexico, India, Iran, New Zealand and Australia; a range of disciplines, including law and policy, economics, psychology, computer science and machine learning, biology, and music. Together we are curating the conversations about our future, creating an inclusive space to help ensure that when cyber-physical systems scale, they are safe, sustainable and responsible.

6.9 Co-Lab: Australian Signals Directorate

The Co-Lab is a 15-year collaboration between the Australian Signals Directorate (ASD), ANU’s Mathematical Sciences Institute (MSI), and DCS. It commenced in 2019. It is located in, and contributed to the cost of, the Hanna Neumann Building.

ASD and ANU staff and students carry out research in Co-Lab on national security problems. Specific areas include data science, cryptography, secure communications, computing, cyber security and vulnerability research. Another focus is on the development of Australia’s science, technology, engineering and mathematics (STEM) workforce.

Chapter 7

ANU: Academic Computing

Chapter 2 covered the story of ANU computing facilities throughout the 1960s and 1970s. After 1980, there was massive growth in overall computing power, and major changes in how it was delivered. Initially, the computing power was concentrated in the Computer Services Centre (CSC), which provided operators, systems and applications programmers, and technical staff. The Director and Assistant Director played a major role in formulating campus computing policies. The CSC User Services Group provided user support and taught courses. The Management Services Group (MSG), supporting ANU's business computing needs and later known as the Administrative Services Group (ASG) was part of the CSC.

Over time, networking, student computing, and microcomputers became much more important while computing power became more diversified and more distributed around campus. Over time the CSC took on more functions and was eventually renamed as Information Technology Services (IT Services). MSG was split out as the Management Information Systems Division (MISD) in 1985.

Over time, MISD was renamed as Information Services and later as Corporate Information Services and took over responsibility for help desk, standard desktop environments, and network, audio visual, and compute infrastructure. You can read more about that in Chapter 10.

The CSC Annual Report for 1980[1], gives full details of the Univac 1100/82 configuration installed during 1980 to replace the 1100/42. This was a well-worthwhile improvement in speed and reliability, but occurred as the dominance of the central computing service was already diminishing. In 1987, the Univac 1100/82 was decommissioned.

Two very influential people in the development of academic computing facilities were Bob Landford, Director of the CSC from 1978 to 1987, and Robin Erskine. Bob had been Systems Engineering Manager for the Canberra branch of IBM. He contributed to development of ANU's IT policy throughout his tenure and led or contributed to all the major IT purchases: Univac 1100/82, DEC KL-10s, Faculties VAXCluster, MICOM terminal concentrators, Administrative FACOM systems, and the first ANU supercomputer (see the photo on Page 199). He left ANU to take up a position at Elders IXL in Adelaide.

Bob told me a number of stories about what it was like working for IBM in those days:

- If IBM found out that you had applied for a job outside the company, you were immediately dismissed and asked to hand in your keys and collect your personal belongings from security.
- Bob had been attending an IBM conference in a city in the US and was about to head to the airport for the flight back to Canberra when he received a message telling him to collect any necessary cash from the company cashier and to travel to a distant US city instead, for a fortnight of training.
- On another occasion he was phoned at home one morning and told to pack his bags for an international trip and bring them into the office for an afternoon departure for the US. Just before departure he was told that plans had fallen through and that he could go home.

[1] https://openresearch-repository.anu.edu.au/bitstream/1885/149187/24/2.1.8.18%281%29-3-1980.pdf

- Bob was faced with an intermittent failure in a mission-critical tape drive at a government IBM installation in Canberra. His engineers were unable to fix the problem but discovered that it was temperature-sensitive. A replacement drive could not be delivered for weeks. What did they do? Packed the tape drive with dry ice, refreshing it as often as necessary until the replacement drive arrived!

Robin arrived as Assistant Director of the CSC in November 1981 and was very influential in the evolution of computing on campus and in the establishment of AARnet. He took over as Director following the resignation of Bob Landford in 1987. In response to the decentralisation of computing facilities he became Director, Computer Services, rather than Director of the Computer Services Centre. Eventually, his role was renamed to Director, IT Services. In recognition of his campus role, he and his administrative assistant Joan Gillham were relocated to the Chancelry, but he much preferred to be close to his staff and to the machines. Robin retired in 2002.

7.1 Computing Policy Committee – Future of University Computing

THE AUSTRALIAN NATIONAL UNIVERSITY
A REPORT TO THE COMPUTING POLICY COMMITTEE

DR. B. P. MOLINARI
Department of Computer Science
Australian National University
Canberra, Australia, 2600

THE FUTURE OF UNIVERSITY COMPUTING 1979~1981

Part of the cover of the 1979 *Report on the Future of University Computing*, compiled by the Computing Policy Committee sub-committee on University Computing Requirements, chaired by Don Faulkner. At the time Max Neutze was Chair of the CPC. Note the very futuristic typeface.

Reflecting the increasing dispersal of computing equipment around the campus, the Computer Centre Advisory Committee was, in 1978, reconstituted as the Computing Policy Committee (CPC). It was authorised to make decisions on all major, centrally funded computing purchases. It commissioned at least two reports on the Future of University Computing, each covering a non-overlapping three year period. Brian Molinari notes that even the 1982–1984 report was typewritten.

In connection with the 1982–1984 report, an ANU delegation visited a number of local and overseas computing sites. On their return, one of them told me about their adventures. A bomb exploded in Frankfurt airport soon after they left it. They watched trees bending and breaking and debris flying during a typhoon in Tokyo. Most frightening of all was the discovery, on landing in Honolulu, that their aircraft had been piloted by a woman!

The Future of University Computing reports analyse each significant area of the university:

- documenting the available computing resources and their usage patterns,
- noting deficiencies and possible remedies with their costs, and
- recording any policy implications.

They also discuss overall trends, such as the increasing cost of software noted in the screenshot below, and make recommendations.

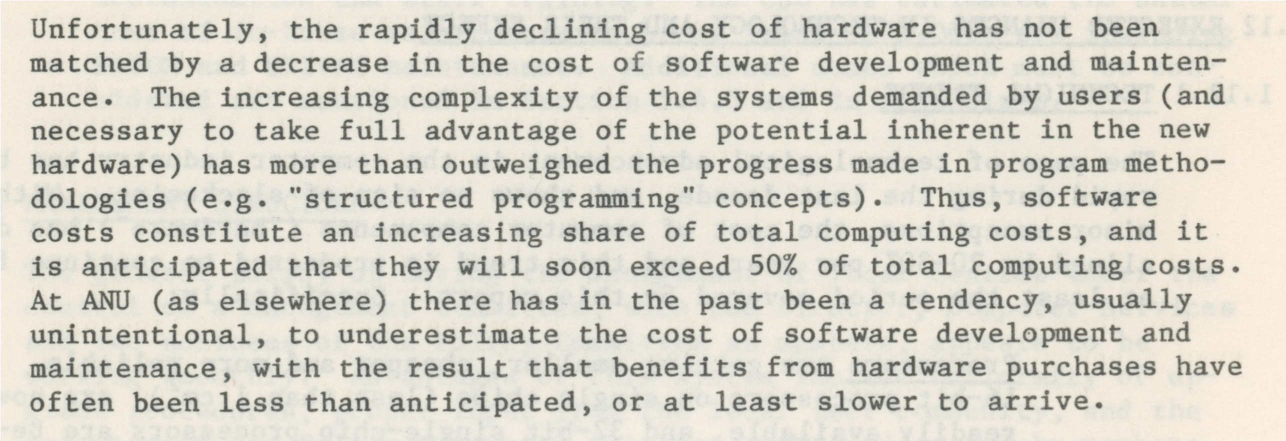

Unfortunately, the rapidly declining cost of hardware has not been matched by a decrease in the cost of software development and maintenance. The increasing complexity of the systems demanded by users (and necessary to take full advantage of the potential inherent in the new hardware) has more than outweighed the progress made in program methodologies (e.g. "structured programming" concepts). Thus, software costs constitute an increasing share of total computing costs, and it is anticipated that they will soon exceed 50% of total computing costs. At ANU (as elsewhere) there has in the past been a tendency, usually unintentional, to underestimate the cost of software development and maintenance, with the result that benefits from hardware purchases have often been less than anticipated, or at least slower to arrive.

The Summary and Recommendations volume of the 1982 – 1984 report notes the increasing cost of software relative to hardware.

A document titled *Strategy for Academic Computing 1986-1990 for the Australian National University* was published by the Review of University Computing in late 1986. Unfortunately, I have been unable to locate a copy.

7.2 1980s – A Decade of Major Change

At the end of 1979, ANU's model of academic computing consisted of a pre-eminent central facility, augmented by medium scale facilities in Research Schools and The Faculties, and accessed by a burgeoning network of dumb terminals.

This model changed dramatically during the 1980s. There was a rapid increase in the compute power accessible to individual academics and a corresponding increase in the convenience of accessing it. Instead of obtaining an account on the Univac and writing a FORTRAN program (or using the SPSS package) to process simple survey data, a sociologist could use a spreadsheet package on their own desktop computer. They could use the same computer to write and typeset their paper, copying and pasting data and graphs from the spreadsheet and printing the document to a local laser printer.

Networking changes such as AARNet and the World Wide Web which both had their beginnings in 1989, are discussed in Chapter 11.

- The importance of VAX/VMS waxed and then waned. At its peak there were several VAXes in the CSC and others in the Faculties, RSC, RSBS, RSPhysS and MSSSO, along with a number of MicroVAX IIs.
- The importance of Unix waxed throughout the decade. When running Unix on a VAX was considered, Bob Landford visited the University of California at Berkeley in order to speak

with the suppliers of BSD Unix. He was shocked to find that all that was there was a couple of ratty academic offices. One of the CSC Vaxes did run BSD 4.2 for a while but eventually two Pyramid 90x systems (CSCUnix and FACUnix) were acquired to provide Unix services, along with the one in DCS.

- There was an explosive growth in Unix servers and workstations. By 1990 the CSC annual report claimed that there were nearly 250 on campus, with over 200 supplied by Sun Microsystems. Not only did these workstations run Unix, but they had more powerful CPUs than personal computers, and they featured megapixel (1152 x 900) bit-mapped graphics displays, some in colour. By contrast, IBM's Colour Graphics Adapter (CGA) for PCs was limited to 640 x 200).
- From 1984, there was an explosion of microcomputing, particularly Apple Macintoshes.[2] ANU signed an Apple Consortium agreement,[3] and ANU itself, plus staff and students, were able to purchase Macs for discounted prices through ANUTech, ANU's commercialisation arm. Several general-access laboratories of Macintoshes were established around the campus, particularly in the libraries. There were a smaller number of general-access Amiga 500s and Windows PCs.[4]
- I was responsible for purchasing three QMS Lasergraphics 1200 laser printers early in the 1980s, but the release of the much more powerful and more affordable Apple LaserWriter in March 1985 opened the floodgates. Dozens appeared on campus. The LaserWriter implemented the PostScript Page Description Language invented by John Warnock of Adobe Systems, and was capable of printing complex graphics, mathematics, and arbitrarily scaled/rotated fonts. ANU people soon started using packages like Adobe PageMaker and QuarkXPress (and of course LaTeX) for typesetting. Desktop publishing had arrived.
- Around 1982, the CSC established a MicroComputer Information Unit, with Gloria Robbins in charge. It established and grew a collection of microcomputer hardware and software for evaluation by members of ANU staff. Among the people who worked in the unit were Rosemary Coggins, Virginia Woodland, Robert Davy, Kevin Dawes, and Melissa Waterford. Around 1986, this was made part of a larger group, headed by me, which provided micro-computer and networking advice around the campus. In that era there were many proprietary networks like AppleTalk, and Novell Netware.
- The ANU Supercomputing Facility was established in 1987, when a Fujitsu vector processing supercomputer was acquired. See Chapter 8.

7.2.1 VAX/VMS

Robin Erskine came with a great deal of knowledge of VAX/VMS systems, having been Director of the Computer Centre at the University of St. Andrews, which relied on them. He was decisive, made good decisions, and was able to keep track of masses of detail. I recall him planning a very complex programme of DEC hardware shuffles across the campus, down to the details of dozens of serial line interfaces of several different types .

Geoff Huston and I both have memories of a disaster on a VAX/VMS system, but our memories differ on whether the VAX in question was at the CSC or in the Faculties.[5] Jeff Brindle needed to remove a whole lot of files from somewhere and typed the VMS command, `DEL *.*;*` and watched the listing, in alphabetic order, of the files as they were deleted. As the names rolled through to the second half of the alphabet, he suddenly realised with horror that he'd been in the wrong directory and was deleting all the system executables. He hastily hit CTRL-C. Unfortunately the

[2] The Maths Department had earlier purchased an Apple Lisa, which tried out many of the ideas in the Macintosh, but it was too expensive for general use.

[3] Roger Clarke remembers attending a meeting of the ANU Finance (Boardman?) Committee with Robin Erskine and Neville Smythe to argue that ANUTech should be allowed to set up PCTech to support the consortium.

[4] Apple released the Macintosh in January, 1984. Microsoft released Windows in November, 1985.

[5] Just about all systems programmers were involved in at least one disaster caused by "finger trouble" but as author I'm able to forget about any I might have had.

critical program for restoring the system was BACKUP.EXE, and it started with a 'B'. There ensued a desperate racking of brains to see whether the system could be restored using only programs starting with 'M' and beyond. I don't remember the outcome.

VAX/VMS supported a wide range of privilege bits controlling access to various operations. To reduce risk of adverse events, whether malicious or accidental, one was supposed to turn on the minimum set of privilege bits necessary to carry out each operation. But that required one to remember what all the privilege bits were and which operations required them. Accordingly, cowboy systems programmers like me got into the habit of starting each session with:

```
set proc/priv bypass
```

VAX/VMS devotees considered it a very user-friendly system with mnemonic command names, guard rails against accidents, and a structured help system. VMS required confirmation of any action likely to destroy data – though it didn't help in the anecdote above. Unix lovers preferred terse commands like cat, ls, awk, and rm. Instead of rm, the VAX/VMS command for removing files was called DELETE but you could abbreviate it to DEL. Unix users poked fun at VAX/VMS users and their need for hand-holding. I remember a T-shirt which provided a Unix version of DEL for people coming from a VMS background. It was an executable script /bin/DEL:

```
    #! /bin/sh
    /bin/rm -rf $*  # Delete all the files, then ask
    /bin/echo -n "Were you sure? "
```

7.2.2 Unix

ANU was late in exploiting the capabilities of Unix, compared to the University of Sydney (Bob Kummerfeld and others), University of New South Wales (notably John Lions) and Melbourne University (Robert Elz, known as KRE). Those departments developed expertise in Unix internals, developed schedulers oriented toward student computing and contributed to the Unix code base. For example, KRE was responsible for the Unix system of managing disk quotas.

A truly remarkable Australian Unix achievement was the port by Richard Miller (with the support of Juris Reinfelds) of Unix version 6 to the Interdata 7/32 at the University of Wollongong. Going into production in July, 1977 this was the world's first port of Unix to a non-PDP platform. Juris Reinfelds's account of the port is well worth a read![6] So is Richard Miller's own account.[7]

As previously mentioned, Pat Keogh installed Unix on an ANU PDP-11 very early on, and Robin Stanton acquired two Unix-based Sun-1 workstations in the early 1980s. DCS continued to acquire Sun workstations for staff use throughout the 1980s. During the 1980s, there was a battle between ATT&T's System V and Berkeley versions of Unix (BSD) , the so-called "Unix Wars". System V tended to be favoured by commercial people and BSD by technical people.

When I returned to DCS in 1988, the Sun machines were running SunOS, a variant of BSD. Soon after that, Sun and AT&T attempted to merge System V and BSD. Sun's version of that was called Solaris. My steep learning curve to get up to speed on the system programmer's view of SunOS was very soon followed by a steep ramp-up on Solaris.

Of course rival vendors refused to accept Solaris as the one true Unix, and developed their own merged versions, forming the Open Software Foundation. And then, in September 1991, Linus Torvalds released Linux. Fortunately, the POSIX standards, which evolved from an initial version in 1988, provided a degree of interoperability.

At the time of writing, Unix has indeed conquered the operating system world, though it has many variants. Apart from many different distributions of Linux, versions of MacOS since 2001 have

[6] https://documents.uow.edu.au/content/groups/public/@web/@inf/@scsse/documents/doc/uow103747.pdf

[7] http://bitsavers.informatik.uni-stuttgart.de/bits/Interdata/32bit/unix/univWollongong_v6/miller.pdf

been based on Unix, Android and iOS phones and tablets are Unix-based, and Windows 10 now has a BASH shell built in.

In 1988, our Sun workstations were running Sun's NeWS graphical user interface. With Solaris, they changed to use OpenWindows, a GUI based on X Windows, and implementing the Open Look look and feel, defined by Sun and AT&T. Naturally, the Open Software Foundation created an alternative, known as MOTIF. By 1993, the threat from Windows caused the major Unix vendors to agree on the Common Open Software Environment (COSE) whose Common Desktop Environment (CDE) was based on Motif. Yet another change!

In the 1980s, a major Sun slogan was, "The network IS the computer." They provided support for all manner of network services, notably RPC (Remote Procedure Calls) and NFS. NFS officially stood for Network File Service, but No File Security was a good alternative.

A regular visitor to Mount Stromlo from the University of Virginia, on one occasion brought with him a Sun workstation. Unfortunately it ran far slower on Mt Stromlo than it had in Charlottesville – Different altitude? Different voltage? Different hemisphere? Hardware fault?

None of the above. The problem was that his `/etc/fstab` mounted the `/usr` partition from his usual server in Charlottesville. When he invoked an editor, compiler or any other system program, everything worked fine, but it took a while to transfer the programs from Virginia to his workstation in order to execute them.

7.3 Computer Services Centre in the 1980s

In 1987, the CSC comprised 50 staff. The break-down of staff numbers shows the diversity of its functions: Director, Assistant Director, Admin (5), User Services (6), Facom Systems Group (5), VMS&Unix Systems Group (4), Microprocessor Support Group (5), FCU (3), Operations (12), Network (3), Technical (3), Supercomputer Facility (2).

In 1987, the last Univac was transferred to the Department of Defence. It was replaced by a VAX cluster comprising a VAX 8700 and a VAX 11/785. The CSC also ran three VAXes in The Faculties, two MicroVAX IIs, two FACOM mainframes, the Fujitsu supercomputer, and two Pyramid 90x Unix systems.

The CSC believed in providing training for its staff. Many of us went off to one-day courses on generic workplace topics, such as *Effective Time Management*. Recommendations from that course included: "Keep a whiteboard and waterproof marker in the shower, and a notepad and pen beside your bed, so that you don't waste any great ideas you have while sleeping or showering." Another course was *Dealing with difficult people* which was mostly attended by difficult people, ..., and me. In an earlier era, the Commonwealth Public Service induction process included training on *How to answer the telephone*.

7.4 School Computing Facilities

The RSPhysS DEC KA-10 was replaced by a DEC VAX 11/780 very early in the 1980s after a procurement exercise in which Prime Computer was a serious competitor to DEC.

In the mid 1980s a University review into computing in the Joint Schools (RSSS and RSPacS) recommended that the DEC KL-10 not be upgraded, but rather augmented by micro-computers and networking. I was seconded for 18 months or so to work with a Management Committee on the implementation of this plan, and had the privilege of discussing requirements and possibilities with all 50 or so of the departments, centres and units within Joint Schools. We bought hundreds of microcomputers in each of two years, and gave recipients the choice of a standard Apple Macintosh or a standard IBM compatible system.

Unfortunately, the test sample of the IBM-compatible we chose in the first year were not representative of the ones delivered. There were endless problems with poor assembly and faulty components, including, in the worst case, a fire. We had aimed to spend a similar amount (about $2.5k) on

Macintoshes and IBM-compatibles. At the time, genuine IBM PCs were less capable and much more expensive. For example, the Psychology department bought an IBM PC XT configuration for more than $11k (nearly $30k in 2019 dollars). After obtaining advice that we were allowed to do so, we bought dozens of education packs for Microsoft Word for both Macs and IBM-compatibles, so that everyone had a licence, and that everyone would be using the same word processor. Despite this, one or two departments insisted on using WordPerfect, licence status unknown.

Unrelated to the unreliability of the IBM-compatibles, it was clear that staff found Macs easier to use. The Macs disappeared into people's offices and, soon afterward, the owners could be seen working away. That was often not the case for the PC compatibles. In the extreme case, one of them was dumped outside the computer room door with a note saying, "This is useless. Take it back."

In the late 1980s, RSC started purchasing Apollo DN10000 workstations, based on a RISC architecture called PRISM. The school's logic was that they could afford one such workstation per year out of their own budget and it would deliver at least as much compute power as they could get from their share of a central facility – and with high resolution graphics at their fingertips. The DN10000 was very little use for computer science teaching since the compilers were highly optimising and thus very slow. Apollo was later taken over by Hewlett-Packard who improved the performance of the DN10000 still further and marketed it as the HP Snake.

The floating point performance of the Snake was vastly better than that of the then current Sun workstations. HP claimed a huge gain in sales when they emblazoned the Snake's floating point specifications on huge billboards outside the Sun User Conference in the US, along with a telephone number for HP Sales.

7.5 The Heyday of Programmers on Campus

In the 1980s many people were employed as programmers across the campus. Central organisations like CSC, MISD, and the libraries employed quite a few and research schools and departments quite a large number in total. There was quite a bit of interaction and mutual support between staff responsible for computers of the same type. For example, the School Computer Unit, RSPhysS and the Faculties Computer Unit used to interact frequently when they both had DEC-10s. Favours were granted in exchange for shared packets of Tim Tams.

The University recognised programming as an employment stream and defined Programmer, Senior Programmer, and Principal Programmer classifications. There were different numbers of salary levels within each of these classifications. When established, the levels were based on Commonwealth Public Service classifications but, over time, ANU's drifted behind. Each year you could expect to automatically rise to the next level within a classification, but once you hit the top of the classification, a reclassification case was needed to get your next increase, based on the increase in the level of duties which you were required to perform. There were descriptors associated with each classification, which were usually very terse and quite subjective. The difference between levels often came down to whether the same portfolio of duties was carried out under "limited", "broad", or "general" direction.

Q: What's the difference between "broad" and "general"?

A: $3,500 a year!

In the 1980s, many programmers read Tuesday's Australian. Not only did it contain many articles with computing news but the vast majority of IT jobs were advertised there. One of the IT recruitment agencies, O.U. Norman Holdings, regularly inserted a prominent display ad for "Anal/Progs". That means Analyst/Programmers in case you were confused. The CSC also subscribed to the ComputerWorld magazine and it was read by some, during coffee breaks in the Leonard Huxley tea room. Interestingly, the issues kept flowing even when the subscription was cancelled, since circulation numbers were critical to advertising revenue.

The distinction between [System] Analysts and Programmers was never felt in the academic part of ANU, though it may have been in the administration side, i.e. the organisation sometimes known as MISD (Management Information Systems Division). In business computing, analysts were the people who elicited requirements, wrote specifications and flow charts, and designed testing programs. Programmers were the less skilled people who translated flow charts or other ways of specifying an algorithm into code. In that world, programmers enthusiastically seek promotion to Analyst roles. Right up to my retirement at age 65, I have been asked by friends in that world, "David, don't tell me that you still write code!!"

In another world, coding skills are highly valued. In the part of Microsoft supporting the Bing search engine, where I worked for four years, proficiency in coding was not only valued but expected of all the people in technical roles. Excluding the Vice Presidents of course, but they were highly technical too. The search quality team at Google told potential recruits, "We don't really care what your PhD was on, because you can learn all that. What we really care about is, 'Can you code?"'

Over the decades at ANU I had to prepare dozens (perhaps even a hundred) of recruitment packages, each including advertisement text, duty statement, and selection criteria. As time went on, the amount of paperwork increased, and rules about what you could say and what you could ask for became tighter and tighter. Rules about the constitution of selection panels and the genders of people who carried out short-listing of applicants became more and more rigorous. These changes were well motivated by desires to improve fairness to women, ethnic minorities, and people with disabilities, and to improve safety in the workplace, but sometimes made the process quite frustrating. One recruitment package I submitted to HR, or was it Personnel, or was it the Staff Office, was considered a perfect example and was shown to others round the campus as an example of how things should be done. The next one I submitted, following exactly the same pattern, was rejected and required masses of changes.

The selection committee rules put a load on the relatively few women who had reached higher levels in the University. Erin Brent, Kathy Handel and Melanie Rooney were frequently called upon to serve on selection panels around the campus. For more senior appointments within the CSC and elsewhere, Sue Serjeantson from the Department of Human Genetics, JCSMR was often called upon. It seemed rather embarrassing to frequently take up the time of a high-flying researcher who went on to become ANU's Deputy Vice-Chancellor when appointing to a post far below her level. (I'm told that it's still the case that female employees at universities spend more time than men on panels and committees.)

In the early days I sat on a selection panel for a CSC Computer operator. The only selection criteria were that the applicant should have an ACT driver's licence and be able to lift 16kg. We had a flood of applications including more than one with a PhD. How do you rank 45 applicants against those criteria? Pick the person who drove the fastest and could lift 150kg?

The status of programmers at ANU seemed to undergo a downward transition over the years. In the 1970s they were generally regarded as holders of valuable knowledge, trusted advisers, and creators / enablers of good things in IT. Then, a change of mood among decision makers led to an attitude of, "it's much better value to buy packages which approximate what's needed rather than to develop our own application," and a view that programmers were much more about cost than about benefit. Fortunately, this generalisation didn't apply in DCS, though there was sometimes friction between academics and programmers. In recent years, programming (now usually known as "coding") is again an honourable profession, with politicians like Bill Shorten encouraging students to learn how to do it.

7.6 Some Notable Characters

David Jobson, when working as an operator at the CSC, had the habit of going to a random office, closing the door, and conspiratorially asking, "What's this I hear about you leaving?" Apparently, several impending but unannounced departures were flushed out by this means.

David Kenny, head of the Systems Group at the CSC, reflected carefully on issues and spoke slowly. One of his pearls of wisdom was, "You can always buy your oats a little cheaper if you don't mind them having been pre-processed by a horse."

1969: Erin and Richard Brent, recently married, on the liner taking them to Stanford. ***Photo supplied by Richard.***

Erin Brent was a significant contributor to computing at ANU, as you may have noted earlier in this book. From 1972, she was Program Librarian at the Computer Centre, and later was an evangelist for graphics, being responsible for advancing and supporting CSC services in plotting and graphics. In around 1992 she became head of IT services at the ANU Library, charged with making the McDonnell Douglas library system accessible via the net. She also worked for a while at the Textronix Canberra branch, originally located in Fyshwick just next to a brothel, but later moved to more salubrious accommodation in Belconnen. She enjoyed the work but Tektronix as a company was going downhill at that time.

Erin had a large repertoire of amusing stories about computing at ANU. Writing this, I wish I could consult her. She was well-informed on technical matters and often took strong positions on how computing should be developed. Her arguments were almost always well-founded and very determinedly argued. One never repeated the mistake of thinking that an argument made in Erin's quiet voice could be ignored.

Erin (then O'Connor) completed an Honours degree in Computation at Melbourne University, one of the first women in Australia to attain such a qualification, and married Richard Brent in 1969. While Richard was completing his PhD at Stanford, Erin worked as a Teaching Assistant (for George Forsythe) and then a Research Assistant (for Gene Golub) in the Computer Science Department. She also worked for the MASCOR start-up in Silicon Valley.

After Richard took up the Chair of Computing Science at Oxford in 1998, Erin worked for the Numerical Algorithms Group (NAG). NAG, formerly "Nottingham Algorithms Group", was a spin-off company from Oxford but not officially part of the University.

The O'Connor family has been tragically blighted by breast cancer. Erin's mother and her four daughters have all been affected by the disease, and Erin succumbed in 2005 after a long illness, during which she had researched the relevant medical literature with her characteristic determination. Although it was a long time since she last worked at ANU, large numbers of ANU people attended the funeral and mourned her loss.

At the funeral a neighbour recounted Erin's attempts to involve Richard in household duties, of which the most successful was to make him responsible for the "breakfast eggs". According to the neighbour, Richard would put the eggs on and think of a problem. If the problem was very easy, the eggs were runny; If it was hard, the eggs were hard-boiled. "That doesn't seem to have been very successful", said the neighbour. "Oh yes", said Erin, "It was fine. I don't eat eggs for breakfast!"

Kathy Handel arrived to start work at the CSC in 1980, driving an Alfa Romeo replete with the bat logo of SUSS (the Sydney University Speleological Society). Ian Simpson hired her but left for England before she arrived. She was the Program Librarian, helping staff and post-grad students with mathematical and statistical programming. When we visited her house once, she played Pachelbel's *Canon in D* followed by Eric Clapton's version of *Cocaine* and I immediately recognised a refined taste in music – just like my own!

After graduating from USyd with a BSc (Hons) degree in Computer Science in 1975, Kathy worked for the Traffic Accident Research Unit, and then the Riverina College of Advanced Education before coming to ANU. In 1986 she became Head, User Services, replacing Nick Van Vucht. (Earlier people in the role included Wayne Davey, and Wayne Naughton.) Some time after I left the CSC in 1988, Kathy also took on responsibility for the Microcomputer Support Group.

2013: Julie Bakalor (L) and Kathy Handel (R) walking the Haute Route from Chamonix, France to Zermatt, Switzerland, with Mont Blanc in the background. The fact that Julie is doing this walk with her arm in a cast seems to testify to a degree of determination! *Photo: Keith France*

In 1990, Kathy left the CSC to be System Controller of the ANU Pay/Personnel system. Then in 1992 she left the ANU to work in IT support at the Trade Practices Commission. She says:

> In 1999 my job was outsourced from the Australian Competition and Consumer Commission to Advantra, where I worked on enterprise management systems for a group of government agencies. After several corporate restructures, I was made redundant again in 2007 and from then until I retired in 2013, I worked for companies contracting to Defence on enterprise management systems and network monitoring systems.

7.7 An Explosion of Davids

Toward the end of my time at CSC, the CSC Management Committee consisted of Robin Erskine as Chair, and **David** Kenny, **David** Jobson, **David** Baldwin[8] and **David** Hawking.

On another occasion I arrived late to a meeting of Faculties IT staff, plonked myself in an empty chair and found that I completed a line of four **David**s.

When the Fujitsu AP1000 (the CAP) was installed in DCS in 1991, a help team christened CAPTAIN (CAP Troubles And Inquiries Network) by Peter Bailey comprised Peter (who had a brother called David), **David** Walsh, **David** Sitsky, and **David** Hawking. What chortling there was when we replied to an application for a CAP account from one **David** vouched for by another **David** and signed the email (as usual), "**David**, **David**, **David**, and **David**'s Brother".

To quote Monty Python, it certainly avoided confusion.

7.8 Academic Computing Facilities in DCS

After the move of DCS to the Crawford[9] building in 1986 (see the invitation on Page 92) Suns started to appear in academic offices and to be interconnected with thickwire ethernet. In those days Ethernet transceivers were very expensive, of the order of $1000 if I remember correctly. They connected to the thickwire co-axial cable by means of "vampire taps" which involved drilling a hole through the insulation and outer conducting sheath into the central conductor. Reflections within the co-axial cable were a problem, and there were cases where Sun A could communicate perfectly with Sun C but not with Sun B, while B and C could talk to each other.

Many of the Suns ran diskless, supported by a Sun 3/470 server known as Vega. (All of the DCS academic Suns were named after stars.) Vega was based on a Motorola 68030 chipset running at 33MHz.

Robin Stanton managed to negotiate a very good deal with Pyramid for the supply of a Pyramid 90x as a general Unix computing resource for departmental staff and research students. It was located in an inner room on the top floor of the Crawford Building and known as Ra (the Sun god). Crawford offices which didn't yet have a Sun workstation, had a terminal connected to Ra. I'm pretty sure that both Ra and Vega had 370MB Fujitsu Eagle disk drives. For the time, Eagles offered very good storage for the money and were very reliable. I think that some of them ran faultlessly for 8 years, performing $3960 \times 1440 \times 365 \times 8 = 1.665 \times 10^{10}$, (i.e. more than 16 billion) revolutions.

I returned to DCS in 1988 as Head Programmer, at a time when thickwire ethernet was being replaced by much more manageable thinwire and Sun workstations were becoming much cheaper. The workstation of choice from 1989 was the Sun 3/80, with an 1100 x 952 monochrome display, 4MB of RAM, and based on a 20MHz Motorola 68030 chip and Motorola floating point unit in a "pizza box" base. The base also contained a small local disk, avoiding the need for diskless operation.

Within a few years, the whole of the department had been wired with thinwire Ethernet and all staff had their own workstation. I did a lot of lifting of ceiling tiles, pulling of cables, and fitting of BNC connectors. The recruitment of Eric White (formerly a technical officer at the CSC) raised the level of professionalism. When surprised, Eric was known to come out with "Well strap mah face to the sahd of a hawg and rollllll me in the murd", in a Southern accent.

7.9 The Current Situation

The default position in 2021 is that every academic at ANU should have an ITS-supported PC or Mac running an ITS-supported OS, with ITS-recommended anti-malware protection. ITS provides support for these environments through its Service Desk, however the level of support is highest for standard Windows environments. The Service Desk is accessible by all staff and students. Every staff

[8] My memory on this is a little rubbery. Perhaps the fourth David was David Jones.

[9] Now renamed after Beryl Rawson

member is issued with a standard `firstname.lastname@anu.edu.au` email account on ANU's Exchange 365 system.

A single sign-on system allows staff to perform routine HR administrative functions through systems like APOLLO, ANUBIS and HORUS.

A webpage operated by ITS allows staff (or more likely Administrators, acting on their behalf) to purchase standard-cofiguration computers from a menu of ITS-supported options.

Thanks to ANU's ubiquitous WiFi network and ubiquitous Ethernet connection points in offices, staff can bring their own smartphones, tablets or non-standard PCs and connect them to the ANU infrastructure. (This is known as BYOD, Bring Your Own Device.) There is nothing to stop staff members from running Linux on their machines, but they're essentially on their own, unless support is provided through Local IT Support Staff (LITSS). CECS used to have an IT Group but it was folded back into ITS in 2020 to reduce costs.

The School of Computing employs three IT support staff: Bob Edwards, Andrew Wilkinson, and Jie Gao. They are responsible for hundreds of student workstations, dozens of college or school servers, and dozens of project machines.

Bob Edwards recently took me on a tour of the machine rooms that host machines used by academics and students for larger scale computations. There is one room on level one of the Brian Anderson Building (split into two, one half for production servers and the other for research machines). Bob tells me that many of the servers contain four, or even eight, GPUs and that some academics have purchased gaming computers to gain lower cost access to GPU power. The other is on level two of the CS&IT Building. In addition to these, many academics run compute-intensive jobs on project-funded machines located in their offices. This practice was dramatically exposed when ANU turned off the power to offices during the COVID-19 campus shut down.

Gaming machines whose GPUs are doing important computation for the Australian Centre for Robotic Vision. They're located in a small server room in Brian Anderson Building along with dozens of other GPU-festooned research servers. *Photo: David Hawking*

It is unlikely that academics in other parts of the university would have even a small fraction of the local compute power that is deployed by CECS. However, all academics can access computing resources on virtual machines on ANU's VMWare cloud. Furthermore, they can apply for time on the NCI facility.

According to `https://services.anu.edu.au/business-units/information-technology-services`, Garry Whatley is ANU's Interim Chief Information Officer (CIO) at the time of writing. Other members of the IT Services executive team include:

- Helen Duke, (Interim) Deputy CIO, and Director, IT Services;
- Lucy Boom, Associate Director, Project Delivery and Engagement;
- Steven Fox, Interim Associate Director, Infrastructure Services;
- Lynne McDonald, Associate Director, Customer Service;
- April Weiss, Associate Director, Application Services.

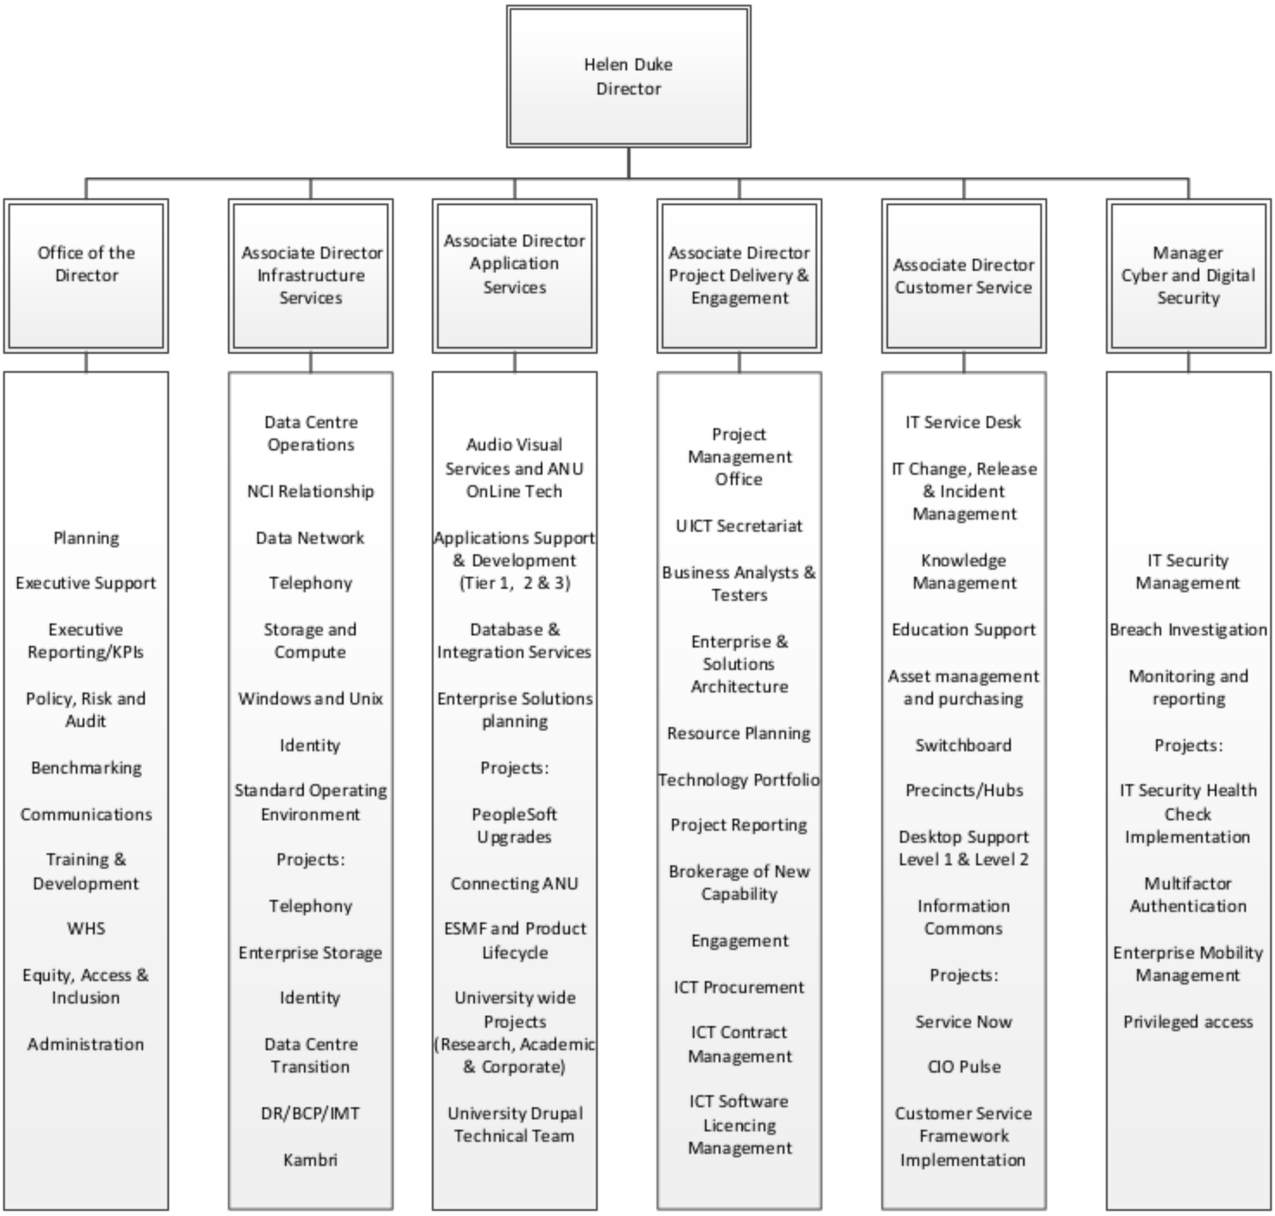
Helen Duke
Director
Office of the Director
Planning
Executive Support
Executive Reporting/KPIs
Policy, Risk and Audit
Benchmarking
Communications
Training & Development
WHS
Equity, Access & Inclusion
Administration
Associate Director Infrastructure Services
Data Centre Operations
NCI Relationship
Data Network
Telephony
Storage and Compute
Windows and Unix
Identity
Standard Operating Environment
Projects:
Telephony
Enterprise Storage
Identity
Data Centre Transition
DR/BCP/IMT
Kambri
Associate Director Application Services
Audio Visual Services and ANU OnLine Tech
Applications Support & Development (Tier 1, 2 & 3)
Database & Integration Services
Enterprise Solutions planning
Projects:
PeopleSoft Upgrades
Connecting ANU
ESMF and Product Lifecycle
University wide Projects (Research, Academic & Corporate)
University Drupal Technical Team
Associate Director Project Delivery & Engagement
Project Management Office
UICT Secretariat
Business Analysts & Testers
Enterprise & Solutions Architecture
Resource Planning
Technology Portfolio
Project Reporting
Brokerage of New Capability
Engagement
ICT Procurement
ICT Contract Management
ICT Software Licencing Management
Associate Director Customer Service
IT Service Desk
IT Change, Release & Incident Management
Knowledge Management
Education Support
Asset management and purchasing
Switchboard
Precincts/Hubs
Desktop Support Level 1 & Level 2
Information Commons
Projects:
Service Now
CIO Pulse
Customer Service Framework Implementation
Manager Cyber and Digital Security
IT Security Management
Breach Investigation
Monitoring and reporting
Projects:
IT Security Health Check Implementation
Multifactor Authentication
Enterprise Mobility Management
Privileged access

ANU IT Services: Functional Organisational Structure, October 2019

Chapter 8

ANU: Supercomputing

A supercomputer has been defined as a machine for turning a CPU-bound problem into an I/O-bound one. In the early years, supercomputers were designed to dramatically accelerate large scale scientific computations, such as those required to forecast weather or simulate molecular and fluid dynamics. ANU's first step toward providing scientific compute power beyond that available from general purpose computers, was the acquisition of SAMs (Scientific Accelerator Modules) for the Univac 1100/82. Since then ANU has acquired a series of machines specifically chosen to speed up research computations.

Standing in front of the Fujitsu VP50, perhaps marking the acquisition of a new Fujitsu mainframe. Left to Right: Bob Landford (former Director, Computer Services), Unidentified Fujitsu, Ken Vine (Head, Management Information Services Division), Ros Dubs (Registrar), Colin Steele (Librarian), Alan Barton (Treasurer), Kosaka-san (Fujitsu Japan), Neville Roach (CEO, Fujitsu Australia), Ian Ross (DVC), Harvey Jones (Head, Finance & Accounting, Robin Erskine (Acting Director, Computer Services). ***Photo: ANU***

Robin Erskine: How the decision was made to buy ANU's first supercomputer

Ian Ross was absolutely critical to the acquisition of a supercomputer. I had the impression that he already had access to the necessary $6 million [Approximately $15 million in 2019 dollars]. The target was a supercomputer to be used for high speed computation for ANU Chemists, Physicists and Astronomers.

A technical committee including me, some CSC staff and some senior academics had considered a number of options, including Fujitsu, Cray, NEC and IBM. Denis Evans and I had visited the overseas laboratories of the main potential suppliers. Our report went to the Computer Policy Committee, chaired by Don Faulkner of MSSSO.

The Computer Policy Committee had a meeting in Room 211 of the Chancelry to discuss the acquisition, and the discussion went on all afternoon, not only about which computer to purchase but also the type of end user. There was a strong lobby which believed that a machine which also satisfied the needs of computer scientists should be considered.

A little after 5pm, Ian Ross came into the room and said, "Well, is ANU going to buy a supercomputer or not? If we are, I want you Robin, and you Don to come up to my office with a recommendation within ten minutes." When we got up to Ian's office on the fourth floor, we were taken to the conference room where we found that there were already six or more senior academics waiting. After preliminary discussion and presentation of what the Policy Committee were recommending, Ian said, "We're all going over to University House for dinner." When we got there, there were a couple of tables set up in a private dining room. During the dinner, Ian circulated around the room, chatting to each of the academics in turn. Finally, he said, "It's getting late, let's go back to the Chancelry and see if we can reach a decision."

And that's how the decision was made – at about 10 o'clock at night. ANU was going to be the first Australian university to get a supercomputer and it was going to be in place for Australia's bicentenary on 26 January, 1988, a symbolic target set a few years earlier. As everyone was leaving, Ian told me he wanted to see me first thing the next morning. When I got there he said, "Well we've made the decision about the supplier, but I would like to see some progress in other areas. In particular:

1. We should have a VP100, not a VP50;
2. ANU and Fujitsu must sign a technology agreement on cooperation and technology transfer;
3. Fujitsu must fund some research fellowships to further ANU research in computer science.

I want you to contact Fujitsu and get them to agree."

The first point was a matter of size of a larger supercomputer, the second and third were to help the development of the strong computer science presence appearing on campus. I proceeded to discuss the requirements with the Fujitsu salesman, Peter MacFarlane. A very short time after that discussion, Fujitsu rang back and said, "Yes. Japan agrees."

Some time later I was chatting with Robin Stanton about the type of the new supercomputer. He still felt that the requirements had been too targeted towards the computation disciplines, to which I said "Yes Robin, but it was only money." That piece of wisdom was quoted back to me many times over the years!

PS. Robin Stanton and Michael MacRobbie had the next bite of the cherry a few years later with investment in a Thinking Machines computer, run for computer scientists. The two supercomputers were housed in the same computer room in the Huxley building and I believe there was no animosity between them.

The advent of the first Fujitsu vector processor in 1987 coincided with the establishment of a group called the ANU Supercomputer Facility (ANUSF). It was headed by Bob Gingold, a former PhD student and research fellow, who had most recently worked on CSIRO's Cyber 205 supercomputer, located across Clunies Ross St from ANU in an architecturally pleasing building (now demolished). Bob kept ducks and one of them was called "Deefa". Ah, I said, "D for Duck". No he said, "D fer Kate". He had a poster on his wall saying, "Guy Fawkes – Only person ever to enter parliament with honest intentions."

The choice of Fujitsu was on the basis that the effectiveness of its vectorising FORTRAN compiler would deliver higher performance in practice than other systems with higher theoretical peak MFlops (Million Floating Point Operations Per Second).

Signing the contract for the Fujitsu VP2200 in March 1991. I hope that the arm in a sling is not the result of arm-twisting negotiation by ANU! Back Row: Robin Stanton (Head, DCS & Director of CISR), Robin Erskine (Director, Computer Services), Peter Kelo (Fujitsu), Bob Gingold (Head, ANU Supercomputer Facility), Stuart Manley (Fujitsu Australia), Michael McRobbie (Executive Director, CISR), Val Mickan (CEO, SGI Australia) Front: Max Neutze (DVC), Neville Roach (CEO, Fujitsu Australia) ***Photo: ANU Public Relations***

Lindsay Botten's *History of NCI*, reproduced in Appendix H, outlines the evolution from the era of ANU Supercomputing Facility (ANUSF), to the Australian Partnership for Advanced Computing led by John O'Callaghan, to the National Computational Infrastructure (NCI). His history includes a table of the principal scientific supercomputers installed over the years at ANU, but it doesn't include the machines associated with the Parallel Computing Research Facility (PCRF) established within the Centre for Information Science Research (CISR): Sequent Symmetry, ICOT PSI-3, Connection Machine CM2, Fujitsu AP1000, and Connection Machine CM5.

From the establishment of ANUSF in 1987, ANU operated a policy of allocating 10% of its supercomputing resources to users from other institutions and of providing supercomputing education to other institutions and to industry. In 1994, ANU, along with UNSW and the University of Adelaide, successfully proposed the funding of an Australian Co-operative Supercomputer Facility (ACSCF). This led to the installation of SGI Power Challenge systems (6 GFlops peak) at each of the three institutions. In 1997, UQ joined ACSCF and acquired a similar SGI system. That year ANU passed on its 32-node Connection Machine CM5 to the University of Adelaide, allowing the construction there of a 128-node CM5. ACSCF was superseded by Australian Partnership for Advanced Computing (APAC) in 2001, and ANU donated its SGI Power Challenge to UNSW, with provision for access from James Cook University.

ANU's track record in providing supercomputer access to external users and participating in ACSCF gave considerable weight to its bids to host APAC and later NCI central facilities.

John O'Callaghan reflects on APAC

APAC involved ANU and six organisations that were formed to provide high-performance computing services in each State. The decision to locate the 'peak' facility at ANU was the result of ANU's leadership and significant investments in HPC systems and services, together with the significant DCS research base that had been demonstrated by ACSys.

The task of starting APAC was made easier by being able to use some of the Federal Government funds for the HPC facility at ANU for support staff at the partners, and an education program. The outstanding capabilities of the ANU HPC group, supported by DCS, was instrumental to the success of the APAC National Facility and the rapid uptake of users around Australia.

As noted in the archived version of the APAC website[1] the APAC National Facility began operation at ANU in April 2001 with the installation of an HP AlphaServer SC computer system. See the table on Page 300.

Interestingly, APAC had first attempted to obtain a leading edge system from Sun Microsystems. In advance of the offered system, Sun delivered two less powerful servers as a stop-gap. Attempts to benchmark the offered system were not successful due to problems with the error-correction hardware in the memory. Identical runs gave different answers! Oddly, the same problems were claimed not to be experienced with other machines of the same type and, the hypothesis was presented that the problems may have related to differences between the southern and northern hemispheres! Sun Microsystems CEO Scott McNealy arrived in his private jet and met with Robin Stanton and John O'Callaghan at Canberra Airport in an attempt to resolve the problems. In the end, however, the system was determined to have not passed acceptance by the due date and the contract was terminated.

Contemporary account of the Sun-APAC contract failure. *ComputerWorld*, 09 March, 2001

Sun won the original $5 million, three-year APAC contract last August. It called for Sun to commission a 200-Gflops system in September, comprising a cluster of four E10000 servers. ... The deal called for the system to have been progressively upgraded by mid-2002 to a 1-TFLOPS system based on Sun's new UltraSPARC III processor technology. ... The deal was scrapped by early November, however, and the four Sun E10000 servers that had been installed were returned.

There are various contemporaneous reports on the web of other Sun systems suffering similar memory-related problems.

John O'Callaghan was the Executive Director of APAC throughout its life from 2001 – 2007. He says that APAC's greatest strength was that it was a partnership. There was a strong sense of an Australian supercomputing community. State governments committed to the partnership and provided funding. This meant that users in every state had access to a local machine and to the education in how to use supercomputing, provided by APAC. Users were also able to make use of software developed by staff at APAC, such as the job queueing and resource sharing system. Users with very large scale computational requirements were then able to easily transfer their applications to the peak facility at ANU.

John says that another strength was the untied way in which funding was allocated from the government. It was a requirement that APAC purchase and operate a peak computing facility clearly more powerful than could be afforded by an individual institution, but there was flexibility in how the budget could be divided between hardware and support staff. As a result, a significant number of support staff were employed, potential users were able to be appropriately assisted, and the peak machine was heavily used from day one.

External supercomputing facilities enviously compared APAC's staff-members-per-teraFlops with their own, and formal reviews of APAC by international panels delivered glowing reports.

It was not inevitable that APAC came to pass. An alternative model discussed was that Australian supercomputing needs could be satisfied by purchasing time on overseas facilities. At the opposite, and regrettably fanciful, extreme, some explored the idea of a supercomputing life cycle in which every stage of the process from the mining of silicon to the fabrication of chips and assembly of

[1] https://web.archive.org/web/20090120211513/http://www.apac.edu.au/

computer arrays would be done locally.

Like ANUSF before it, and NCI after it, resources on the APAC facilities were allocated on merit by a Merit Allocation Committee (MAC). The MAC was charged with assessing the likely scientific value of a request for time and also the suitability of the supercomputer for running the proposed algorithm. (Of course, many small requests were waved through to enable code development and prototyping.) John thinks that merit allocation (also practised in the USA) was a much superior mechanism to the one used in the UK, where bids for research grants could request allocation of vouchers to buy time on British supercomputers.

I hadn't realised it, but John O'Callaghan is an ANU PhD graduate. His PhD, under Steve Kaneff in the department of Engineering Physics, was on the topic of hand-writing recognition, and his code ran on the IBM 360/50. He says that the method he developed prior to graduation in 1969 is essentially in the same class as the method used on modern smart phones and tablets. After graduation he joined CSIRO's Division of Computing Research and stayed with CSIRO until his role at APAC.

Roger Edberg, User Support Manager at ANU's National Computational Infrastructure (NCI), has experience of ANU's supercomputers since the early days. He says that, early on, the supercomputing user population comprised relatively few and relatively sophisticated users, concentrated in a small number of disciplines such as computational chemistry and astronomy. Now the demand for time on NCI facilities is spread across a much broader range of disciplines and many of the users are less sophisticated, mostly running packages developed externally. The demand is very heavy, achieving a utilisation of 94%, which compares with the 70% typical of other high performance computer centres. The number of researchers using the machine is of the order of 5000!

Reading the ANU Supercomputer Facility reports from 1987 – 2001, increasing diversity of the disciplines using the facility is obvious from the breakdown of time use, and from the increasing range of software packages supported.

Roger says that 30-40% of code run on the NCI facilities is still written in FORTRAN. He says that a major reason for this is that PhD supervisors used FORTRAN and they influence their students to use and/or modify the same codes.

An ongoing problem for supercomputing is getting the data to where the computational resources are. A current example is the Autism CRC (`https://www.autismcrc.com.au/`) which produces 10 teraBytes of genomic sequence data per day and wishes to analyse it on the NCI infrastructure.

8.0.1 Current Facilities at NCI

In October 2020, Allan Williams, Associate Director, Services and Technology, was kind enough to take me on a tour of NCI facilities. To a lapsed designer of small-scale computer rooms like me, the scale of the facilities is mind-boggling.

- The area of the computer room is approximately 930m^2. That's approximately the size of two residential blocks in Gungahlin (average size 491 m^2).
- The facility uses an average of 2 megaWatts of power, or 17% of ANU's electricity. You may think it's a bit ironic that such energy intensive facilities are used to model the impacts of climate change, but ACT's electricity supply has been net zero emissions since 2019.
- The new machine Gadi (which means "to search for" in the local Ngunnawal language) consists of 3024 compute nodes, each with two 24-core CPUs and 384GB of RAM spread across six memory channels. Some of the nodes have an additional 1.5TB of Optane memory. There are 640 V100 NVIDIA GPUs.
- The peak performance from the 145,152 CPU cores is about 9 petaFlops on the HPL (High Performance Linpack) benchmark. Total performance above 14 petaFlops is possible if the GPUs as well as the CPUs can be utilised.
- Interestingly, no-one ever uses the entire machine at once. Such a job would be exceedingly difficult to schedule because of the 5000 other users. The largest job run so far (up to late 2020) used 40,000 cores.

- Of course there's massive amounts of storage, and tape libraries.
- Gadi racks are very heavy. The false floor under it has been strengthened to support a load of 2 tonnes per square metre.
- Sections of the Raijin machine which Gadi replaced are still in use, providing a cloud and for additional computing.
- NCI has direct connections to CSIRO and the University of New South Wales. Since the Bureau of Meteorology was hacked, the direct connection to BOM has been cut. (To protect NCI?)
- NCI has a 100 gigaBit/sec link to AARnet.

October 2020: Me standing at the entrance to a corridor of racks containing NCI Gadi nodes. Note my COVID hairstyle. The Gadi artwork is by local indiginous artist Lynnice Letty Church. ***Photo: Allan Williams.***

Largely because of emissions, the use of energy is an increasing concern for both computer centre managers and for algorithm developers. The efficiency of algorithms is increasingly being measured "per Watt". GPUs are generally much more energy efficient than CPUs. At NCI, the major computer systems are cooled via chilled water, and careful design has allowed the use of evaporative rather then refrigerative water chillers.

8.0.2 ANU-Fujitsu collaboration

In the late 1980s, Japanese supercomputer vendors experienced enormous difficulty in selling their equipment in western markets, particularly in the United States. For example, NEC offered to sell an SX2 (or was it SX4) configuration to MIT for $US1, but their offer was refused, despite the machine being claimed to be the fastest in the world. (The SX2 is claimed to be the first machine to achieve a performance of 1 gigaFlops.)

Email conversation with Rika Kobayashi about computational chemistry

Dave: Back when I was running the Fujitsu AP1000, the feeling of chemists like Denis Evans seemed to be that, although parallel machines like the AP1000 potentially had a lot more power than the vector processors, it was too difficult to convert chemistry codes to run efficiently on such machines. With the VP, you could rely on the vectorising compiler to give you a big speed-up and you didn't have to re-cast the basic algorithm.

Rika: This is DEFINITELY an ongoing conversation such that the conference I am preparing for is purporting to address exactly that as you will see from their "mission statement."[a] They call it sustainable software in recognition that they need to have a programming framework (people as well as tools) to adapt to changes in architecture because "it was too difficult to convert chemistry codes to run efficiently."

Dave: But now the scale of the parallel machines like NCI's has left vector processing way behind. I'd really be interested to hear how things have changed with respect to chemistry software.

Rika: What goes round come round. From the early 2000s MPP machines did take over as specialised vector machines just weren't competitive costwise. BUT with the latest chips – think accelerators (Xeon Phis, GPUs, ARM) it turns out that many of the vector processing tricks you needed in the 90s are making a comeback.

Dave: Are there now tools for converting old sequential codes into parallel form?

Rika: There has been work done using autoparallelisation through compilers which we used to cover when I was involved with NCI training. I haven't been involved for about a decade now but from my experience with that I haven't felt compelled to touch it. On the other hand, compilers ARE getting cleverer. People are modularising code more to take advantage of common optimised libraries.

Dave: Does porting involve re-casting the algorithm into a more parallel form?

Rika: Porting involves just getting the code to work on new architectures. BUT part of the Area 3 project is performance improvements and the ideal is to recast algorithms, but time constraints make that a wishlist item. I have been trying to improve the Gaussian coupled cluster algorithm since 2003 (was going to get back to it this year until COVID hit and I got distracted with remote learning initiatives) but have never had the continued time required to work on it.

Also, over the last 20 years all chemistry programs have incrementally been parallelised, but this had to be done by hand because for quantum chemistry, it was the memory parallelisation (data) that was the obstacle not the vector loops.

Dave: Are there new chemistry packages written specifically for parallel machines?

Rika: Yes, definitely. I was amongst the first to do this with the NWChem package, which is acknowledged to be one of the first quantum chemistry packages for MPP. Molecular modelling programs do MUCH better. Most new programs are aiming for ease of programming through Python, which I think is a step backwards.

Dave: Are some modern chemistry codes still written in FORTRAN?

Rika: Programs still use FORTRAN for legacy reasons but new stuff mostly doesn't. This isn't necessarily the case in Astronomy because I believe Physics does programming at undergraduate level better than chemists.

Dave: The other day I noticed a lab of students running Gaussian. Is that still hugely important?

Rika: Yes, it is to Fujitsu. Possibly not so much for the rest of the world. In the old days Gaussian sold computers and so every hardware vendor had a comp chem on staff mainly to deal with Gaussian and other computational chemistry packages. There are very few hardware vendors who do this – the only ones I know now are Fujitsu, NVidia and Intel (except not for Gaussian but VASP – I think VASP is the new program that sells machines).

Fujitsu Japan still pay Area 3 to get Gaussian working and there was a heavy push to have Gaussian available for their A64FX platform which is the same as Fugaku (415 petaFlops) in time for the announcement of it reaching #1 in June's Top 500 list so it definitely still matters to them.

I was asked to teach CHEM3208 this year and we needed a computer lab. I based it on Gaussian and got to deliver a Virtual Lab through CS&IT's VDI which is why we ended up in your building. That was my Computer Labs of the Future pet project.

[a] `https://www.edison.re.kr/web/nhiss/letter`

A close relationship between ANU and Fujitsu developed slowly over a long period of time. ANU bought a FACOM 160F in 1980 to handle administrative data processing, and upgraded several times over the next decade. When ANU bought a VP100 supercomputer in 1987, this was a very important moment for Fujitsu. At the ceremony marking the installation of the VP, the very senior person from Fujitsu made a significant verbal slip:

– "The acquisition of the VP by Australian National University is very significant for U.S. ... so sorry, 'very significant for **us!**"'

In the following years, the relationship became even stronger. ANU improved the attractiveness of the VP supercomputers to other potential Fujitsu customers by porting important 'codes' (software) to the platform, particularly in the area of chemistry. ANU earned substantial amounts of money through this porting activity.

ANU eventually upgraded the VP to a VPP300.

The ANU-Fujitsu relationship expanded over time to encompass four areas:

Area 1 – Research fellowships

Area 2 – Parallel computing – i.e. the CAP project. See Section 8.1.2.

Area 3 – Chemistry and molecular modelling. Ross Nobes headed up this project from 1989, porting chemistry codes to run on the Fujitsu VP machines. The project is still going in 2020.

Area 4 – Mathematical Subroutine Library – A paper by Richard Brent et al (`https://dl.acm.org/doi/10.1145/277830.277837`) describes ANU work on producing a library of mathematical subroutines for the vector-parallel Fujitsu VPP300.

Rika Kobayashi and Roger Amos continue work on porting chemistry packages to run on highly parallel machines like Gadi. Rika works on porting Gaussian, Roger works on porting VASP, and other colleagues work on GAMESS (Ivan Rostov) and AMBER (Vladislav Vassiliev).

The panel on Page 205 records an email conversation with Rika about how computational chemistry has evolved to accommodate to changes in computer architecture.

8.0.3 What Happens to Supercomputers When They Die?

Bob Gingold, former Head of ANUSF, tells me that Fujitsu Australia (FAL) took away the VP100 and VP2200 as part of their provision of successor machines. When the VPP300 came to its end of life, there was no Fujitsu successor, and no arrangement for disposal of the old machine. An "exuberant" young Academic Consultant in ANUSF came to Bob and asked, "Why don't we sell it on e-Bay?" Bob continues the story:

> " I explained all the reasons why that would not be a good idea. But unbeknown to me, the CSC/ITS administrator at the time overheard our conversation and, also unknown to me, she wrote to the VC asking for permission to sell it on eBay. Her note came back stamped 'Approved'. I couldn't hold out any longer. The ad was put up with no reserve – we just wanted to avoid having to pay for removing/junking it.
>
> It led to some publicity.[2]
>
> Within a day or two or even less (can't recall) I got a phone call from Fujitsu Japan at around 11pm, not sure who it was but probably someone I didn't know well or else I'd remember. He expressed concern, making it clear FJ weren't happy but formally only telling me not to sell it to North Korea, Iraq and Iran (as per contract – I think we could have sold it to Cuba). Next day FAL were blowing a fuse after a blast from Fujitsu Japan and we withdrew the ad. FAL paid to take it away. I don't know what they did with it. From an article in the Age it seems we got one bid (for $3100) before we took it down.
>
> We advertised it without an operating system ... we didn't own the operating system to sell it. Hence it was really just a museum piece. You can't buy the operating system."

[2] See for example `https://www.theregister.com/2003/03/12/buy_this_supercomputer_on_ebay/`

According to the Age the advertisement said:

> FOR SALE: Fujitsu VPP300 supercomputer, 28.6 gigaFlops, 14GB of memory, one owner, size of four big fridges. Make us an offer.

8.0.4 Kspace at the National Museum of Australia (NMA)

According to an NMA History[3], Kspace, created jointly by NMA and ANUSF, "was a new and exciting digital experience for kids. Visitors used the Kspace computers to help build futuristic rockets and houses. The personalised designs were then presented as part of a space-age Australian city in a three-dimensional theatre."

Drew Whitehouse, ANU's visualisation expert, and his team, "created the virtual city based on children's drawings of an imagined city of the future."

8.0.5 *One Australia* Sinking

One Australia, an Australian boat competing in the 1995 Americas Cup was created by John Bertrand and the Fluid Thinking company. Fluid Thinking used supercomputer modeling to dramatically speed up the traditional design process. *One Australia* was sponsored by Sun Microsystems and arrangements were made to use ANU supercomputers for some of the design work.

Unfortunately, according to a media report,[4] "... in a match race between One Australia and Team New Zealand, the Australian boat split and sank within two minutes off San Diego, CA."

Bob Gingold says,

> "The US designer using the VP2200 was responsible for the hydrodynamics not the structure of the yacht, so its breakup wasn't down to him. But I recall he moaned a lot about having to use the VP2200 instead of his usual Cray. Especially when it gave him different answers for the same calculation. One of our people (maybe Margaret Kahn, Murray Dow or Dave Singleton) analysed his code and found potential numerical instabilities which might cast some doubt on the reliability of the code in some circumstances (on both platforms)."

8.1 The Centre for Information Science Research (CISR)

The Centre for Information Science Research (CISR)[5] was established in 1987. Key players were Michael McRobbie, Robin Stanton, Richard Brent, and Bob Gingold. Ian Ross, ANU's DVC, was a supporter – many critical meetings took place after 5pm in his office in the Chancelry over a dram or two of whisky. I've heard that Michael McRobbie once took to such a meeting, a bottle of whisky on which he had marked various levels labelled with key stages in the meeting, culminating with "Decision taken" near the bottom.

Michael McRobbie. A double-page feature in the Australian claimed that "without McRobbie there would be no computing at ANU." That somewhat understates the contribution of one or two others, but there's no doubt that Michael had a huge influence. He played important roles in the ANU-Fujitsu relationship, and was the force behind the acquisition of the Connection Machines, the Sequent Symmetry, and the iCOT machine. As noted in Section 6.2, Michael completed a PhD in philosophy at ANU, and returned to found the Automated Reasoning Project with Bob Meyer.

Michael (known from his undergraduate days as Zorba) was very security conscious. He had a large office in H Block with access protected by two levels of reception. He had a Class B safe[6] and once had his office swept for bugs.

[3] `https://www.nma.gov.au/learn/kspace/about-kspace/kspace-history`

[4] `https://www.sailingscuttlebutt.com/2019/03/05/americas-cup-sinking-one-australia/`

[5] Robin Stanton said that it should have been called the Federation of Information Science Research, since there was no Centre, but the acronym wasn't suitable.

[6] A Class B high-security safe is equipped with a 13mm solid steel door and 6.5mm solid steel walls on all five sides.

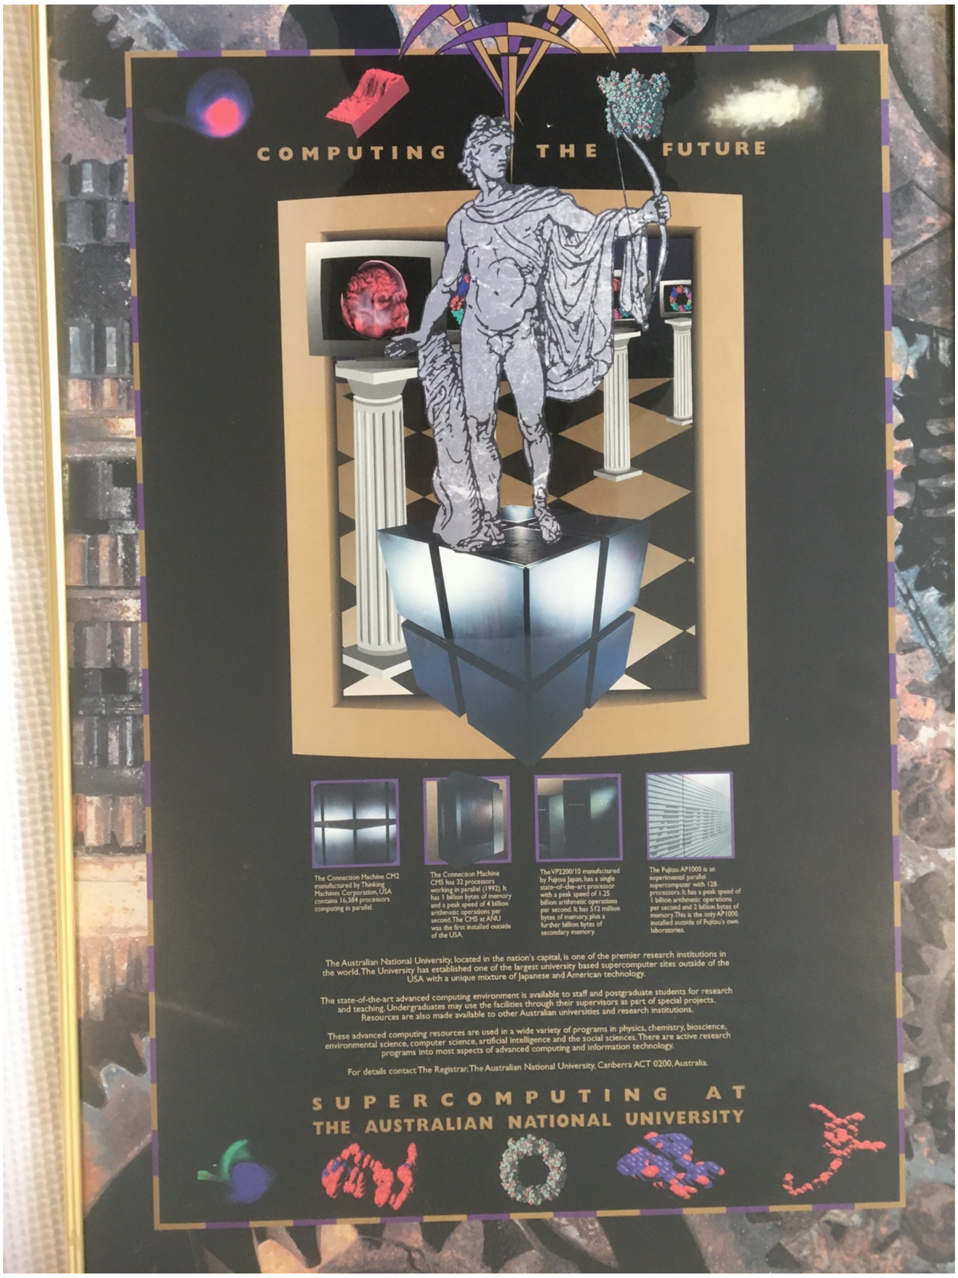

The Connection Machine CM2 manufactured by Thinking Machines Corporation, USA contains 16,384 processors computing in parallel.

The Connection Machine CM5 has 32 processors working in parallel (1992). It has 1 billion bytes of memory and a peak speed of 4 billion arithmetic operations per second. The CM5 at ANU was the first installed outside of the USA.

The VP2200/10 manufactured by Fujitsu Japan, has a single state-of-the-art processor with a peak speed of 1.25 billion arithmetic operations per second. It has 512 million bytes of memory, plus a further billion bytes of secondary memory.

The Fujitsu AP1000 is an experimental parallel supercomputer with 128 processors. It has a peak speed of 1 billion arithmetic operations per second and 2 billion bytes of memory. This is the only AP1000 installed outside of Fujitsu's own laboratories.

The Australian National University, located in the nation's capital, is one of the premier research institutions in the world. The University has established one of the largest university based supercomputer sites outside of the USA with a unique mixture of Japanese and American technology.

The state-of-the-art advanced computing environment is available to staff and postgraduate students for research and teaching. Undergraduates may use the facilities through their supervisors as part of special projects. Resources are also made available to other Australian universities and research institutions.

These advanced computing resources are used in a wide variety of programs in physics, chemistry, bioscience, environmental science, computer science, artificial intelligence and the social sciences. There are active research programs into most aspects of advanced computing and information technology.

For details contact The Registrar, The Australian National University, Canberra ACT 0200, Australia.

A rather odd poster, presumably aiming to attract new staff and students to ANU, based on ANU's 1993 quartet of supercomputers. *Poster: ANU*

CISR established what it called the PCRF (Parallel Computing Research Facility) starting with a Sequent Symmetry shared-memory machine based on Intel 80386 chips. Soon afterward a 16,384 node Connection Machine CM2 was obtained. It was so stunning in appearance, and considered such a coup that we initially set it up in the National Science and Technology Centre, in the central drum. Robert Whaley from Thinking Machines was based at ANU for a year and he trained John Barlow, Lindsay Hood, Terry Bossomaier, and me as Application Engineers. As well as writing programs in C*, operating and fault diagnosing the CM2, we learned how to diagnose, replace and heal faulty disks in the massive 10GB highly parallel DataVault.

CISR's main computational facilities were located in the Leonard Huxley Building and in DCS, but CISR's Director Michael McRobbie and a small number of staff were located in I block. I block also housed some people from the Automated Reasoning Project (see section 6.2) and some technical people: Gustav Meglicki, John Barlow, and the Thinking Machines applications engineer Robert Whaley. I believe that the Sequent Symmetry and the ICOT PSI-3 machines were also located there.

John Barlow (front) and a Thinking Machines person, moving a component of the Connection Machine CM2 into the National Science and Technology Centre. *Australian National University Archives:, ANUA 226-137-4.* `https://openresearch-repository.anu.edu.au/handle/1885/216165`

The CM2 was a beautiful machine – See the photograph on Page 210. The front face of a full 65,536 node version of the CM2 consisted of 4 squares, each with 16,384 red flashing lights on a smokey black background. Internally the nodes were connected as a 12-dimensional hypercube.

The CM2 was capable of achieving 27 gigaFlops (on a highly optimised calculation of the dot product of two exceedingly long vectors), however, if programmed naively to multiply matrices, achieved only about 50 megaFlops. Bridging the gap between these performance extremes required extensive knowledge of the hardware. The real experts at Thinking Machines who wrote the fundamental libraries were said to "program the wires" of the hypercube.

Excitingly, ANU was able to claim, using Bob Gingold's tongue-in-cheek words, that it had "the two fastest computers in all of the countries not starting with U or J."

A senior executive from Thinking Machines presented a pitch at ANU to persuade people to start using the CM2. His argument was based on examples of exciting applications to which CM2s had been put. One of them was a parallel text searching system for Dow Jones. This presentation ultimately led to my subsequent research career in Information Retrieval.

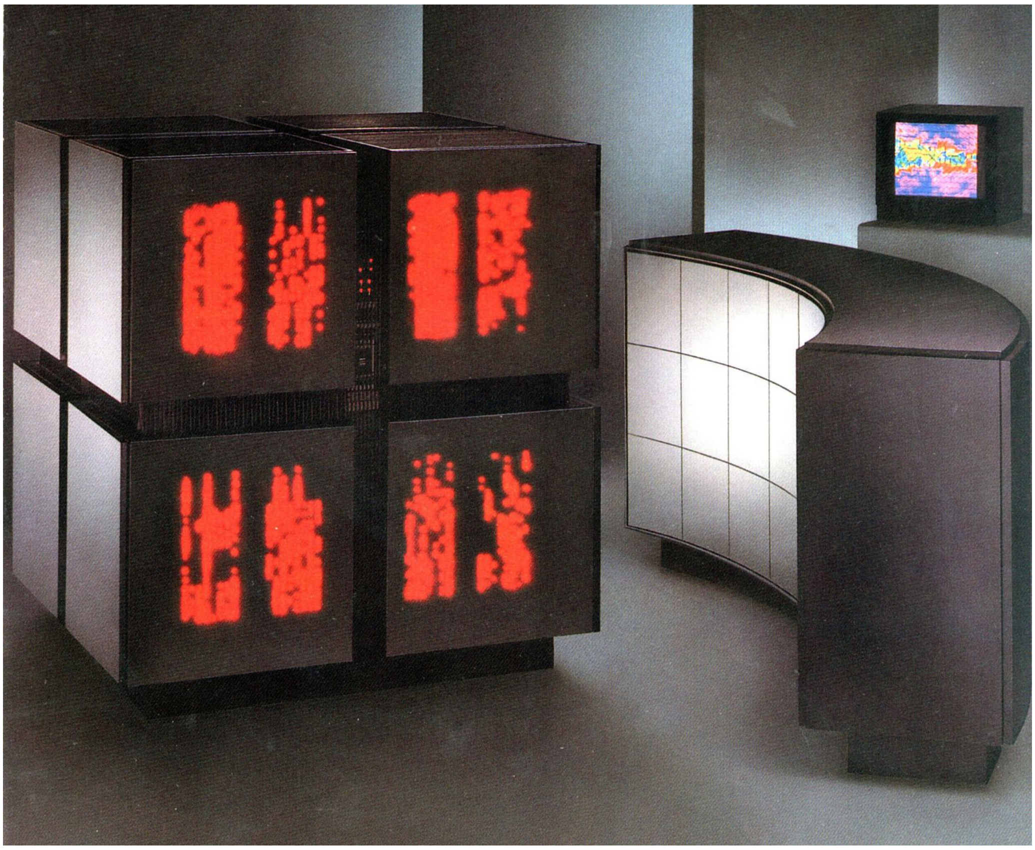

Connection Machine CM2, unfortunately not ANU's. Note the curved DataVault and the high performance graphics display. *Photo: Thinking Machines Marketing*

An extract from a Michael McRobbie report on ANU/ICOT cooperation

In 1991 John Slaney of ANU and Ewing Lusk of Argonne National Laboratory (ANL) in the United States devised a new automated theorem proving system called SCOTT for Semantically Constrained Otter – Otter being a high performance theorem prover written by Bill McCune at Argonne National Laboratory which is generally regarded as the fastest and best general purpose theorem prover presently in existence. On a wide variety of problems SCOTT is about twice as fast as Otter.

8.1.1 ANU's Hidden Supercomputer

From the late 1980s, ANU's fleet of powerful Unix workstations[7] were connected together by Ethernet. If a person with an "embarrassingly parallel" program had login access to all of them, this collection of workstations could function as another ANU supercomputer. At the 1991 Australian Supercomputer Conference, ANU's Russell Standish presented a paper entitled, *The Burnett Problem: Harnessing ANU's Hidden Supercomputer*.

Within DCS, both Richard Brent and Brendan McKay were avid users of the computing resources represented by Sun workstations in the department. They used different versions of monitoring programs to determine whether a workstation was active or inactive. Once deemed inactive, the background computation would spring into life, suspending itself when the owner of the workstation returned to its keyboard. Brendan thinks that the code he ran was probably to prove theorems in combinatorics and graph theory.[8]

As DCS "Head Programmer"[9] in about 1989 I was called to respond to a complaint that when people came into work early in the morning, response from the Sun workstations was incredibly slow. Counter-intuitively, as more people logged in and the workload increased, the workstations

[7] 400 of them, according to one CSC Annual Report

[8] These days he uses NCI and also has a 32-core box he built himself.

[9] Our friend Huw Price thought it funny that my partner Kathy was a neuropsychologist and I was a head programmer.

would gradually gather speed! I eventually tracked it down to a background job of Richard's. After a long period of execution, his background jobs would increase their memory demand beyond the size of the workstation's RAM, causing paging. Because the workstations at that time were diskless, the paging happened over the Ethernet and onto the server Vega, clogging the ethernet and Vega. Richard confesses:

> I've run a lot of computations like that, so not sure exactly which one it was, but most likely it was trying to factor large numbers by the elliptic curve method, or maybe trying to prove the non-existence of odd perfect numbers.
>
> If the memory eventually grew larger than RAM, then I was probably using Magma[10] which (at least in those early days) suffered from memory leaks. Anything that I wrote myself (usually in C) could not possibly have a memory leak. ☺

8.1.2 ANU-Fujitsu CAP project

Australian National University
Official Opening
FUJITSU
AP - 1000 Experimental Parallel Supercomputer
ANU

13 May 1991: Deane Terrell (ANU VC), Kim Beasley (Federal Minister for Education and Employment), Michael McRobbie (Director, CISR), Robin Stanton (Head, DCS), Shigeru Sato (President, Fujitsu Laboratories), Richard Brent (Head, Computer Sciences Laboratory, RSISE), Neville Roach (Managing Director, Fujitsu Australia). *Photo: ANU Archives.* **`https://openresearch-repository.anu.edu.au/handle/1885/216202`**

In about 1990 a delegation from ANU visited Fujitsu Laboratories in Kawasaki, Japan. Robin Stanton showed an interest in Fujitsu's Cellular Array Processor (CAP) project and developed a relationship with Mitsuo Ishii of Fujitsu Labs. Of particular interest was the CAP-II, a system with

[10] `http://magma.maths.usyd.edu.au/magma/`

up to 1024 SPARC nodes connected by a 25Mbyte/sec 2-D torus network (T-net) and a 50Mbyte/sec broadcast network (B-net). The machine had a remarkably high communication to computation speed ratio. Each node ran `CellOS` and programs were broadcast to the nodes from a front-end running Unix.

In 1990, the first ANU-Fujitsu CAP Workshop was held in Kawasaki. It was attended by many from DCS.

Fujitsu agreed to lend ANU a 128-node AP1000 (CAP-II) with 16 option boards and disks, and to provide funding to support research on the machine. This was very significant, since the list price of our configuration was about $A3 million ($5.83 million in 2019). As a machine with a peak speed above one gigaFlops, and 2 GB of RAM – the largest RAM configuration in the southern hemisphere – the AP1000 was subject to US-imposed export controls designed to prevent advanced computers falling into the wrong hands. We were required to deny access to citizens on a list of banned countries.

An advance delegation from Fujitsu brought CAP simulator software (CASIM) and gave lectures on the hardware and how to program it. ANU bought a Sun machine to front-end the AP1000, and, in the meantime it was used to run CASIM.

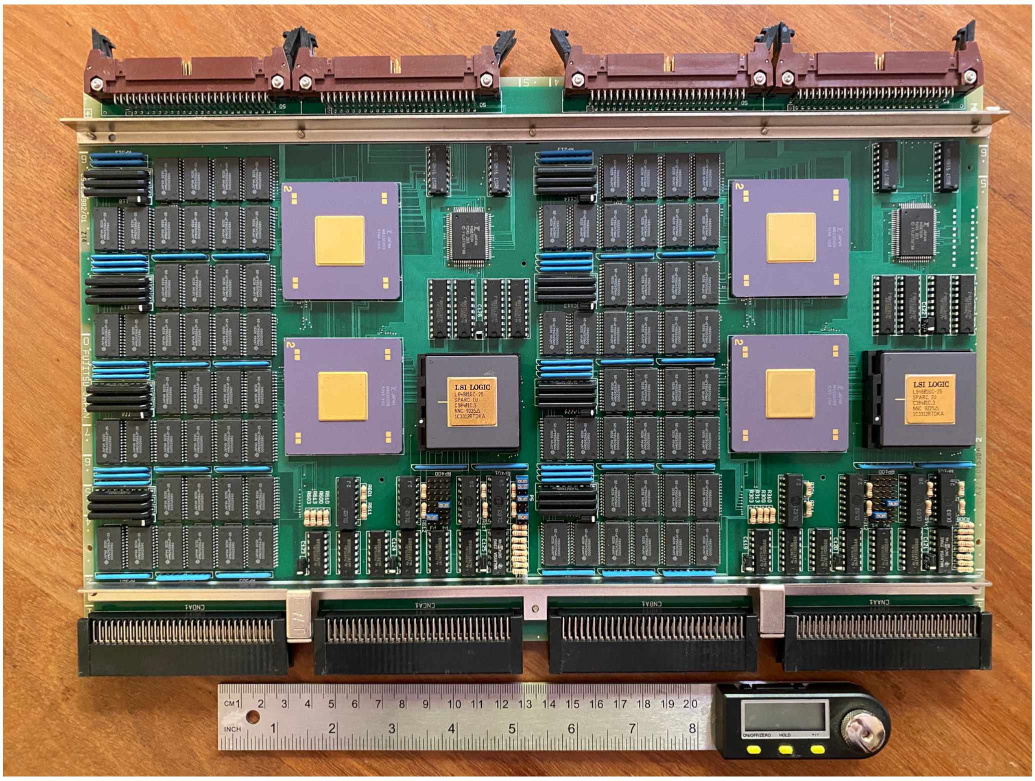

A CPU board (two SPARC nodes) from ANU's AP1000, with a ruler to indicate scale. Note the high quality of assembly. *Photo: David Hawking*

Hardware training for the Fujisu AP1000

In March, 1991, Paul Mackerras and I, with two technicians from Fujitsu Canberra, flew to Japan for training on the hardware of the AP1000. We enjoyed two days of lectures which mostly consisted of a Fujitsu Labs person orally translating the content of the manuals, which were only available in Japanese. Paul had learned katakana before our arrival. Katakana is a phonetic alphabet which is used to represent foreign words imported into Japanese. For example, the English word, "taxi" is rendered as the katakana for the sounds, "ta - ku - shi". Paul discovered that a high proportion of words in the manuals were actually imported words and was able to understand them by sounding them out. Thus, "Bus Arbitration" presented no problem, but simple words like "red" were unrecognisable, because they were written in non-phonetic kanji.

1991: Marketing photo for Fujitsu AP1000, featuring me and Paul Mackerras. *Photo: Fujitsu Laboratories*

1991: AP1000 Hardware Training, Kawasaki, Japan. Back row: Shiraishi, Ishii, Ishihata, Sato, Mackerras. Front row: Alford, Judge (Fujitsu Australia) *Photo: David Hawking*

Fujitsu put Paul and me up in the very expensive, very tall, hopefully earthquake-proof, Shinjuku Hilton, in which the rooms were small but very elegant. Across from my room, I could see people on intravenous drips in a multi-storey building opposite – presumably it was a hospital. At about 09.00 on the first morning, Shimizu-san came to collect us to take us to Kawasaki on the train. We walked through a pedestrian subway toward the world's busiest railway station.[11] Coming toward us was a mighty river of people rushing toward the twin towers of the Tokyo Metropolitan Government. The effect was so much like perching on a rock in the middle of the river just above Niagara Falls, that I found it difficult to maintain balance.

Walking around the Kawasaki campus, Paul asked Sato-san whether he drove to work. "Of course not. There would be nowhere to park." Coming across a small group of cars, Paul asked,

[11] Shinjuku. Said at the time to handle 2 million passengers per day.

– "So, if you are senior enough, you can get a parking spot here?"

– "Yes, and a car, and a driver."

We were also taken to Numazu, a Fujitsu factory close to Mt Fuji and set among tea plantations. We were professionally photographed next to an AP1000 so that Fujitsu could use the photos in sales brochures. In my case an air brush was used extensively to reduce the volume of my curly, and very un-Japanese-looking hair. In Numazu, we were given training on how to diagnose hardware faults, and asked to diagnose a number of faults. Paul very quickly realised that the "faults" had been deliberately created by soldering jumper wires onto the circuit boards, and proposed a simpler diagnostic procedure: "Find the board with the extra blue coloured jumper wire, and replace it with a new board."

I claim to be the only person on the planet to have ever received hardware training on both a Fujitsu AP1000 and a Connection Machine CM2.

During our training, Fujitsu looked after us very well with hospitality during the day, and at least one extraordinary banquet of about 17 courses. On a free evening, Paul and I took a train from Kawasaki to Yokohama and wandered along the coast and among flower gardens, before looking for a sushi restaurant. When we found one we sat at a table, but were immediately summoned up to the counter where everyone else was sitting. They moved coats and bags and shuffled along so we could sit up and watch the chef at work. We were unable to communicate with anyone but they made us feel very welcome. The chef wrote on a piece of paper:

```
A -- 1500 yen
B -- 2200 yen
C -- 3000 yen
```

We chose B, and ate very well.

Heading back to Canberra, Sato-san gave us clear directions and helped us buy tickets on the newly launched Narita Express. We stopped in a restaurant in Tokyo Central Station and then had to wait for ages to pay. Then we discovered that the Narita Express left from a very distant platform.

Paul: I do recall being in Tokyo with you and intending to catch the Narita Express to the airport to come home, going to the wrong platform, getting to the right platform only to see the train pull out as we were at the top of the last set of stairs, taking a slower train, eventually getting to the airport, running up lots of escalators, and getting to the check-in only just in time...
Dave: Yes, the person supervising Qantas check-in, said, "Qantas!! Come with me!". People only a minute behind us were told, "This flight is closed." When we got to Sydney, we ran for the Canberra flight, and I made it, but you didn't.

Since that experience I've allowed lots more time for connections!

1991: Crawford Building, ANU. The AP1000 has just been installed. Back row: Stanton, Hawking, Mackerras, Shimizu, ?, ?, ?. Front row: Ishihata, Shiraishi, Horie *Photo: Unknown ANU*

Installation of the AP1000 in Crawford: April 1991

The AP1000 was installed on the top floor of the Crawford Building in April, 1991, after we installed high current power outlets, confirmed adequacy of air-conditioning, confirmed adequacy of floor loadings, and obtained and installed CAFE workstation (CAP Front-End), etc.

Why Brian Corrie says he came to ANU. (We Canberrans suspect that his real reason was to introduce Canberra to the sport of Ultimate, played with a flying disc, and to live for a while in a place where you could play it all year round.)

What brings a Canadian to ANU?

As an early career researcher, what better opportunity both personally and academically, to move to the other side of the world (Canberra) and work on a project with leading researchers and state of the art hardware. In fact, ANU and the Department of Computer Science drew me to Australia not once, but twice. The first time was in 1991 for a two year stint to work on the ANU-Fujitsu CAP project. Writing parallel visualization algorithms that took advantage of a machine with 128 processors (yes 128 processors in a single computer!), and access to a machine with 1024 processors (at Fujitsu Japan), was pretty well unheard of at the time except at large national research labs in places like the US. What an opportunity, and plus, I was in Australia for goodness sake! The second time was in 1995. After two short years back in Canada, clearly I couldn't get enough of Australia. I was drawn back to ANU to work within the ACSys CRC program, this time on state of the art virtual environment platforms like the Haptic Workbench, the ImmersaDesk, and the home grown ANU Wedge. After three years working with the ACSys CRC, I once again returned to Canada, bringing my experiences in virtual environments to work at Canada's National Research Council. My time at ANU set the direction for the rest of my career, focusing on scientific visualization, scientific collaboration, and computational science. Thanks ANU and DCS!

As may be seen in the photo on Page 214 a team came from Fujitsu Japan, with Shiraishi-san as leader, Ishihata, Horie, and Shimizu from the AP1000 team at Fujitsu Laboratories, and three technicians from the Numazu factory.

Although essentially an experimental machine, the AP1000 was engineered to an extremely high standard. When installed in the Crawford building, we wanted to leave the front doors open, so that users in the lab next door could watch the LED lights through the glass. The installers were horrified and, through Shiraishi-san, asked us to keep them shut to improve reliability. When we inquired what the loss in reliability would be, the answer was that reliability would drop from 99.999% to 99.995% – We kept the doors open. In several years of operation, I don't remember any hardware errors, apart from corrected single-bit RAM errors.

David Sitsky ported MPI and later MPI-2 to the AP1000. MPI is a generalised message passing architecture which can simplify the writing of parallel programs for a wide range of parallel and distributed architectures. Avoiding the need to use non-standard CellOS calls by having MPI for the AP series greatly assisted porting of external codes to the platform.

In fact, the AP1000 was not much used by people outside the computer science area. One exception was David Hansen, whose ANU chemistry PhD relied on large scale simulations on the AP1000. My understanding was that his supervisor Denis Evans was using David's PhD work to explore the usefulness of a highly parallel machine for chemistry work. I think the conclusion was that it was much easier to port existing codes to the VP than the AP because of the vectorising FORTRAN compiler. Peiyi Tang (now at the University of Arkansas in Little Rock) worked on parallelising compilers, but progress didn't occur quickly enough for the chemists.

Inspired by Thinking Machine's parallel text retrieval system for the CM2, I developed PADRE (the Parallel Document Retrieval Engine) using MPI on the AP1000. It loaded the entire textbase into the memory of the nodes and because of the high degree of parallelism was able to achieve rapid retrieval without an index, while supporting complex operations such as regular expression matching. E.g. Find all the palindromes in the text of the Oxford English Dictionary. Peter Bailey worked with me on this project and we had support from Harriet Michell and the Australian National Dictionary Centre. Eventually Richard Jones told us about the TREC (Text Retrieval Conference) evaluation campaigns run by NIST in the US, and we started participating in 1994. Paul Thistlewaite developed software to automatically turn NIST topic statements into PADRE queries and I also manually constructed massively complicated queries involving patterns with hundreds of alternatives.

Andrew Tridgell, on the occasion of his award of an honorary DSc, with Elanor Huntington, Dean of CECS in the robing room. ***Photo: CECS Marketing***

The Boyer-Moore-Gosper algorithm allows for considerable speed-up of string pattern matching based on a skip table. I wrote my own multi-alternate version to handle patterns including multiple alternates, but as the number of patterns increased, the speed-up decreased. I proudly presented the work at a Friday lunchtime gathering of the CAP team. Tridge said, "you could use bigrams!"

Two days later he presented me with BMG2, a version of multi-alternate BMG in which the skip table is indexed by letter bigrams rather than single letters. The code he wrote included many tricks to reduce space and increase speed, and included an option to use trigrams rather than bigrams. Trigrams turned out to be unnecessary. Even when the number of alternate strings is very large, the speed-up achieved by BMG2 hardly drops. BMG2 has been in use ever since and no bugs have been found.

Fujitsu AP+

The further loan by Fujitsu to ANU in 1995 of a Fujitsu AP+ was highly significant. Although there were only 16 nodes, the nodes were much more powerful, each had 64MB of RAM and, best of all, each had a memory management unit (MMU). This allowed the possibility of running a conventional operating system on the nodes. As described by Paul Mackerras in the panel below, a "battle" ensued between "gun coders", Paulus and Tridge, to get a version of Unix going on the AP+.

Paul Mackerras describes the battle between NETBSD and Linux for the Fujitsu AP+

That was when the AP1000+ first arrived, which was early 1995 if I recall correctly. Since it had SuperSPARC processors which included an MMU, it was possible to run a real operating system kernel on each node. Since I was using NetBSD on my home system (which was based on a board I constructed myself with a Motorola 68030 processor) and Tridge was a Linux fan, we decided that we would each port our favourite OS to the AP1000+ nodes. Tridge got Linux running in a very short time, from memory only a week or two, whereas I never finished a NetBSD port – my excuse is that I had two small kids at the time, about 5 and 1.5 years old. I recall that Tridge initially booted with init=/usr/bin/emacs on the kernel command line, i.e. he was using emacs as the main system initialization and command line interface for the system. ☺ I don't recall the specific reason why he used emacs rather than bash.

The port was achieved in an amazingly short time.

Steve Blackburn recalls Tridge debugging AP/Linux

Tridge had been working on it nearly 24/7 for a number of weeks! My office was right opposite the machine room and I remember often opening the door and seeing Tridge in there with ugg boots and a heavy coat debugging by watching the LEDs on the boards.

Dave Sitsky's memory of it.

... Andrew worked passionately throughout the night, and when we saw him arrive at work, a slightly bleary-eyed Andrew showed us the next amazing milestone he had reached. It wasn't all easy sailing for Andrew though – I remember coming into his office and seeing him smashing his keyboard in frustration!

A few short weeks after installation of the AP+, Paulus and Tridge visited Japan and presented their work to Fujitsu:

> **FJ:** "But you must have implemented a driver for the broadcast network. How long did that take?"
>
> **Paulus:** "Yes, I did that. It took one week."
>
> **FJ:** "Aaah. We had budgeted six months."

A communications benchmark for parallel machines is the daisy-chain benchmark. One node transmits a small data packet to a neighbour which then forwards it to another, and so on, until the packet is back at the original source, having passed through all the nodes. When there is no co-ordination between nodes, the benchmark may take ages, due to scheduling delays on each node.

Tridge wrote a simple "loose-gang" scheduler for AP/Linux. Whenever a node was about to schedule a process, it checked whether it had a runnable process with the same id as the running process on node 0. If so, it immediately ran that. This resulted in a huge speed-up of daisy-chain.

Fujitsu AP3000

One of the difficulties faced by Fujitsu in achieving sales of the AP series machines was that, in that era, the speed of individual chips was increasing rapidly. If an n-node AP1000 based on a particular SPARC chip were released simultaneously with the release of Sun workstations using that chip, then the maximum n-fold speed-up might justify the cost and the complication of writing parallel code to use it. Unfortunately, design and development by Fujitsu of AP boards took a considerable time, and by the time customers received their new machine, the maximum speed-up relative to a Sun workstation was only, say $n/2$. Furthermore, most applications achieve only a fraction of the maximum speed-up. Finally, by the time your AP was three years old the speed-up relative to newly released Sun workstations would have reduced much further.

In order to speed up development time, the AP3000 successor to the AP+ used Sun SparcStation boxes as nodes. The machine consisted of large cabinets with racks into which Sun pizza boxes were slotted. Fujitsu's value-add was in the network which connected them, and the software.

ANU also acquired an AP3000, but by that stage the CAP Project was running out of steam. The cost of the system was a considerable barrier to purchasers and the number of applications for which it was compellingly superior to other options, such as ethernetted workstations or PCs, was small.

Section 8.1.3 on Page 218 describes a low-cost, DCS-built machine of similar architecture to the AP3000.

ANU-Fujitsu CAP Project Lessons

I should note that the CAP project was folded into the ACSys CRC once it got off the ground in 1993.

The CAP project was an amazing opportunity for the ANU people who participated. It was a rare occasion when researchers with very diverse research interests felt a common interest, and worked as a loosely coupled research team. As Steve Blackburn says, Steve Blackburn's memories of the CAP Project

I think the great Linux / BSD battle captures a number of interesting aspects what we were up to, the enthusiasms and competitive nature of our friendships. I fondly remember many years of lunches in the union with a large gang of people which always finished with Brian Corrie wandering up to the union shop and coming back with a packet of two Reece's Peanut Butter Cups, which apparently kept him in touch with North America. My recollection is that Brian introduced us to ultimate frisbee while we were in the Crawford building and we played on the lawn in front of Chifley. One time Leon Smith stumbled upon us and showed us how to play properly and, before too long, ultimate became established in Canberra.

Dave Sitsky's time with the ANU-Fujitsu CAP Project

Post graduation, I was tempted to continue to do a PhD, however I also wanted to build software that would benefit people. I was fortunate to be offered a position in the CAP Research Program, a collaboration between the Computer Science Department and Fujitsu Laboratories in Kawasaki, with David Hawking as my boss. As a recent graduate, suddenly having access to exotic machines like the Fujitsu AP1000, effectively 128 simple computers connected by a high-speed network was a dream.

Under the tutelage of David Walsh, I got up to speed on how to program the AP1000 and what traps to be aware of. The CellOS operating system that ran on the Fujitsu AP1000 was a relatively simple embedded operating system that allowed for the execution of a single parallel program, with special calls for message passing between nodes. David was working on an extension to CellOS called "message domains" so that message passing libraries could use individual channels to avoid cross-talk with other libraries – effectively a namespacing mechanism. Peter Bailey was also working on a parallel version of the ML language that had "channel communications" that could run on the AP1000.

I started to get more exposure to real-world scientific modelling programs that run on these exotic supercomputers. Chris Johnson was enthusiastic about PVM, which at the time was a defacto standard for message passing applications. A short time later, the MPI (Message Passing Interface) standard was created. During this time, multi-core machines didn't exist, so parallel computation existed either on COWs (cluster of workstations) or on dedicated supercomputers like the AP1000.

As a part of my work on the CAP Research Program, I created an implementation of MPI for the AP1000. This allowed many existing MPI programs to run on the AP1000 and received a lot of interest from Fujitsu labs. I was fortunate to attend conferences with other MPI implementers and share notes. This led to two stays at Fujitsu labs in Kawasaki, Japan, on the order of three months each, where I was able to collaborate closely with the Fujitsu hardware and software engineers to optimise the MPI implementation for the AP1000+ and later the AP3000.

During my time on the CAP Research Program I was extremely fortunate to collaborate with two very talented software engineers – Paul Mackerras and Andrew Tridgell. Paul introduced me to the NetBSD networking code, with which he was very familiar, having written the PPP implementation for it, and showed me many aspects of the FDDI (Fiber Distributed Data Interface) he designed both in hardware and software for fast data transfers to/from the AP1000. I learnt an immense amount from Paul and Andrew in an unofficial "advanced operating systems and programming" course.

My time ended at the CAP Research Program by the end of 1998. In 1999 I had a "gap year" with my future wife teaching English in Nepal and travelling in Asia, before settling in Sydney in 2000.

For most of us, the CAP project was a once-in-a-lifetime opportunity to have such close interaction with, and input to, the development of a commercial supercomputer. A lot of good software was developed and we learned lots. In my case, developing a parallel document retrieval engine (PADRE) for the AP1000 was the first step in a 25-year research and commercialisation career. If there is a criticism to be made of the project, it is that we weren't as good as we might have been in turning CAP-related work into published outputs or external impact.

8.1.3 The Bunyip

In 2000 Bob Edwards began constructing Bunyip out of low-cost PC boxes running Linux. He made a supercomputer out of library shelves! In the panel on Page 219 Bob Edwards explains where the money came from.

SC2000
GORDON BELL PRIZE

Robert Edwards

Winner, Price/Performance Category

92¢/MFlops/s, Ultra-Large-Scale Neural-Network Training on a PIII Cluster

SC2000 CONFERENCE • NOVEMBER 04-10, 2000 • DALLAS, TEXAS USA

LOUIS TURCOTTE, SC2000 GENERAL CHAIR

JAMES McGRAW, CONFERENCE VICE CHAIR, TECHNICAL PROGRAM

Bob Edwards with his SC2000 Gordon Bell Award for Price Performance, achieved with Bunyip. ***Photo: David Hawking***

The Bunyip comprised 96 Cougar[12] nodes with dual CPUs per node (192 x 550MHz Pentium III). The CPUs on a node shared 384MB of RAM. Each node had three 100MB/sec network interfaces, connected in a tetrahedral mesh. Bisection bandwidth was approximately 28.8 Gbits/sec. In November 2000, Bob along with Doug Aberdeen, and Jon Baxter were awarded the SC2000 Gordon Bell Prize in the Price/Performance category. As Bob says:

Bunyip's Gordon Bell award

By that time, we had managed to get performance to 92c per MFlops for a "non-trivial" parallel programming task (Kanji character recognition, from memory).

From the e-mail trail, the project was funded jointly by RSISE (John Lloyd as head) and by DCS (Chris Johnson as head). May have been more complicated than that (eg. an ANU Major Equipment grant), but my role was more technical/design – others did the budgeting stuff – so I don't really know.

We spent just shy of $A250k (most went to Cougar, some to HP for the network switches, there were other bits and pieces as well, notably 288 colour-coded Cat 5 network cables...).

[12] Cougar were a small Canberra company who assembled custom IBM-compatible computers

Chapter 9

ANU: Student Computing

Student computing up until the end of the 1970s was supported on computers whose main role was academic and/or administrative computing. That era is covered in Chapter 2.

The commencement of operations of the SGS DEC-10 in second semester 1979, coupled with conversion of Copland G5 into a terminal room, marked the first time that ANU provided computing resources specifically for students. Later that DEC KA-10 was replaced by a more powerful DEC-10, then by VAXes, then by Unix servers, and later Unix workstation laboratories. For a period X-terminals were used and, over time, laboratories of public access Macintosh and IBM compatible (later Windows) workstations were created.

The small group which supported the initial DEC-10 was known by the same name as the computer – "The SGS DEC-10". When the School of General Studies became "The Faculties" the group became known as the Faculties Computer Unit (FCU). Later, its scope was broadened as the Teaching and Learning Technology Support Unit (TLTSU).

9.1 The DEC KA-10

I have previously described (see Page 73) the arrival of the hand-me-down DEC KA-10 in the Copland Building. As I reported to the DEC-10 Management Committee meeting on 31 July 1979, the KA-10 had been operating since 18 July, despite on-going carpentry and other works:

- 28 (out of a possible 32) terminals were connected, and
- 30 individuals and 4 courses had been allocated PPNs (accounts). By March 1980 the number exceeded 600.
- The total memory available to users was initially 84K 36-bit words (144K – 60K words for the TOPS-10 monitor[1] but removal of unnecessary features in the monitor increased available user memory to 95K words.
- Stringent quotas were applied to memory, disk space and printouts. The limits were relaxed a little between 02.00 and 10.00 each day.
- Game playing was banned by the Committee and the games normally installed on a DEC-10 were deleted. ... The most popular game was *Adventure*. ☺
- At 10.00 each weekday, disk contents were backed up. On Mondays this was a full backup, and the system was unavailable to users. On the other weekdays only files changed since the previous day were backed up, and users were allowed.

Reliability of the KA-10 was quite poor for a long time due to unsatisfactory air conditioning and power spikes. The air conditioning units were designed for domestic use, and kept icing up. They also caused rapid fluctuations in temperature.

[1] Monitor == operating system.

The system was taken down for ten days while better airconditioning was installed. A thermal shut-off switch was installed to protect against further airconditioning failures.

My October 1979 report to the Management Committee reported that 170 first and second year computer science students had been given terminal access from 01 October that year.

By the November 1979 Management Committee meeting, SGS had agreed to purchase a 64K word memory module from RSPhysS which increased user memory by 67%. The module in its very large cabinet was delivered by bare-footed removalists from Tommy Tortoise and the old unit was left outside the computer room.

Unfortunately, the next day was the Friday of Bush Week, and the memory module was appropriated by students and carried off to Union Court where Kambri now is. We were fortunately able to retrieve it undamaged.

ANU old-timers will be familiar with the poplar trees which used to line University avenue. In late September or early October they produced large quantities of white fluff which covered the ground like snow drifts. Student legend had it that if you hadn't started studying for the end-of-year exams by then, you were in trouble.

For a period, the lowest level of Union Court featured a long, shallow drain stretching about 100 metres from the bookshop to the pharmacy, and covered by metal grating. This drain would sometimes fill up with poplar fluff, and I was told that one dry October, a student checked to see whether the fluff was inflammable. He lit it near the pharmacy and a sheet of flame shot along the drain, causing minor burns to the legs of someone stepping over it near the bookshop.

Response time at peak times on the KA-10 was quite poor. Not surprising really, considering the number of students squeezed onto a machine, which was by today's standards, very weak indeed.

David Berriman, a student who went on to a distinguished career in IT around the campus, describes how he and a few colleagues avoided slow DEC-10 response in 1980 when doing assignments in MACRO-10 (DEC-10 assembler). They would spend the evening in the union bar, head to the Dolly's fast food van in the carpark behind Copland at 1am, then spend the 2am to 5am period working on their assignments on the KA-10.

Once first year computer science students were granted interactive access to the KA-10, later year CS courses were transferred to the Univac, accessed via the CreasyLink and later via an ANUNet node. By mid 1981 proposals for major upgrade of the student computing facilities were being considered by the University's funding committees.

When the RSPhysS KA-10 was decommissioned at the end of 1981, the student KA-10 was upgraded to its maximum of 256K words, and an extra eight terminal lines became available.

The DEC-10 Management Committee was chaired by Ray Jarvis and initially included Ray Byron (Statistics), David Hawking (SGS DEC-10 Site Manager), Ken Johnson (Geography), and Bob Landford (CSC). A representative from the Department of Pure Mathematics joined later. Initially that was Neville Smythe, later it was Martin Ward. Soon afterward, Arthur McHugh (Accounting & Public Finance) and Bob Cushing (Sociology) joined the committee.

At the second meeting of the Committee, a request was received from Prehistory & Anthropology, Geography, and Sociology for student access to a small version of the SPSS statistical package to be installed on the KA-10.

A major topic of discussion for the Committee over its first year and beyond was the provision of a mini-network to provide terminal access to Dedman, A.D. Hope and other areas of The Faculties. There were initially only 32 terminal lines on the KA-10, increasing to 40 with the addition of further hand-me-down equipment from RSPhysS. Since lecturers needed access to KA-10 terminals in order to prepare lessons and assignments, there were sometimes tense discussions about allocation of terminal lines to departments.

Two operators were initially employed to supervise the SGS DEC-10: Leah Cox and Catherine Thorpe. It soon became clear that only one operator would be needed. That was Leah. After she left, Inta Skriveris transferred from other duties in the CSC. Other CSC operators came over when necessary to relieve the usual operator.

The KA-10 was equipped with a seven-track tape drive and two nine-track drives. We had the only seven-track drive on campus and provided a seven to nine track tape conversion service, inherited from the Joint Schools. On one occasion, we prepared a 7-track boot tape for the DEC-10 at Auckland University, which only had 7-track drives.

Resourcefulness rescues expensive scientific data

Every week Doug Christie from the Research School of Earth Sciences (RSES) would bring over seismological recordings from Tennant Creek, NT, on seven-track tape. He also brought an equivalent number of brand new tapes, and the DEC-10 operator would transcribe the data from seven to nine track. From memory there were two or three tapes transcribed per week, meaning that the number of new tapes transcribed over the period we had the KA-10 would have been close to five hundred!

One day, the batch of seven-track tapes could not be read. This was disastrous because the recordings should have included the results of a hugely expensive Australian-American experiment in which a large quantity of explosives was detonated by a chartered ship in the Timor sea.

Nothing daunted, Doug's team measured the spacing between the inter-record gaps on the tape and realised that the recorder in Tennant Creek had been running slow, but fortunately at a consistent speed. (A capacitor in the recorder had failed.)

DEC field service engineers were summoned – could we increase the speed of the motor by 14%? to compensate for the lower recording density on the tape? Unfortunately not!

To the rescue came the RSES workshop, who manufactured a new capstan for the tape drive, larger in diameter by just the right amount. Problem solved. Millions of dollars invested in the experiment had not been wasted.

Up until the 1980s it was normal for computer rooms to have an Emergency Stop button. One of the more interesting of the many causes (mostly related to airconditioning failures) of downtime on the SGS DEC-10 resulted when a ladder kept in the computer room to facilitate diagnosis of airconditioning faults, was bumped and fell against the Emergency Stop.

9.2 Students v. DEC-10 Staff

Most operator tasks required that the operator log in as `[1,2]`, the super-user, which of course required a secure password, changed regularly. To communicate the current password to the relieving operators, we taped it to the inside of a filing cabinet and left a humorous note in the log book saying, "If you need the password, look inside my drawers. Signed, the filing cabinet." After a while, suspicions grew that some students were making use of the `[1,2]` account. We eventually suspected that they were using a card or a blade to pick the lock of the computer room door. To test the theory we wrote an incorrect password in the usual place and monitored the logs. DEC-10 logs recorded not only the account associated with a failed login attempt but the password which had been used. A record of the bogus password in the log confirmed our suspicions. Realising that the game was up, the perpetrators left a pair of underpants (drawers) in the computer room.

The DEC-10 logs also helped us identify the perpetrator in the case where a student unplugged the serial cable from a terminal in Copland G5, ripped off the connector and pushed the wires into a power outlet, causing expensive damage to parts of the DEC-10. The incident happened just after closing time at the Union Bar, after the end of semester, and the perpetrator attempted to log in to an account which was no longer active. However we could tell that he used the correct password for the account. Consultation with DCS revealed that the account belonged to a student who failed the unit he was enrolled in after being penalised for cheating. What we thought was compelling evidence was presented to the person in the Chancelry responsible for discipline, and ....., nothing happened.

Logs were not able to identify the person who poked a fire hose through the ventilation louvres in one of the doors into Copland G5 and turned on the water.

In the panel on Page 223, Robert Cohen reveals several successful attempts by him to circumvent controls placed on students accessing the DEC-10. We weren't even aware that these tricks were being played. Well, we knew that students were reactivating the Setup function on VT100s but I don't remember hearing about the mercury switch![2]

[2] Geoff Huston remembers the mercury switch and thinks Dean Spire actually found it.

Robert Cohen recounts stories of the arms race between students and DEC-10 SysAdmins.

Back in 1983, ANU Comp Sci was using a DEC KL-10 for undergraduate teaching. The undergrad terminal room was filled with VT100s connected to serial lines. The equipment was capable of doing 9600 baud, but there was a concern that the serial processors couldn't handle the number of terminals in the room at full speed. So the terminals were turned down to 1200 baud. Which was agonisingly slow. A screen refresh took about 10 seconds.

However, the terminals had a `setup` button which could be used to set the speed back up to 9600 baud and the serial controller would autosense the speed of that line. So the setup buttons were all disabled by cutting a trace on the circuit board inside the keyboard.

For a while there was a game of cat and mouse where an enterprising student would repair one of the keyboards by soldering back across the gap. And students in the know would hunt around the room for the "good" keyboard to fix their terminal. And every so often the staff would do a sweep and disable any keyboards found to have a working setup key.

However after a few months of this the students came up with a brilliant workaround. A mercury switch was used to bridge the gap. So the setup key only worked if the keyboard was held upside down. The staff could hunt for working setup keys to their hearts content, but never figured out how the speeds kept getting increased.

There was a concern about students stealing other students' work. While the OS did have permissions that could be used to stop students seeing each other's home directories, students could [and did] accidentally [or intentionally] give their files the wrong permissions.

The TOPS-10 OS used by the KL-10 had a concept of privileges that could be assigned to users. So a local kernel hack was put in place to implement a special "privilege." To prevent students stealing each other's work, a special privilege was associated with the student accounts which meant they could only access files in their home directory and system directories. But never other students home directories – regardless of file permissions.

Unfortunately, it was implemented as an *unprivilege* i.e. being granted this reduced your access. Whereas it should have been done as a privilege given to everyone else.

Now, while I had no need to look at other students work, I of course treated the restriction as a challenge to get around. I made a lucky discovery. The OS included a command which allowed you to arbitrarily give up privileges. So revoking all your privileges granted you the ability to roam the filesystem to your hearts content (permissions permitting). Then I made another happy discovery. Revoking your privileges had another side effect. The ANU at that time was imposing connect time, CPU time and printer quotas. However it turned out there was a privilege needed to update quota files and record your usage. So revoking your privileges also freed you from the constraints of quotas.

The next year, it was decided that giving students command line access was just too dangerous. So the students were restricted to a menu system named RUFUS (Restricted Usage For Undergraduate Students).[a] It let you edit files in your home directory and compile and run your programs but nothing else. However it turned out there was a timing bug. If you logged into two terminals simultaneously, there was a chance that one of the terminals would get a normal command prompt instead of RUFUS.

Student attempts to escape the strait-jacket imposed by PASTIE and later RUFUS, were not always successful. We had a good chuckle when one student escaped from the PASTIE menu system, but found themselves in no-man's land – unable to to do anything, even to logout.

Despite all the subversive activities, Robert was fundamentally an honest person. When an operator inadvertently sent the entire plain-text password file of the DEC-10 to the printer in the student terminal room, Robert returned it! Robert ran a Canberra ISP for a number of years before working in the Unix Support Group at ANU ITS for 20 years, and very recently moving to NCI.

[a] RUFUS, developed by Peter Hobson, was a generalised version of PASTIE.

Stretching of the bounds of acceptable behaviour continued long after the DEC-10. Hugh Fisher recounts one such case from around 1996:

> Being one of the computer system administrators at ANU Computer Science was always challenging and interesting. Staff and students were always trying to do something new, and sometimes they created new problems as well.

One day an entire room of about twenty student computers suddenly crashed. This was over a century ago in Internet years, so all the computers were Sun workstations running a Unix OS. The desktop PCs and Macs of that era were underpowered and unable to cope with misbehaving programs, but these Unix workstations were supposed to be industrial quality and rock solid. So an entire roomful crashing at once was very disconcerting for us.

As we frantically wondered what could have happened, a student appeared at the door of our shared office space. Not wanting to be rude, we said we'd be happy to help but now was not a good time because a lot of computers had suddenly crashed. The student replied that they might know something about that.

That morning a computer security mailing list had announced a new “Ping of Death”,[3] a carefully designed IP network packet that caused a computer receiving it to crash. And this was a multicast ping of death, so could be sent to every computer on a local network in a single operation rather than to just one. Our student had read this and thought it was ridiculous and couldn't possibly work.

Well, we wanted students to be curious and explore their interests outside the curriculum, and no actual damage had been done. We thanked the student for their honesty, and asked them to please check with us before carrying out any similar experiments in future.

9.3 Copland G5: The First Student Terminal Room

Brian Molinari recounts the creation of Copland G5.

> In 1979 a DEC 10-KA mainframe was installed in the Copland building by the Computer Services Centre and used to support teaching. It was administered by CSC staff assigned to a Faculties Computing Unit. A nearby tutorial room Copland G5 was converted to a “terminal room” and terminals together with a line-printer were installed. The card punches were retired.
> It was hardly nirvana. In my words from a subsequent report:
> This lead to the Kafkaesque Copland G5, which in its infamous heyday had around 50 terminals crammed into its 80 square meters. The ambience of the room was established by the television security cameras and neatly set off by the complete absence of windows and natural light. The security cameras were sabotaged by students (and one could hardly blame them).

9.4 Faculties DEC KL-10

In 1982, it was clear that the Faculties DEC KA-10 system would be unable to meet the pent-up demand for interactive access by students. The Computing Policy Committee approved a substantial upgrade, in paper 3.19.28B, 18 March 1982, with the following caveat:

> (b) No additional contributions to the system shall be provided from central funds for shared general purpose computing within such period after installation (say 5–7 years) as appears appropriate in the light of growth in teaching-based needs of The Faculties.

Wow! No upgrades for five to seven years!

A procurement exercise was undertaken, led by a committee chaired by Bob Watts. Its four other members included William Dunsmuir from Statistics, Ray Jarvis from DCS and Bob Landford and Robin Erskine from the CSC.

The budget was $750,000, a handsome sum because at the time $A1 was worth about $US1.27. Warned by the failure of some previous benchmarking exercises, I put a lot of effort into developing an accurate model of projected student computing load. The benchmark tested whether a system could maintain fast interactive response, despite a realistic heavy load from statistical packages, CAL (Computer Aided Learning), and Edit-Compile-Execute cycles in various languages. At the same time it revealed whether the necessary packages and languages were available.

[3] https://en.wikipedia.org/wiki/Ping_of_death

There was huge vendor interest in the opportunity – 36 copies of the Request For Proposals were given to suppliers, and 17 attendees at the bidders meeting took away benchmarking tapes.

As Technical Adviser to the committee, I travelled the country, spending time at universities using FACOM, DEC-10s, VAX/Unix, Univac, etc. for student computing, and reporting back to the committee. Prime and Honeywell were also in the mix. Robin Erskine remembers travelling with me to Flinders University, passing through Melbourne, soon after his arrival from Scotland. "Coming home that night after the day's travel I said to Enid that I had done the equivalent of Edinburgh – London – Geneva return in one day. I had traveled to Geneva frequently for my PhD and work but never in one day."

Digital Equipment Corporation submitted two bids, one based on a pair of VAX 11/780s running VAX/VMS, and the other on a DEC KL-10 running TOPS-10. The benchmark showed that the KL-10 gave significantly faster interactive response than a VAX running half the load, and both of them outperformed the other benchmarked systems.

The dual VAX/VMS proposal had a problem: DCS feared that it would be forced to use only one of the VAXes, and the other departments feared that DCS would get a VAX to itself.

Several vendors didn't attempt to run the benchmark, or ran parts of it which were meaningless on their own. IBM ignored the procurement process and went straight to the Vice-Chancellor with an offer to solve all of the University's computing problems (incidentally including student computing). Since I would be the one fielding the complaints from disgruntled students and lecturers if the KA-10 replacement didn't perform, I was pleased and relieved that this unbenchmarked offer wasn't taken up.

The procurement committee was determined that there would be a play-off between two short-listed vendors. Bob Watts insisted that DEC should withdraw one of their bids. They chose to continue with the KL-10, and it was played off against a new Honeywell system, DPS-6, running GCOS6, which hadn't run the benchmark. Honeywell ran a non-interactive and essentially irrelevant benchmark against an older version of their proposed system, and then projected what the timings would be on the new system. During negotiations they offered to beef up the RAM configuration of their proposed system to 16MB, which they claimed would completely obviate the need for paging or swapping, and thereby achieve good response.

Since the Honeywell option hadn't been benchmarked and hadn't been used to support teaching in any comparable university, I considered that their offer entailed an unacceptable degree of risk. Negotiations dragged on and hadn't reached a decision by the time I left on an extended overseas holiday. When I returned, I discovered that the committee had chosen the lower-risk DEC KL-10. My team's job of switching over from one physically large computer system to another within a confined space during a teaching break and getting all the required software up and running, was thus much easier and entailed much lower risk.

Geoff Huston stood in for me while I took leave in Europe and in June 1983, went on four months leave himself. Bob Landford advised that when he returned he would not return to The Faculties but would join the network group in the CSC.

The decision to upgrade the student computing facility had ramifications:

- All usage had to be interactive, because no card reader was purchased;
- Thirty new VT100 compatible terminals were purchased;
- Copland G18 was converted into a new terminal laboratory.

Several operational innovations occurred in parallel with the KL-10. An enrolment tape from the university administration allowed us to dramatically speed up the issue of student PPNs (accounts). We also had payslip-type stationery printed to allow secure issuing of passwords to students. Finally, the new `KANGEL` program allowed teaching staff godly powers over their students, allowing them to inspect files, look at usage data, and refresh quotas.

I met with all five Faculty Deans, to communicate the expanded capability, and to scope out likely demand from students in each faculty. By March 1993, there were nearly 2000 accounts on the KL-10.

I must have come back from leave with a light heart. My report to the Management Committee included the following:

> **Unseemly haste** – Six months ago we requested The Faculties for a small room in which to store line printer paper. Quick action is expected some day.

9.5 Faculties VAXes

After about a year of service, DEC anounced that its Jupiter project to create a successor to the KL-10 had been abandoned and that the company would consolidate its efforts on the VAX line. ANU, led by Bob Landford, Director of Computer Services, claimed that ANU had been sold the DEC-10 on false premises and managed to negotiate the replacement of the KL-10 by two VAX 11/780s configured as a cluster and a smaller VAX 11/750. However, the VAX 11/750 would not be part of the KL-10 replacement but would be located in the Leonard Huxley building to provide a Unix computing resource, 50% for a new undergraduate course taught by Ron Bird in Commerce, and 50% to provide a resource for researchers in JCSMR.

A significant battle ensued between the Faculties and the CSC about the location of the new VAX 11/780s. The CSC argued that newly released MICOM equipment, interconnected with high speed links, would permit the VAXes to be located in the Leonard Huxley Building, comfortably supporting the required number of student terminals in the Copland area. Past experience with cut cables at ANU,[4] and lack of experience with MICOMs, caused great concern in the Faculties about potential unreliability.

Changeover to the VAXes needed to take place in the mid-year break, and the short time available for change-over would not permit them to housed in the room occupied by the KL-10. A fly in the ointment for locating them in the Faculties area was thus the need to find a room which could be converted to a machine room. Robin Stanton launched into a forensic analysis of lecture and tutorial timetables and proved that all classes could be comfortably accommodated with one less teaching room. Copland G4 was chosen, next door to the big (48 terminals) Copland G5 terminal room.

Micoms were installed anyway and provided a big upgrade to staff access to central computers, but the VAXes came to G4.

Based on results of benchmarking in the procurement exercise I was very worried that performance of the VAXes (called VAXen by cognoscenti) would be inadequate. I designed a new computer room for them in Copland G4 (see sketch plan on Page 227), and oversaw the considerable task of providing all the facilities needed by the courses on the DEC-10. At the same time I asked to be relieved of the position of Head of the FCU, and given a new role in the Computer Services Centre, once the VAXes were up and running. The two FCU Vaxes were known as FAC1 and FAC2.

My September 1984 report to the Faculties Computer Management Committee (see extract on Page 228) shows that unreliability of the new VAXCluster was of more concern than response time.

Over time a series of upgrades to the VAX configuration were undertaken, but through the decade, DEC focus on the university sector diminished, and the value-for-money of its offerings declined relative to more recently established vendors such as Sun Microsystems, Apollo, and Silicon Graphics.

Demand for student computing continued to increase through the 1980s and FAC1 and FAC2 were eventually augmented by two Sun 4/280s known as FAC3 and FAC4. Brendan McKay's letter reproduced on Page 229 highlights the still-increasing demand for terminal access by students, and the continuing problems with interactive response on the VAX Cluster.

[4] In one case a major communications cable was cut and not repaired for days because of a nest of jumper ants in the pit.

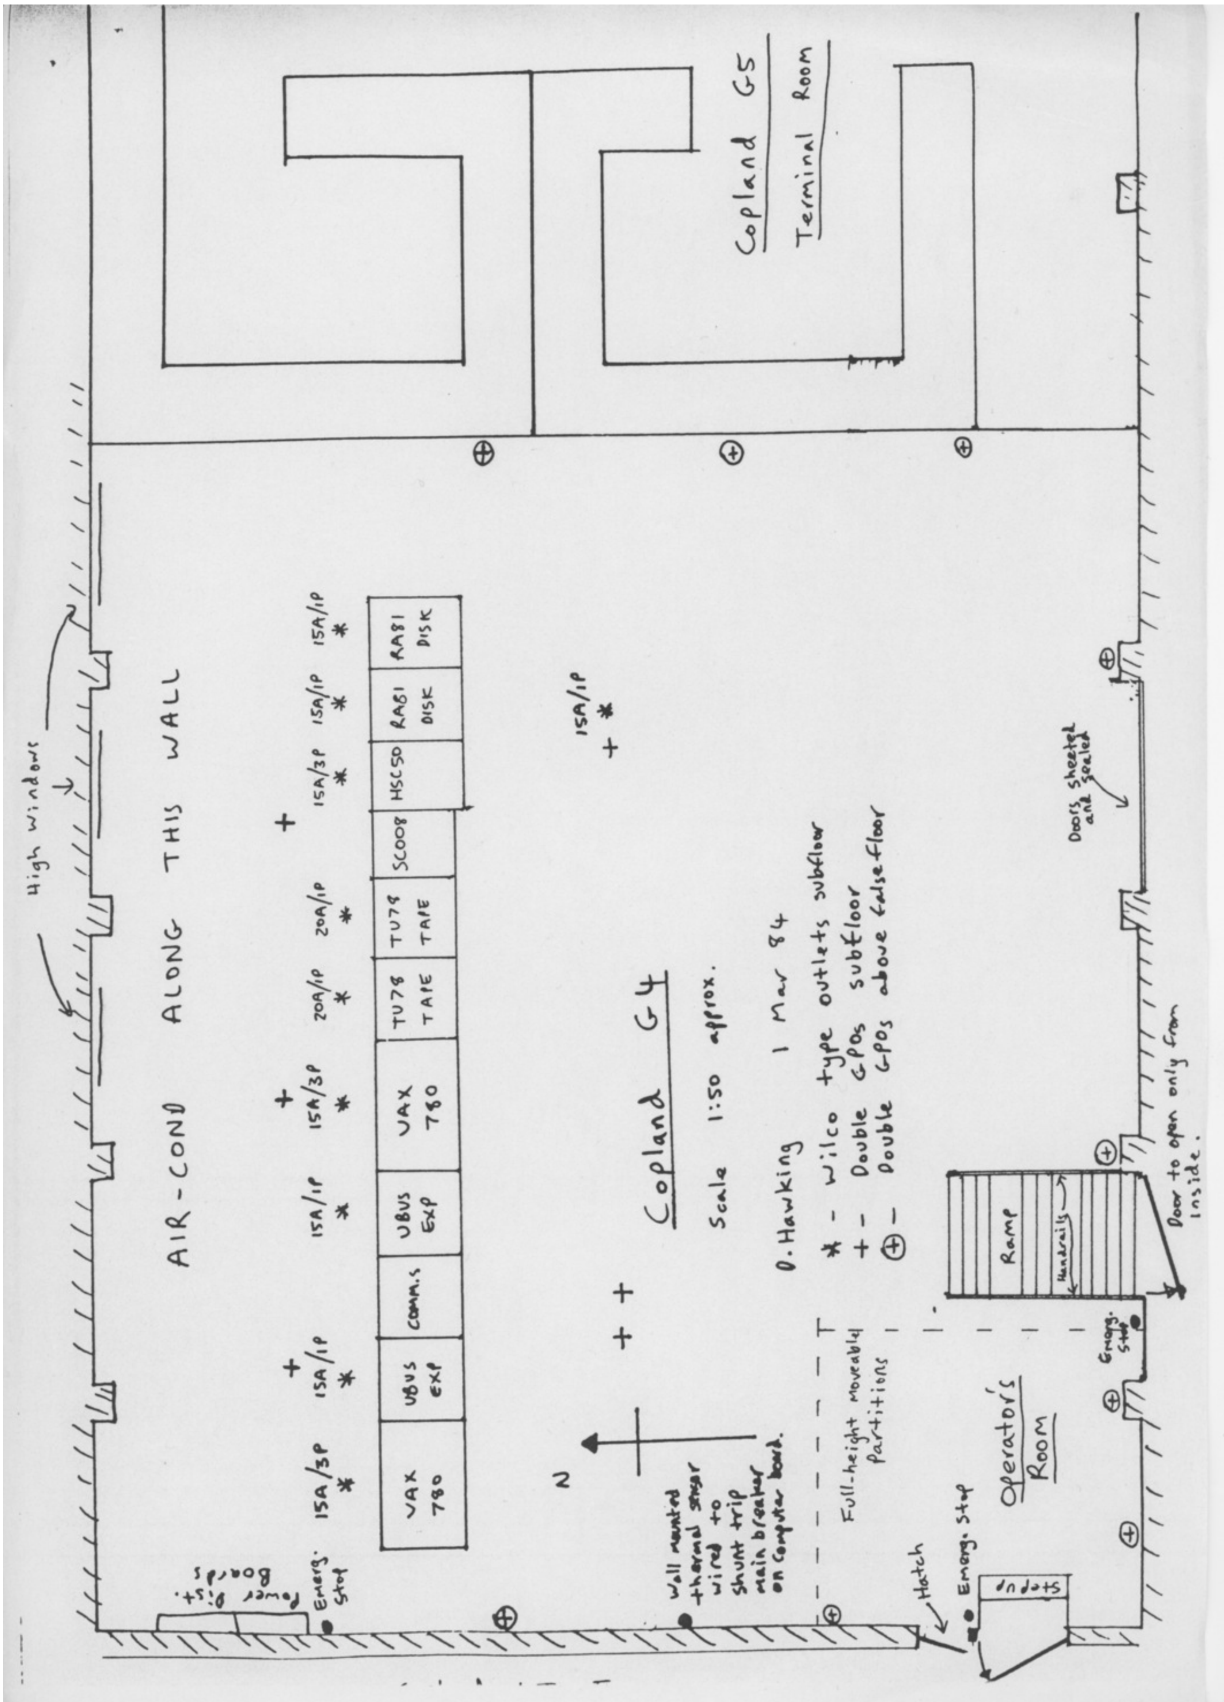

1984: My design for the conversion of lecture room Copland G4 into a machine room for the FCU VAXes. If John Warnock's PostScript language and Apple LaserWriters had been available then, I would have drawn this as a PostScript program. The interesting layout of benches which can be seen next door in the Copland G5 terminal room was due to Ray Jarvis. That terminal room held a maximum of 48 terminals and came into service soon after the arrival of the KA-10.

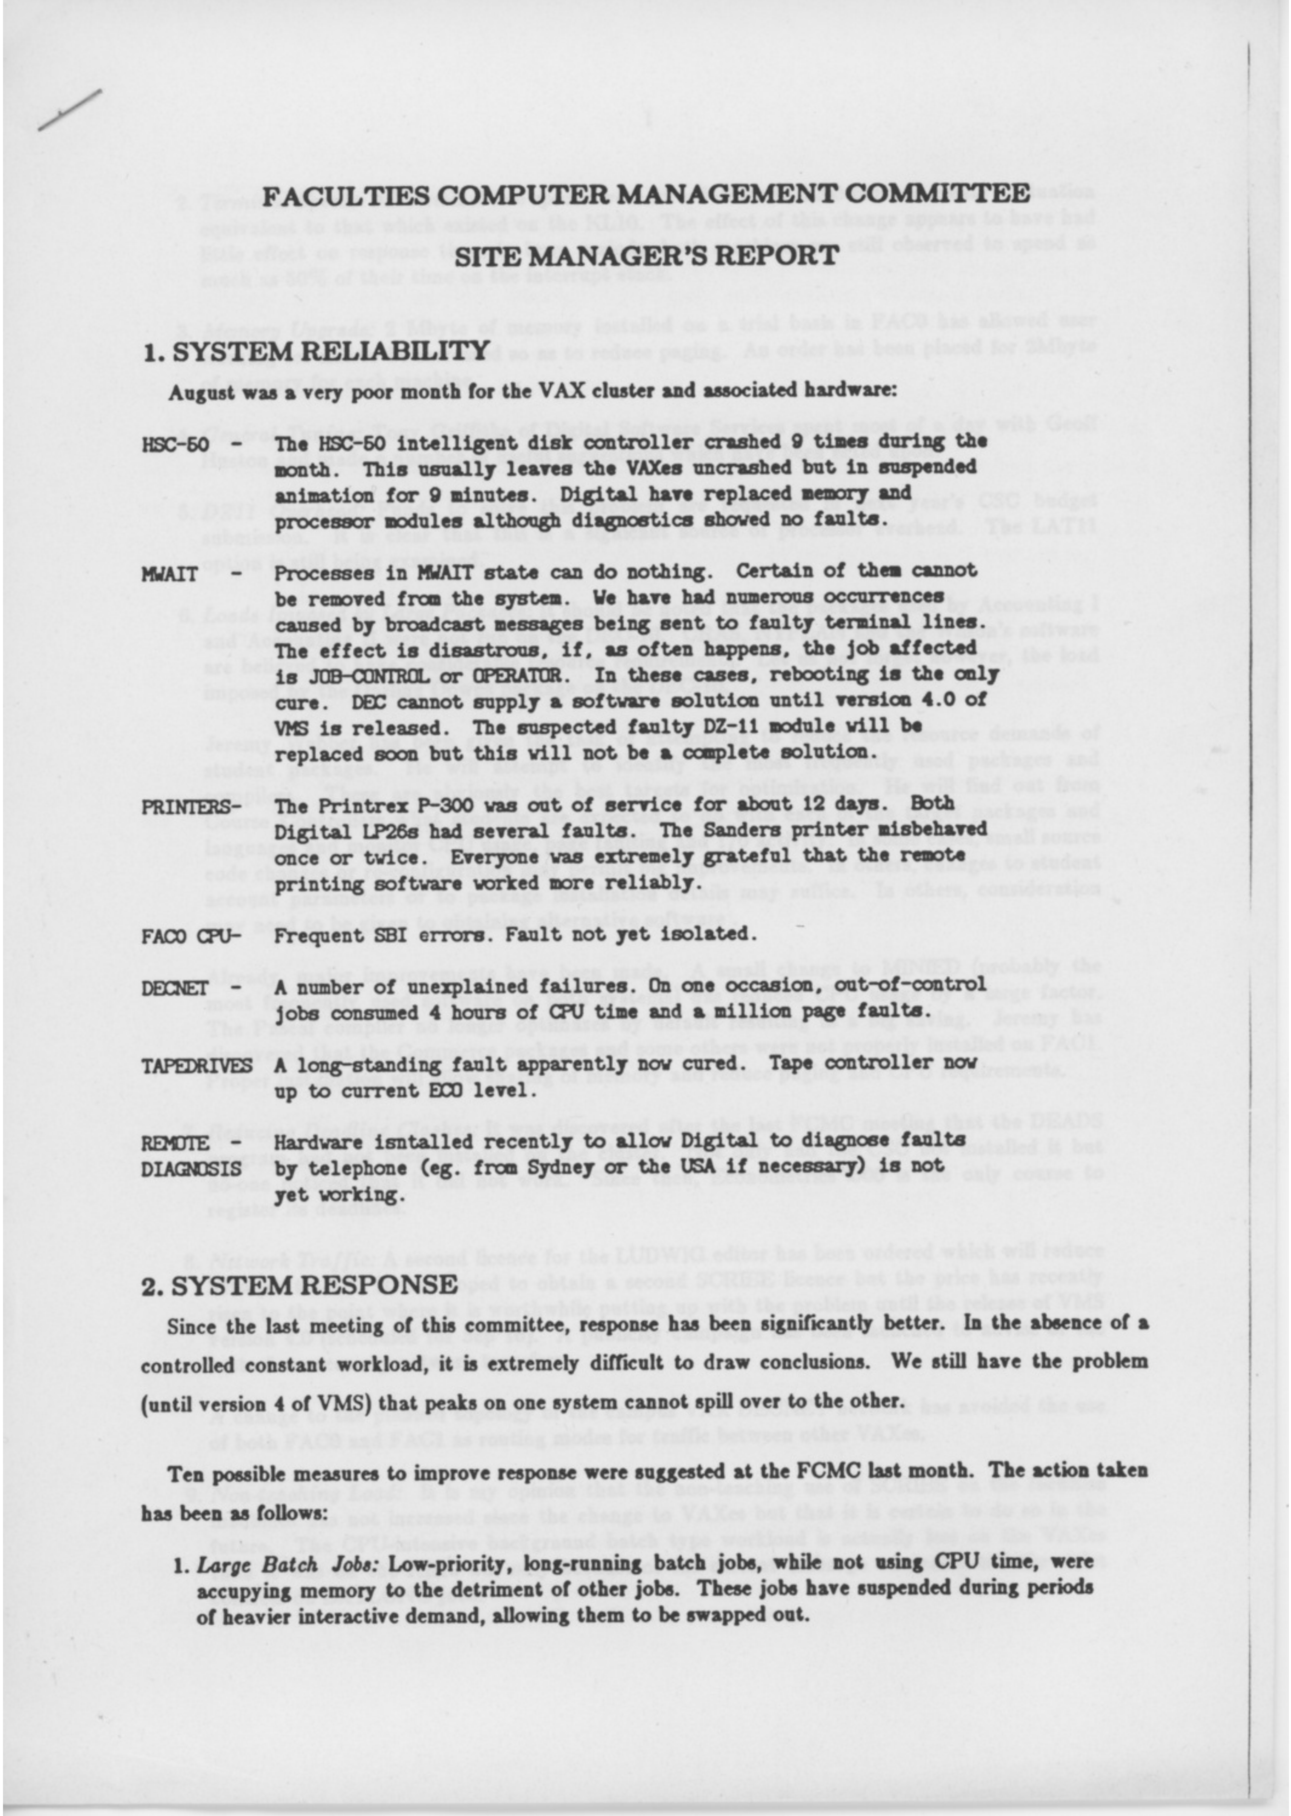

FACULTIES COMPUTER MANAGEMENT COMMITTEE

SITE MANAGER'S REPORT

1. SYSTEM RELIABILITY

August was a very poor month for the VAX cluster and associated hardware:

HSC-50 - The HSC-50 intelligent disk controller crashed 9 times during the month. This usually leaves the VAXes uncrashed but in suspended animation for 9 minutes. Digital have replaced memory and processor modules although diagnostics showed no faults.

MWAIT - Processes in MWAIT state can do nothing. Certain of them cannot be removed from the system. We have had numerous occurrences caused by broadcast messages being sent to faulty terminal lines. The effect is disastrous, if, as often happens, the job affected is JOB-CONTROL or OPERATOR. In these cases, rebooting is the only cure. DEC cannot supply a software solution until version 4.0 of VMS is released. The suspected faulty DZ-11 module will be replaced soon but this will not be a complete solution.

PRINTERS- The Printrex P-300 was out of service for about 12 days. Both Digital LP26s had several faults. The Sanders printer misbehaved once or twice. Everyone was extremely grateful that the remote printing software worked more reliably.

FAC0 CPU- Frequent SBI errors. Fault not yet isolated.

DECNET - A number of unexplained failures. On one occasion, out-of-control jobs consumed 4 hours of CPU time and a million page faults.

TAPEDRIVES A long-standing fault apparently now cured. Tape controller now up to current ECO level.

REMOTE - DIAGNOSIS Hardware isntalled recently to allow Digital to diagnose faults by telephone (eg. from Sydney or the USA if necessary) is not yet working.

2. SYSTEM RESPONSE

Since the last meeting of this committee, response has been significantly better. In the absence of a controlled constant workload, it is extremely difficult to draw conclusions. We still have the problem (until version 4 of VMS) that peaks on one system cannot spill over to the other.

Ten possible measures to improve response were suggested at the FCMC last month. The action taken has been as follows:

1. *Large Batch Jobs:* Low-priority, long-running batch jobs, while not using CPU time, were accupying memory to the detriment of other jobs. These jobs have suspended during periods of heavier interactive demand, allowing them to be swapped out.

1984: The first report to the Faculties Computer Management Committee following the change-over to VAXes.

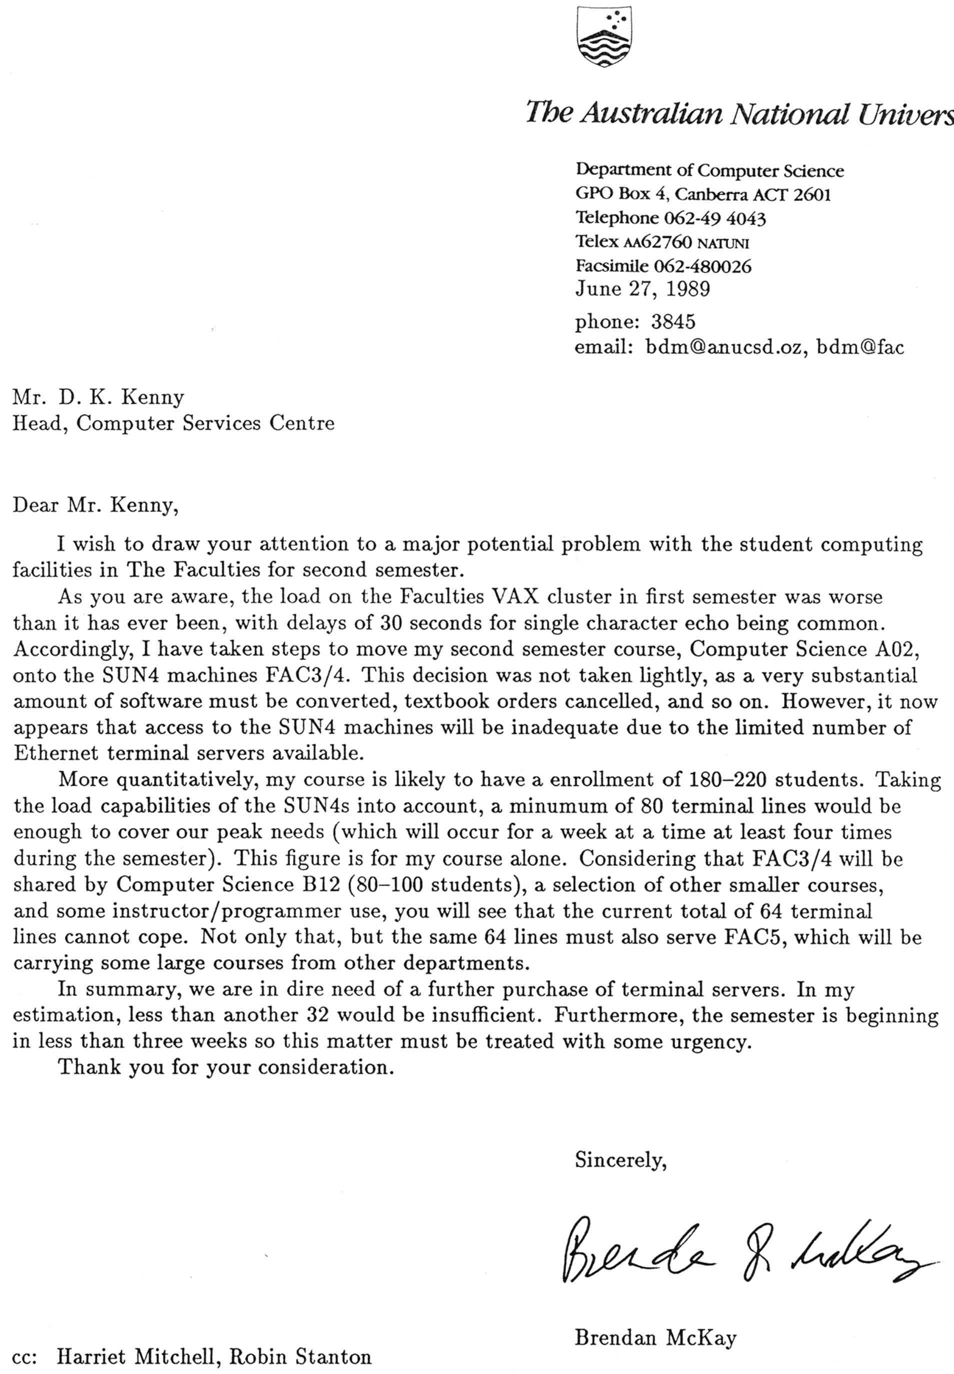

The Australian National University

Department of Computer Science
GPO Box 4, Canberra ACT 2601
Telephone 062-49 4043
Telex AA62760 NATUNI
Facsimile 062-480026
June 27, 1989

phone: 3845
email: bdm@anucsd.oz, bdm@fac

Mr. D. K. Kenny
Head, Computer Services Centre

Dear Mr. Kenny,

I wish to draw your attention to a major potential problem with the student computing facilities in The Faculties for second semester.

As you are aware, the load on the Faculties VAX cluster in first semester was worse than it has ever been, with delays of 30 seconds for single character echo being common. Accordingly, I have taken steps to move my second semester course, Computer Science A02, onto the SUN4 machines FAC3/4. This decision was not taken lightly, as a very substantial amount of software must be converted, textbook orders cancelled, and so on. However, it now appears that access to the SUN4 machines will be inadequate due to the limited number of Ethernet terminal servers available.

More quantitatively, my course is likely to have a enrollment of 180–220 students. Taking the load capabilities of the SUN4s into account, a minumum of 80 terminal lines would be enough to cover our peak needs (which will occur for a week at a time at least four times during the semester). This figure is for my course alone. Considering that FAC3/4 will be shared by Computer Science B12 (80–100 students), a selection of other smaller courses, and some instructor/programmer use, you will see that the current total of 64 terminal lines cannot cope. Not only that, but the same 64 lines must also serve FAC5, which will be carrying some large courses from other departments.

In summary, we are in dire need of a further purchase of terminal servers. In my estimation, less than another 32 would be insufficient. Furthermore, the semester is beginning in less than three weeks so this matter must be treated with some urgency.

Thank you for your consideration.

Sincerely,

Brendan D McKay

Brendan McKay

cc: Harriet Mitchell, Robin Stanton

1989: DCS complaint about response times on student computing facilities.

Geoff Huston reminds me that the G4 computer room flooded:

A flood in the Copland G4 computer room

It was a winter morning and when someone raised a floor tile of the false floor in the machine room they saw their head reflected in the water that had flooded in. I think it was one of the few times that the red emergency stop button was pushed deliberately (aside from the time that Brian Jaggers lent on it). The problem was then how to get the water out. The room was not designed to flood and it did not have a drain! I recall we managed to scrounge a wet vacuum cleaner and spent the rest of the day sucking the water out!

In my defence as designer of the room, I present the following:

> From the *Works Proposals for Copland G4 Computer Room, 7 Mar 1984*
>
> Sub-floor drainage -- A drain-hole and pipe through the North wall is required to provide a fail-safe escape route for water condensed or leaked from air-conditioning units in the machine room.

Up until this era, many computer rooms were equipped with carbon dioxide flooding for fire suppression. Unfortunately, because carbon dioxide forms a pool at the bottom of the room, it is possible for a person in the room to suffocate if they fall over. The G4 computer room was fitted with the safer Halon 1301 (now banned because of its ozone-depleting properties) system. Although the Halon reaches a concentration of only 7% of the atmosphere in the room, its large molecules interfere with combustion.

During commissioning of the computer room, a test of the flooding system was required. It was quite dramatic! First, the fire detection system was manually triggered. That initiated a series of audible evacuation alarms over half a minute or so, and then power was turned off to the computers and the airconditioning. Finally, the Halon was released, shrieking as it issued from the nozzles, and filling the room with the clouds of Venus within seconds. We were all invited to enter the room and smell the “healthier option” gas. It was very unpleasant; headaches ensued.

Sub-floor drainage is even more important in computer rooms fitted with water sprinklers. A couple of other computer rooms I designed were fitted with so-called dry-head sprinklers (instead of Halon). Normal sprinklers can flood a room if they happen to be hit, (by a ladder for example) but dry-head versions have an additional valve which remains shut until smoke is detected. It’s apparently OK for computer circuitry to be soaked in water, as long as the power is off, and not restored until the components are dried.

The computer room on Level 2 of the CS&IT building was fitted with a VESDA (Very Early Smoke Detection Alarm) designed to switch off power at the slightest indication of an overheated component. That didn’t work out so well when bushfire smoke penetrated the room.

Computer Aided Learning (CAL / CAI)

One of the science departments, Chemistry I think, used a computer aided learning tool which ran on microcomputers. A small laboratory of Hitachi Peach computer systems was set up to support this. It was managed by Gloria Robbins before she set up the MicroComputer Information Unit at the CSC. I have a vague memory that the CAL system ran on Toshiba optical disks.

Another department to use computer aided learning was Accounting in the Faculty of Economics. The people responsible were Roger Debreceny (now at the University of Hawaii) and Ron Bird (now at the University of Waikato). They used the GNOSIS CAL program, written in SIMULA 67 and running on the Faculties KL-10. One of the jobs necessitated by the move to VAXes was to get GNOSIS running on the VAXes. We obtained a version of Simula 67 for the VAX, but it was a portable version based on the S-machine. From memory it was two orders of magnitude slower in compiling and a factor of 10 slower in execution than the native DEC-10 version. It also lacked the very rich SIMULA 67 library available on the DEC-10. Consequently the conversion was non-trivial.

9.6 1980s: Evolution of the Faculties Computer Unit

After Geoff Huston left FCU in 1982, I was joined by Peter Hobson (now a Solutions Architect at QSuper). In a job interview he was once asked the standard question, “Where do you see yourself in five years time?” His answer was, he claimed, “In five years it will be Saturday, so I will be drinking beer!” Peter was replaced by Jeremy Webber who introduced me to the Unix command `less`, explaining that `less` is more or less `more` but actually more than `more`.

After I left the FCU in 1984, my role as Head was taken on by Geoff Huston, until approximately 1988, when he went to head the campus network group. Harriet Michell then took over until 1990, when she was seconded to the Australian National Dictionary Centre (ANDC). Harriet was a colourful, very energetic and lively person who had previously been known as Pat Jeffress. Very sadly she was killed by a dangerous driver near Nimmitabel in the mid 1990s. Melanie Rooney (then Melanie Bleeze) worked with Harriet between 1988 and December 1990. She remembers Harriet as one of the most inspiring people she has worked with. She recounts:

> One of the most rewarding things Harriet and I did in our time together at the FCU was to work with a then new undergraduate student in 1989 to implement software that helped that student convert required course readings/written text to the spoken word on her computer. That student went on to become the first totally blind student to graduate with honours in Law from the ANU. That's IT that makes a real difference in my mind!
>
> In 1990 the FCU implemented a visual impairment printing project and a Braille embossing printer was installed in one of the microcomputer laboratories.

After Harriet moved to ANDC she was briefly replaced by Monica Berko who then went on parental leave, and returned later to CNIP (Centre for Networked Information and Publishing) to work with Mark Corbould, and then followed Mark to join the National Library IT section. She was replaced by David Baldwin, a mad tri-athlete now at the Australian Institute of Sport, and finally John Tucker. At some point FCU was expanded in scope and renamed the Teaching and Learning Technology Support Unit (TLTSU).

David Baldwin was a student in DCS from 1984 to 1986. He recalls that on the last day of operation of the DEC KL-10, students were permitted an all-night session playing `vttrek` on 9600 baud terminals. [I must have agreed to that, but I don't remember it.[5]] David remembers some significant innovations during his tenure as Head, TLTSU:

- The demise of terminal rooms.
- The expansion of student laboratories fitted with X-terminals (initially LabTam, later Tektronix), Macs, and Windows PCs.
- The advent of Student Consultants, who sat in the library computer labs and assisted students. Duncan McIntyre (now a First Assistant Secretary in the Department of Industry, Science, Energy and Resources) and Jamie Norton (now Chief Information Security Officer at the ATO) served in this role.
- The expansion of computer labs in the halls of residence, as a campus safety measure (to avoid individuals walking around the campus late at night). These labs were partly populated with superseded, but still viable, equipment from the main TLTSU labs.
- The introduction of a unified student card, with a swipe stripe, photo and barcode, giving access to computer rooms, buildings and the library. David stood up an LDAP identity server in connection with this.
- The LDAP server was also used to implement quotas on internet downloads. This could be done because all Web access was via ANU's proxy server.
- The development of smart lecture theatres with video projection (initially 3-gun, later LCD technology), PA systems, facilities for lecture recording etc.

Mention of video projection reminds me of the rash of thefts of video projectors from ANU lecture theatres in the late 1990s. Student labs were also targets of theft. One idiot stole an external drive from a Macintosh by cutting its cable. Good luck with finding the appropriate connector and soldering it to the very thick cable! A mass theft later occurred in which thieves used a van or ute to tear the security grille off a window to an Engineering Department computer lab and made off with a whole lot of computers. Arguably, theft of equipment was less serious than a concerted campaign of malevolent hacking by one student, let's call him X.

[5] Geoff Huston says that not only did I authorise it, but I wrote a program to create havoc by randomly reassigning jobs to different terminals!

X ran password crackers, took over accounts, installed root kits and totally disrupted the operation of school computers around campus, causing a massive waste of staff time, and loss of productivity. Many machines had to be re-installed from scratch to eliminate the malware. In gathering evidence against X, ANU staff and police monitored dial-in terminal lines and security cameras. On one occasion, X was the only user in one half of the Copland G5 terminal room. A real-time log of his keystrokes showed that he was engaged in nefarious activity. To make the evidence absolutely watertight, ANU staff climbed up on the desks in the other half of G5 to look through the glass divider in order to be able to provide eye witness evidence.

In one session X (a caucasian male) was logged in as a female Asian student and was writing an email requesting increased quotas and privileges in order to be able to complete assignments, when he realised that he didn't know how to spell the name of the student whose account he was operating. He suspended the email, typed the command to print the details of the account, then resumed the email and signed his victim's name. All this activity was logged.

A totally compelling (to my eyes) brief of evidence was prepared for a prosecution but it was not proceeded with, because in order to try to halt the damage, Robin Erskine had called X in (with his mother) and ordered him to cease and desist. Apparently this was a "punishment" and he couldn't be punished twice for the same crime.

9.7 Laboratories of Student Computers

The first laboratory of personal computers for student use was the original MacLab in the Chifley Library. Set up in 1985, it comprised 11 Macintoshes, two dot-matrix printers and an Apple LaserWriter. Apple donated half of the equipment.

IIn 1986, a second micro laboratory with five IBM compatibles and five Macintoshes plus printers was set up in the Hancock Library. Within two years, the IBM compatibles were replaced with Macintoshes, and a charging system was introduced for laser printing. A document I wrote in 1988 reports that both Statistics and Commerce departments had set up their own departmental microcomputing laboratories.

Brian Molinari was instrumental in changing the way in which student computing was funded and operated around 1989. Brian explains:

> It is convenient to think of 1989 as a pivotal year in the provision of computing resources in The Faculties.
>
> The key committee was FCAC (Faculties Computing Advisory Committee). Each year it put together a bid for computing resources (equipment and support). This bid was taken by the Director, Computer Services Centre, as part of his overall bid to the central Computing Equipment Budgetary Committee. The bid also went to the Faculties Resources Committee. In a tradeoff the latter funded much of the support while the former funded much of the equipment.
>
> In 1989 the process changed. The Faculties would make the bid directly to CEBC, rather than being part of the Computer Services Centre bid. A strategies document, addressing three to five years into the future, had to accompany the first bid. The resulting document, *Strategic Plan for Computing in The Faculties: 1990-1992*, provided a detailed review of the situation in 1989, and spelled out a path ahead.
>
> In 1989:
>
> - There were something like 240 workplaces available to students, the majority (some 75%) having been provided centrally through FCAC. Half of the workplaces were terminals connected to multi-user systems (Vax, Sun). Almost all were provided in bookable laboratories (with less than 20 places per laboratory), rather then the 50-place horror of the original terminal room.
> - There were about 400 academic staff and 300 general staff in The Faculties. It was estimated that some 130 had computers/terminals on their desks. On the other hand, of the 19 buildings associated with Faculties departments, only five were connected to the campus backbone network and only two had an internal LAN connected to the network. Almost all of these desktop machines were thus stand-alone, and did not provide access to email.
>
> The report presented a strategy with the following elements (amongst others):

- The responsibility for the student computing environment fell to The Faculties, not to individual departments.
- Resources for the environment would be provided on a needs basis (see next panel).
- Staff and students would be provided with a networked environment. Note that widespread use of the Web was still some years away.

The strategy was adopted, and largely implemented over time. I am unaware that a subsequent strategy document was written. On the other hand, a decade later the University established the Division of Information which brought almost all information technology on campus into a uniform planning and management framework.

Brian Molinari : And the currency was “bookable laboratories.”

The resources needed for the year ahead were tallied unit-by-unit, noting the student enrolment, and reduced to the equivalent number of ”bookable laboratories” of size no greater than 20 places. (A demonstrator could react with a group of that size.) In 1990 the need was for 16 such laboratories, the expected need in 1992 was for 26.

This system gained credibility. Predictions for the year ahead were based on student enrolment, and hence quite precise. Very high laboratory utilisation could be demonstrated. Student satisfaction was noted.

The mechanism carried an unstated threat to the Faculties: “If the indicated needs are not substantially met, then the teaching enterprise will be compromised.” There was a sweetener, however, in that it guaranteed that the software packages in the environment were properly licensed. Being found using unlicensed software had just emerged as a major risk to an institution’s reputation, and the mechanism managed that risk.

The initial discussion with department X would go as follows:

X: Why should I fill this in?
FCAC: Do any of your units next year require student computing?
X: Yes, but we have only three PC’s available!
FCAC: Fill this in, and space in a bookable laboratory will be provided.
X: (suspiciously) Will it cost me anything?
FCAC: Not directly. And we will make sure it runs.
X: Where do I sign?

Before X came off their initial high, Murray Napier (FCAC) would extract any underutilised space in the department for the next laboratory.

DCS was a major beneficiary, in that the financing of the student computing environment was not a call on the department budget. Initially we were the biggest user in the environment, though that declined over time as student use of IT became ubiquitous.

In practice, the support of the DCS-related laboratories was delegated to the DCS programmer group. When we moved to the CSIT Building in 1995, this mechanism fitted out five 20-place laboratories. These were, I think, the first purpose-built student laboratories on campus. In principle these laboratories were bookable by other users, but in practise this rarely happened.

The ANU Reporter article reproduced on Page 234 celebrates the advent of a shared 12 person lab, “the Tennis Court”, in which the Sun workstations were named after tennis players, and the powerful server, which operated over the net, was called Boris (Becker). I believe that I may have been responsible for the naming. The three-year Sun research fellowship mentioned in the article employed Raju Karia, who worked on scientific visualisation, starting in March 1993.

As foreshadowed in the ANU Reporter article, the tennis player lab was expanded to two labs of 18 workstations each. To improve out-of-hours access, all 36 workstations were relocated during a two-week teaching break to the old cartography labs in the basement of the Dedman building (now demolished). In the absence of airconditioning, Boris was placed in a bathroom, cooled by contact with the earth behind and beneath the room. Boris ran reliably at a constant temperature of around 35C.

The tennis player workstations were all Sun 3/80s based on Motorola 68030 CPUs. They contained 104MB local disk drives which used a poorly-chosen spindle lubricant. It solidified when cold. When power was lost to a lab, and then restored, only half of the machines would restart. A gentle tap sometimes worked, but a hairdryer was the ultimate remedy. Fortunately, power outages in Canberra are quite rare.

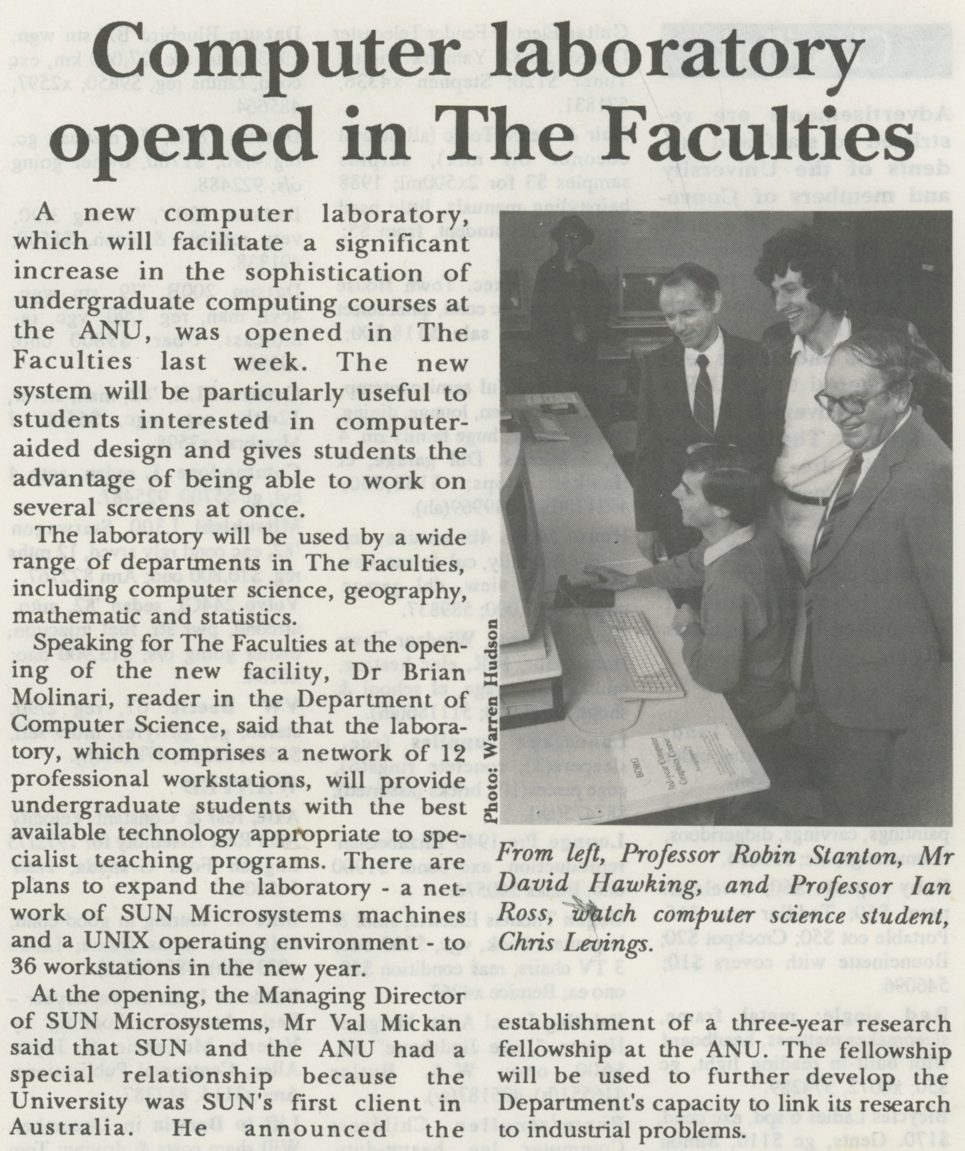

Computer laboratory opened in The Faculties

A new computer laboratory, which will facilitate a significant increase in the sophistication of undergraduate computing courses at the ANU, was opened in The Faculties last week. The new system will be particularly useful to students interested in computer-aided design and gives students the advantage of being able to work on several screens at once.

The laboratory will be used by a wide range of departments in The Faculties, including computer science, geography, mathematic and statistics.

Speaking for The Faculties at the opening of the new facility, Dr Brian Molinari, reader in the Department of Computer Science, said that the laboratory, which comprises a network of 12 professional workstations, will provide undergraduate students with the best available technology appropriate to specialist teaching programs. There are plans to expand the laboratory - a network of SUN Microsystems machines and a UNIX operating environment - to 36 workstations in the new year.

At the opening, the Managing Director of SUN Microsystems, Mr Val Mickan said that SUN and the ANU had a special relationship because the University was SUN's first client in Australia. He announced the establishment of a three-year research fellowship at the ANU. The fellowship will be used to further develop the Department's capacity to link its research to industrial problems.

Photo: Warren Hudson

From left, Professor Robin Stanton, Mr David Hawking, and Professor Ian Ross, watch computer science student, Chris Levings.

1989: ANU Reporter article recording the opening of the first shared Faculties student workstation (as opposed to terminal or PC) laboratory. Workstations were named after tennis players and the server was called Boris, after Boris Becker who was a powerful server. The lab was initially located on the top floor of the Crawford Bldg.

Because the workstation labs were at least nominally general access, the CSC budgeted for their maintenance. One year, I noticed that the amount that the CSC had budgeted to pay Sun for maintenance on the 36 workstations was equivalent to the cost of purchasing quite a large number of Sun SLCs. SLCs (Sexy Little Computers) were SPARC based and outperformed the 3/80s. Accordingly, I proposed that CSC contract with DCS, rather than Sun, to provide maintenance on the student workstations. We would use the money to buy SLCs and use them to replace Sun 3/80s on academic desks. Some of the thus-freed 3/80s were passed on to graduate students on a "while not otherwise required" basis. We undertook to keep one of the freed 3/80 as a hot spare and another one as a source of spare parts, and to provide more as needed. If the 3/80s turned out to be reliable (which they did), then the graduate students would be much better off.

This scheme was so obviously beneficial to ANU that it went ahead for two years, and all academics ended up with SPARCstations.

From the point of view of Sun's field service division, it was not at all beneficial. Our local field service technician, Tommy Ivansen, was called up to Sydney to role play with his boss. Tommy played the part of ANU, and was asked to respond to increasingly attractive maintenance offers from his boss. Once the boss had gone as far as he was prepared to, and Tommy said that it was still in ANU's interest to self-maintain, the boss said, "You're a bastard!" Tommy said, "Yes, they are!"

As an aside, maintenance contracts were sometimes highly lucrative for vendors, because business managers tended to carry forward last year's budget item for computer maintenance without too much question. Several years after the last DEC computer was purchased by ANU, Robin Erskine found that ANU was still paying a total of over half a million dollars a year in maintenance to DEC. In several cases the payments were not justifiable, because the equipment was no longer relied upon, in some cases not even switched on.

DCS's move to the CS&IT building in February, 1995, saw the creation of several new 24-hour access teaching laboratories on level one. Each of the labs could accommodate 20 working students, but economics dictated that the workstations were LabTam X-terminals supported from very powerful SPARC servers named after conductors: Karajan, Iwaki, Bohm, etc. When an even more powerful server was acquired later on, Bob Edwards named it "Yttrium" because, as every physicist knows, yttrium is a super conductor.

9.8 The Situation in 2021

In 2021 there are a large number of Windows and Macintosh laboratories across the campus operated by IT services and available for teaching and student use. The ensemble is known as the *Information Commons*. To quote the relevant ANU web page:[6]

> The Information Commons are a collection of computer labs, Group Study Rooms and Teaching and Learning spaces located within ANU Library locations, residences, and buildings across campus.
>
> There are over 1,500 computers located within the Information Commons. These computers provide students and staff with access to Internet, wireless, software, and are connected to a shared printing system.
>
> ...
>
> ANU has launched at-home access to the software applications students rely on for their studies by establishing an ANU Virtual Information Commons Remote Desktop solution. This solution is now ready for student use.

The situation in CECS is somewhat different. Bob Edwards tells me that there are 270 all-in-one[7] Linux machines in student laboratories across the college. There are another 32 such machines kept in a store room in the (new) Hanna Neumann Building which can form a pop-up laboratory in the CS&IT Seminar Room when needed. (See the photo on Page 236.)

Software costs for Linux labs are very low compared to others because of the use of open source software.

All these machines are set up to use PXE boot. That is, when booted, they automatically restore a standard Linux image from a server and a VDI (Virtual Desktop Infrastructure) desktop. Bob Edwards tells me that he creates an image and ships it to an organisation in Perth who test it before propagating it to ANU's image servers.

Student home directories are mounted from a redundant fileserver. The file server and some image servers are located in a small computer room on level two of the CS&IT building. Some

[6] https://services.anu.edu.au/information-technology/software-systems/information-commons

[7] An all-in-one system has only a keyboard and screen; the "computer" is inside the monitor.

other image servers are actually virtual machines in ANU's VMWare cloud in the Leonard Huxley Building.

This system is amazingly flexible. The School of Engineering has two labs of student machines running Windows. If needed, for example to run an exam, these machines can be configured to run the standard Linux images. Further, when ANU students were banished from campus during the COVID-19 pandemic in 2020, they were able to install the standard image on their own machines wherever they happened to be.

The 32 (yes, 32, they're two deep) Linux machines for a pop-up lab, in a crowded store room in the Hanna Neumann Building. *Photo: David Hawking*

The small server room on Level 2 of the CS&IT Building which houses the file servers and some image servers for the student labs, as well as some other production servers. *Photo: David Hawking*

Chapter 10

ANU: Administrative Computing

There seem to have been three fairly distinct eras in the history of ANU administrative computing:

Era	Approx. years	Hardware	Software	Data acquisition
First	1965 – 1980	IBM m'frame	Assembler, PL/1	Paperwork, data punching
Second	1980 – 1998	Fujitsu m'frame	Adabas/Natural	Administrative FNA terminals
Third	1998 –	Windows, VMs	PeopleSoft + web	web interfaces

10.1 First Era

As reported by Lorrel Sherar (see Page 272) ANU administration used "Unit Record Equipment", i.e. card sorters, collaters, and tabulators prior to the arrival of the IBM 360/50.

Pat Wilson, Head of the Management Systems Group, in 1975. *Photo: Gab Carpay.*
`https://openresearch-repository.anu.edu.au/handle/1885/224454`

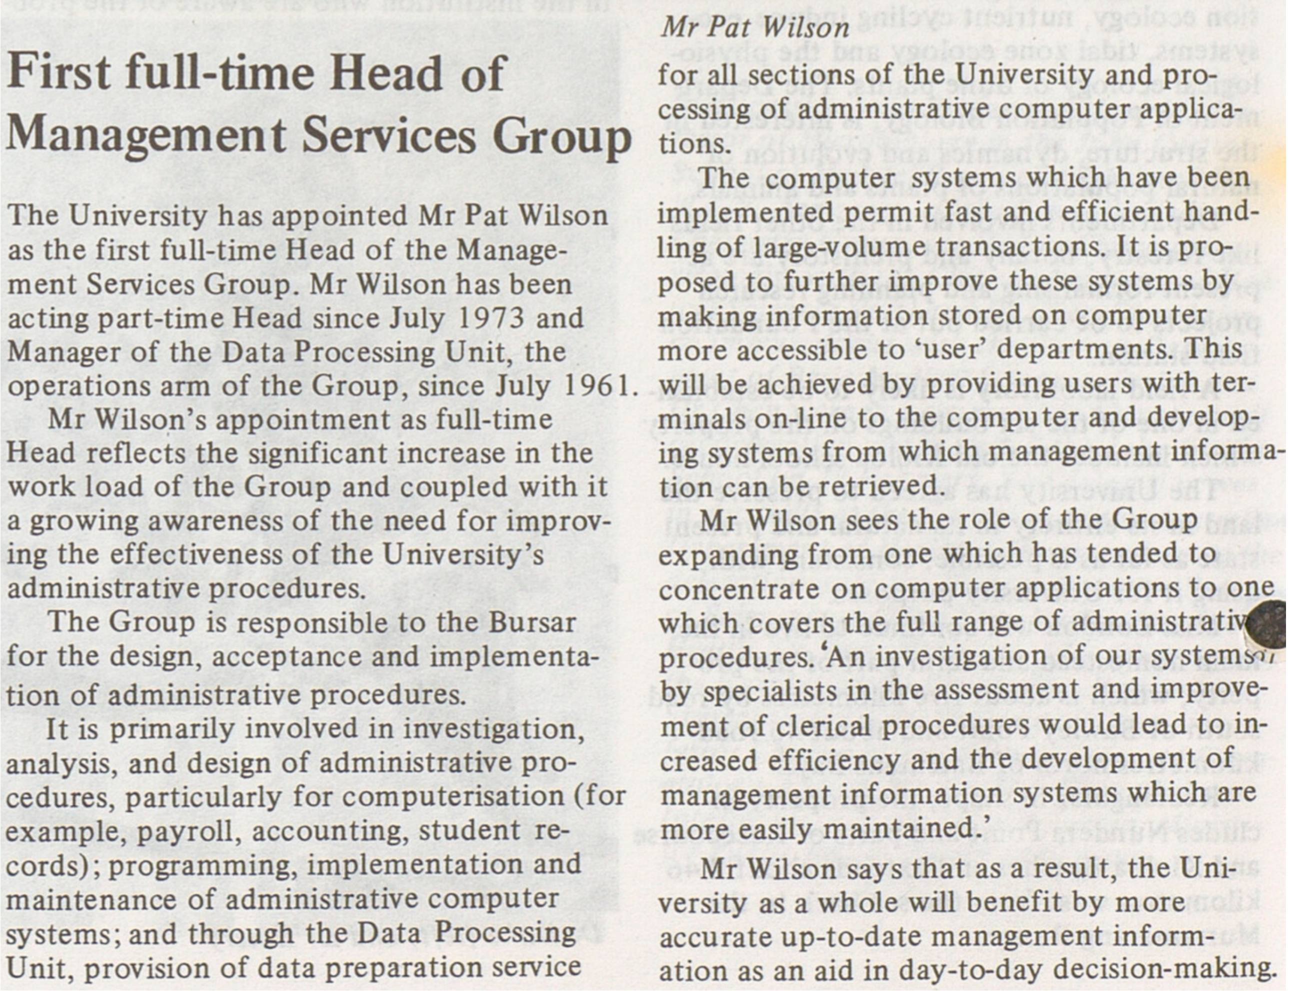

First full-time Head of Management Services Group

The University has appointed Mr Pat Wilson as the first full-time Head of the Management Services Group. Mr Wilson has been acting part-time Head since July 1973 and Manager of the Data Processing Unit, the operations arm of the Group, since July 1961.

Mr Wilson's appointment as full-time Head reflects the significant increase in the work load of the Group and coupled with it a growing awareness of the need for improving the effectiveness of the University's administrative procedures.

The Group is responsible to the Bursar for the design, acceptance and implementation of administrative procedures.

It is primarily involved in investigation, analysis, and design of administrative procedures, particularly for computerisation (for example, payroll, accounting, student records); programming, implementation and maintenance of administrative computer systems; and through the Data Processing Unit, provision of data preparation service for all sections of the University and processing of administrative computer applications.

Mr Pat Wilson

The computer systems which have been implemented permit fast and efficient handling of large-volume transactions. It is proposed to further improve these systems by making information stored on computer more accessible to 'user' departments. This will be achieved by providing users with terminals on-line to the computer and developing systems from which management information can be retrieved.

Mr Wilson sees the role of the Group expanding from one which has tended to concentrate on computer applications to one which covers the full range of administrative procedures. 'An investigation of our systems by specialists in the assessment and improvement of clerical procedures would lead to increased efficiency and the development of management information systems which are more easily maintained.'

Mr Wilson says that as a result, the University as a whole will benefit by more accurate up-to-date management information as an aid in day-to-day decision-making.

Pat Wilson becomes full-time Head, Management Services Group. *ANU Reporter, 14 March 1975.*
https://openresearch-repository.anu.edu.au/handle/1885/141596

The acquisition of the IBM 360/50 in May,1966[1] allowed ANU to commence the process of computerising some of its administrative processes, such as payroll and accounting. Most of the important administrative applications, like the payroll, were written in IBM 360 Assembler, though there was a move toward use of PL/1.

I believe the data needed for electronic processing would have been coded on coding sheets and sent for punching to cards by the Data Processing Unit (DPU). DPU employed a considerable number of data punch operators (all female) located in F Block in the old administrative area. I understand that DPU staff were responsible for operating the 360/50 at certain times of the day, and for running the administrative jobs such as printing the payroll.

Pat Wilson headed both DPU and the Management Services Group (MSG), which was located in an old cottage on Balmain Lane, next to the student counselling facility which MSG staff called the "stress cottage." When Barbara Bankovsky joined MSG in 1975 as an assistant programmer, she was struck not only by the trees in the back yard which produced peaches and enormous figs, but by how cold the house was in winter until the boiler was fired up and got up to temperature. Programmers gave her PL/1 coding sheets which she punched onto cards and sent by internal mail for processing on the 360/50 – with a four-hour turnaround.

When Pat Wilson left ANU, Skaidrite Darius, DPU Supervisor, took on more responsibility. According to Barbara Bankovsky she used to backup the IBM 360 every evening and took the tapes home. See Page 274 for more of Skaidrite's remarkable story.

When the Univac 1108 was acquired in 1972, there was a project to migrate the administrative

[1] See the photo of an IBM 360 system on Page 34.

systems to Univac COBOL, but this wasn't completed. MSG programmers were sent to a five-day COBOL course in Canberra and Barbara's request to join them was accepted. She was promoted to Programmer, and wrote programs to generate reports for Bob Jeremy who was in charge of Student Systems. Barbara remembers that, around this time, ANU bought an Inforex 5000 system at large cost, to take load from the mainframe and to simplify administrative processes. It was used to run Sundry Debtors (Accounts Receivable), Phone Directory and a number of other small systems. Barbara was in charge of the Inforex and delighted in being able to enter transactions from a monitor rather than from a card deck. The motivation for acquiring the Inforex is easily inferred from the relevant Wikipedia article.[2]

Because some administrative applications were still running on the IBM 360/50, ANU passed up an opportunity to sell it second-hand in the Middle East and kept it running until it failed. After that an IBM 370/138 was leased until the acquisition of a FACOM (Fujitsu Automatic Computers).

On 28 August 1978, Bob Landford became Director, Computer Services Centre (CSC), having been Systems Engineering Manager (Canberra) for IBM Australia. Soon after, on 26 October 1978, the Management Services Group, headed by Ian MacNicol, were incorporated into the CSC and moved into the Huxley Building. Later, MSG moved to a new demountable building, "K Block" near the Law School. DPU remained in the old administrative building (F block?)

Extract from Computer Centre Annual Report for 1966, reporting the arrival of the IBM 360/50

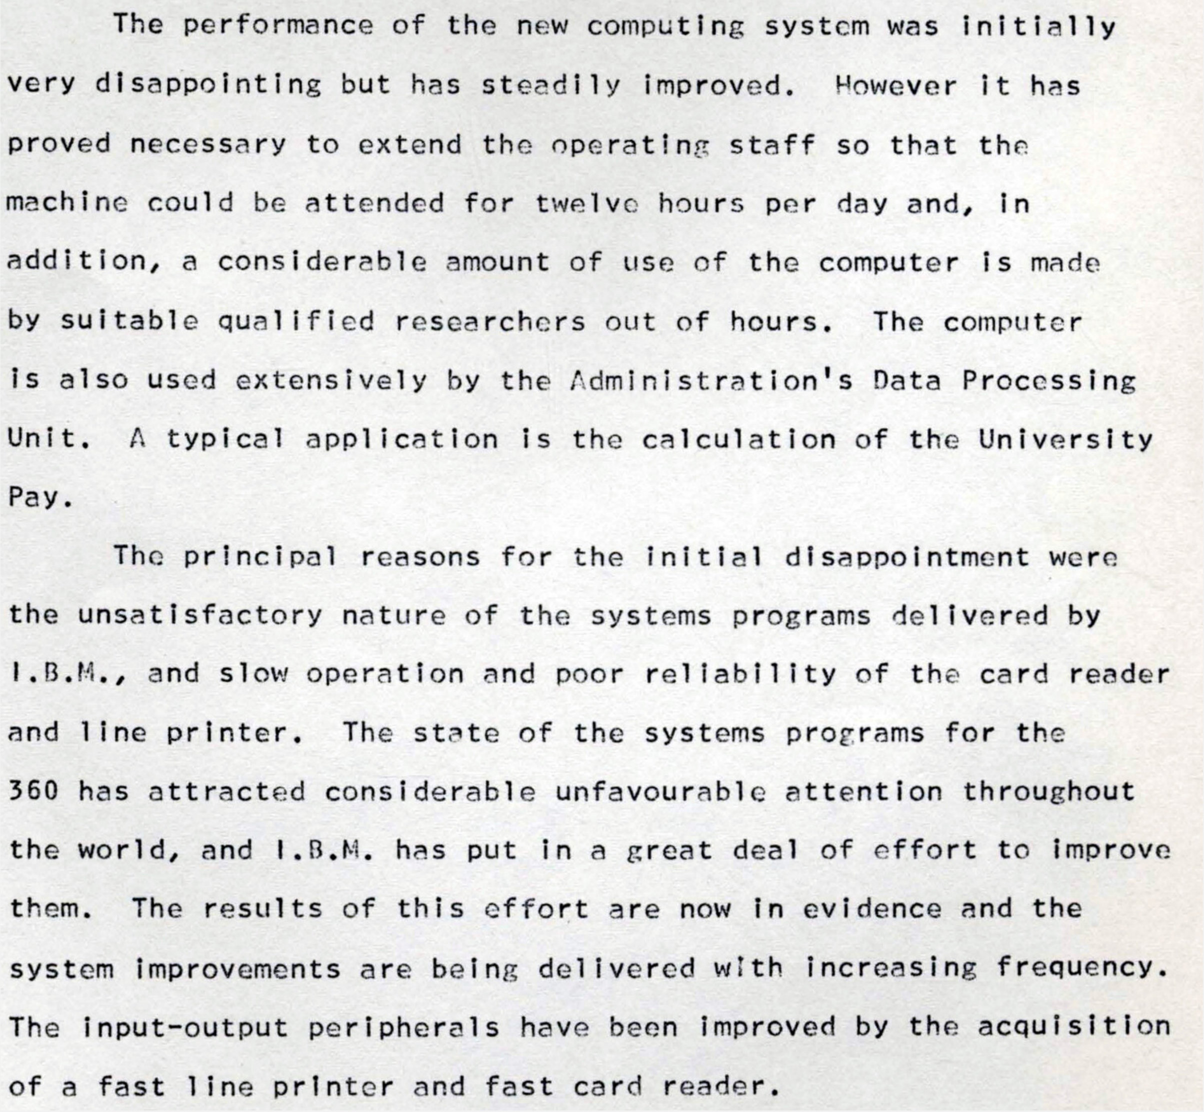

The performance of the new computing system was initially very disappointing but has steadily improved. However it has proved necessary to extend the operating staff so that the machine could be attended for twelve hours per day and, in addition, a considerable amount of use of the computer is made by suitable qualified researchers out of hours. The computer is also used extensively by the Administration's Data Processing Unit. A typical application is the calculation of the University Pay.

The principal reasons for the initial disappointment were the unsatisfactory nature of the systems programs delivered by I.B.M., and slow operation and poor reliability of the card reader and line printer. The state of the systems programs for the 360 has attracted considerable unfavourable attention throughout the world, and I.B.M. has put in a great deal of effort to improve them. The results of this effort are now in evidence and the system improvements are being delivered with increasing frequency. The input-output peripherals have been improved by the acquisition of a fast line printer and fast card reader.

In an odd little twist, an ANU graduate and long-term IBM employee, John Dravnieks, became

[2] https://en.wikipedia.org/wiki/Inforex_1300_Systems

responsible for maintaining IBM's assembler from 1990 – 2006. A panel on Page 242 records why the assembler for a 40 year old machine still required and received maintenance and development.

10.2 Second Era

The Computer Services Centre Annual Report for 1980[3] documents the acquisition of a FACOM M160F system to handle the ANU's administrative computing load, replacing the leased IBM 370/138. At the same time, a project was embarked upon to rewrite key administrative systems using the `Adabas` database and the `Natural` fourth-generation language. Unfortunately, when first run, the payroll system, which had taken half-an-hour to run on the much slower IBM 360/50, took more than 24 hours on the new FACOM. George Kacer, who had been the maintainer of the IBM assembler program, was brought out of retirement to run the assembler version on the new FACOM, until the performance of the Adabas/Natural version was improved.

Barbara Bankovsky told me that George Kacer was living in Czechoslovakia at the end of WWII. In the short gap between the departure of the Nazis and the arrival of the Russians, George and his wife walked over the mountains to the West, George carrying his (drugged) daughter on his back. He spent time in internment camps prior to arrival in Australia. George used to sell hams to his colleagues at Christmas. One year the money from the sales was stolen from his office, and his colleagues took up a collection to replace it. In his spare time, George was a professional photographer.

Adabas was a non-relational, inverted list database management system (DBMS) sold by Software AG since 1971. Natural was a proprietary "Fourth generation language" (4GL) from the same vendor. It was first introduced in 1979.

A new Fujitsu Network Architecture (FNA, modeled on IBM's SNA) network based on coaxial cable was developed and extended during this era. It supported Fujitsu's equivalent of IBM 3270 block-oriented terminals. Eventually all departmental administrators had one of these. Applications developed in Natural would display a series of interactive screens on these administrative terminals.

Having a completely separate administrative network based and special terminals helped isolate distributed administrative functions from security risks on the main network. See Page 266 for information on when and how ANU networks were eventually consolidated into one.

A report in 1982 saw no role for central computing facilities in providing text processing and formatting. Departmental administrators in The Faculties typically used Wordplex word processors, which didn't communicate with the computers used by academics, or with the FNA terminals.

After the consolidation of networks, it was possible to emulate the FNA terminals on the PCs then used for word processing and other office tasks, and to allow some distributed processing of administrative data. As David Berriman recalls:

> MIS utilised a product called Entire Connection (a 'smart' 3270 terminal emulator) to replace the multitude of stand-alone mainframe VDUs (Visual Display Units, i.e. terminals) with PC desktops. My team (Desktop Applications) developed scripts that allowed the PCs to transfer data from the mainframe directly to the desktop for use in applications like MS Excel and Word — a feature that until then was only possible by overnight download requests run centrally. We also developed VBA macros that automatically formatted the downloaded data for reporting requirements, or to utilise point-of-sale (POS) devices (cashdrawers and cheque readers) for distributed payment collection. This approach won several accolades and awards at Software AG conferences.
>
> One specific use of this technology was for the automated generation of the ANU testamur[4]. Up until then this was laboriously done by having an old stencil machine produce blank testamurs for every degree and then employing a calligrapher to manually write in the names and dates. This was then passed onto the senior executives (VC and Registrar) who had to hand sign each one. This took months of planning and organisation, and if any errors (or smudges) occurred it took days to get replacements. If, during the registration for a graduation ceremony, an error was

[3] `https://openresearch-repository.anu.edu.au/bitstream/1885/149187/24/2.1.8.18%281%29-3-1980.pdf`

[4] Degree certificate.

discovered, then typically the student was presented on stage with an empty testamur tube. (A corrected version was then produced over the coming weeks for the student to pick up).

The new automated system (which for the first time used digital signatures) took only days to produce all the testamurs needed for a graduation ceremony and minutes to fix any issues. For the very first ceremony which used this new technology, it was discovered that one student was going to be presented with the wrong degree (an administrative error in their faculty office). We managed to reprint the correct testamur and place it in the tube in time for the student to receive. (It took 20 mins including the time it takes to run between Melville Hall and Llewellyn Hall).

The new technology was not without controversy and a letter to the ANU Reporter lamented the use of automated technology to give students a soulless A4 sheet of paper (an approach used by most other universities at the time). ANU still used the large goatskin parchment – finding a printer that could handle it was one of the biggest challenges of the day. I invited the complainant to come see the final results for himself (and compare the old and new testamurs). He was so won over (the new ones looked so much better) that he wrote his own rebuttal in the next Reporter!

The administrative FACOM system was successively upgraded to a Fujitsu M760. After the upgrade to the M760, the superseded M360 became the front-end for the Fujitsu VP100 supercomputer. Finally, in 1995, ANU acquired its last mainframe computer, a Fujitsu GS8400 which was in service until 2006.

Rick Van Haeften's recollection of marking the demise of ANU's last mainframe.

I decided to do something to mark the occasion and settled on getting a component from the old machine, having it put in a casket and burying the casket in the front garden area of the Leonard Huxley building. With the help of the folks in Facilities, we got hold of a Wollemi pine sapling and planted it on top of the 'grave'. I had a plaque set into a concrete plinth with some words of remembrance. We even had a small gathering and invited some influential people. To this day, I do not know what was in the casket. I hope the plaque is still there.

The concrete plinth is still there but unfortunately, the Wollemi pine seems to have died.

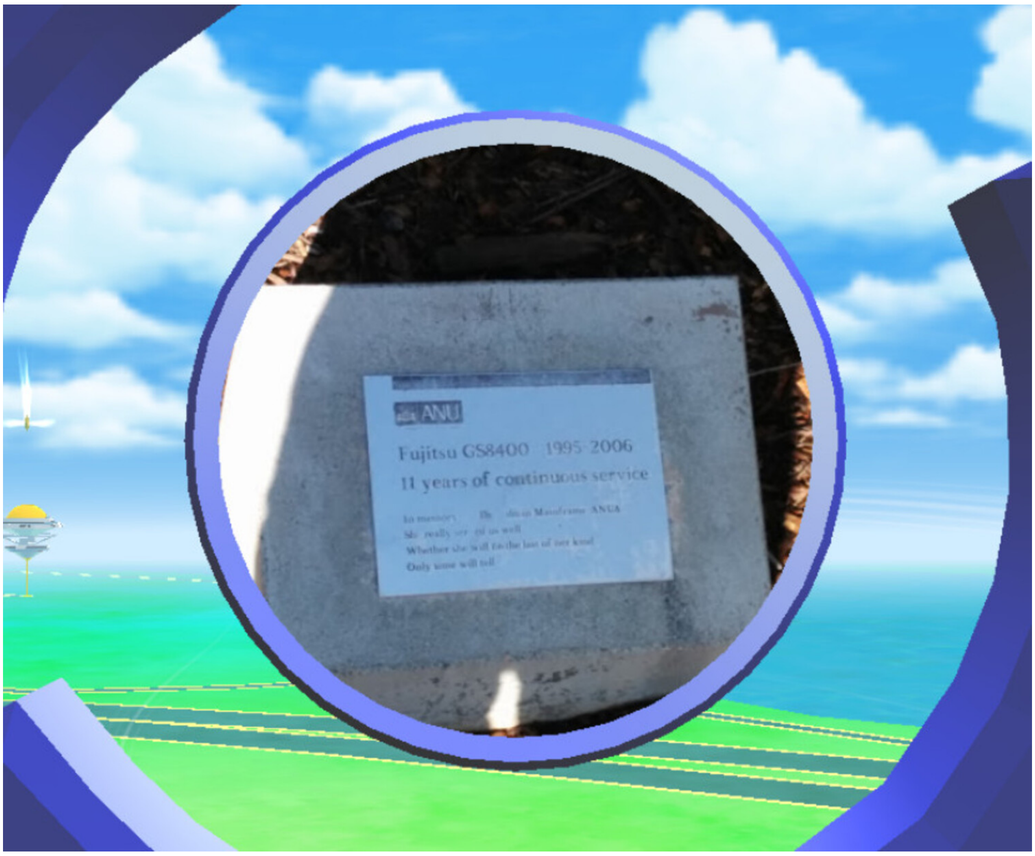

A photo of the memorial plaque, before the wording wore off. *Photo: Rika Kobayashi via Pokémon Go.*

Epitaph for ANU's last mainframe

Fujitsu GS8400 1995 – 2006
11 years of continuous service

In memory of the Admin Mainframe ANUA
She really served us well.
Whether she will be the last of her kind.
Only time will tell.

John Dravnieks's work for IBM on the 360 assembler, 1990 – 2006

John: When we started on in it was known as Assembler G and had been in a moribund state for years. We created High Level Assembler (HLASM) R1 versions 1 to 6 over the years – I fear it has now dropped back into care and maintenance.

Dave: What prompted development?

John: The assembler group of SHARE (IBM users group) had a list of requirements, including that mods and extensions developed by the Stanford Linear Accelerator Centre (SLAC) should be incorporated in the product. (The source of the old assembler was made available to users.) And there may have been pressure from the hardware group to support new instructions (this area had fallen by the wayside as part of the "No one writes in assembler" mantra.) IBM at the time had four operating systems – all with different assemblers – one aim of HLASM was to provide the same program to them all.

Dave: Who writes new code in IBM assembler?

John: Not many users these days – the main writers are the software companies like IBM, 21st Century Software, Phoenix Software, Rocket to name a few. As a rule they create what is called middleware – a lot of the utility programs get written in assembler. Though nowadays the IBM C compiler allows a `__asm` directive and you can write hardware instructions directly. Programs written in C/C++ rely on a runtime environment to operate though.

There is also a product called Metal C which has no runtime dependency and allows for system services to be called directly – it is intended for writing operating system code.

IBM for many years has used an internal language called PL/X – it creates assembler code which then has to be assembled (instead of creating the object text directly as most compilers do – I think Metal C does this as well. I do know that PL/X were exploiting many of the changes we put into the assembler.

Dave: What sort of capabilities were added to make "High-Level" assembler?

John: I will relate a story about this that someone else used when asked this – Way back in the 1970s, an American airline had a fleet of 727s and their slogan was "whisper quiet jets" (or something like that). When someone asked what was different about their 727s compared to the ones flown by others, what extra feature did they have, an executive from the company quietly replied, "Marketing."

More to the point, we incorporated all the SLAC mods, and added lots of little things like allowing mixed case in the source. (Before then it had to all be in UPPER case). We also added in exit points for the I/O and that allowed users to improve the product even more; the I/O was changed from BSAM to QSAM (Basic versus Queued Sequential Access Methods). And we changed how new instructions were added, making it almost clerical.

Dave: Does IBM still sell machines with the same instruction set as a 360?

John: The short answer is yes. IBM has a policy of backwards compatibility for user programs – a user program is one that uses standard operating system interfaces and only operates in what is called *problem* program state (as opposed to *supervisor* state – there are many instructions that require supervisor state to work – things that the operating system would use. Any problem state program written for a 360 would operate correctly today on the latest hardware.

Dave: Can the assembler target different architectures?

John: Not directly. There is however a very powerful macro language in the assembler – it does a lot more than just a copy and symbolic substitution – and over the years people have used this to create cross assemblers. The macro language allowed the creation of a set of structured programming macros for the assembler – IF/ELSE/ENDIF, DO loops etc.

10.3 Third Era

The initial move into the third Era came because of the perceived need to address Y2K issues and to be able to handle GST calculations within the finance systems. The old Adabas systems stored years as only two digits. The mainframe system was also nearing end of life.

After a process led by Fay Gibbons, (Head, MISD) PeopleSoft, then sold by PeopleSoft, Inc, prior to its takeover in 2005 by Oracle, was selected. In its favour was that it was the only option that had an integrated HR & Students systems (through the Campus Community package)[5] According to Stuart Kendall:

> PeopleSoft also provided a cheaper education costing model. Having one infrastructure and development platform was seen as a positive. PeopleSoft also provided a development platform to adjust and add to the delivered product which was required to fit ANU needs.

PeopleSoft provided a Windows front-end and maintained its data in a Microsoft SQL database. The ISIS and HORUS systems (See Page 244) were built to access the database directly.

Stuart Kendall worked on the implementation of PeopleSoft Finance for January 1998 and an upgrade to allow GST for June 1998. He later led the PeopleSoft IT team from 2006 to 2018. I asked Stuart how much human effort was required to implement the new system:

> Quite a lot. Joint team of functional experts (5) and developers (4) plus a team building infrastructure (3). Anderson Consulting also assisted.

After the initial implementation of PeopleSoft, the mainframe was run to allow historical data to be accessed until the machine was decommissioned, but there was no parallel operation. Stuart says:

> Can't even remember if we converted any data. At the time only current and previous years were available in the old system

Yvonne Kulesza worked on the HR (2000) and Student Administration (SA, 2002) implementations. A check digit was added to employee numbers during conversion and student IDs were merged or re-allocated. Old and new systems were not run in parallel, but there were many attempts to convert old data to new. Yvonne recalls:

> PeopleSoft suite offered new, modern applications, and was open to modifications and additions. Initially ANU was keen to keep everything as it was on the mainframe. That resulted in a lot of modifications to the delivered product. It was a mistake because after every PeopleSoft patch we had to re-apply ANU specific modifications. Later we either used delivered functionality with special ANU settings or, if impossible, we would develop ANU module beside the main product, so patches would not affect our work.
>
> HR (Human Resources) and SA (Student Administration), though separate applications, shared the People (employees and students) database. Initially both applications were placed on the same database, so updates were instantly 'visible' for both. Later PeopleSoft (Oracle) separated them, so a special arrangement had to be developed to keep both databases in sync. For HR it was a very flexible arrangement, which allowed multiple concurrent (part time/casual) jobs for employees as well as students to be employees under one ID. HR is relatively universal all over the world, so the delivered functionality was sufficient for our then needs, simplifying or eliminating many manual procedures.
>
> With time we added 'our' subsystems:
>
> - Automated calculation of wages for special casual employees (simulated patients, invigilators etc.).
> - Recording and monitoring of annual leave for academics, which require a different set of rules from professional employees.
>
> On the other hand, functionality delivered by the SA package was very limited, as every university aims to be different. It was a paradise for both functional and programming staff. For SA we developed subsystems like Graduations (which was later used by Oracle to be included in the SA product), and Research Students modules — just to mention a couple.

[5] It has since been separated.

Barbara Bankovsky says that MISD staff often attended PeopleSoft conferences. She went to one in Las Vegas and stayed in the MGM Grand. In the front lobby was an actual lion, claimed to be a descendant of the one which growls in the MGM logo.

David Berriman thinks that the Parking Infringement System might still be running on Adabas/Natural. It was ported off the mainframe onto a Windows server running Adabas/Natural. David's team rewrote the Parking Permit System into a web based application (ANUBIS), but never the infringement part.

10.3.1 Disaster Averted

At some other universities the change-over to enterprise applications like PeopleSoft was far more expensive than at ANU and resulted in embarrassing loss of service and huge waste of money.

Stuart Kendall: Why did ANU's change to enterprise applications go more smoothly than it did elsewhere?

Dave: Was PeopleSoft a good choice?

Stuart: Yes.

Dave: Where there problems customising it to ANU requirements?

Stuart: No, but patching and upgrades became more complex over time.

Dave: Does PeopleSoft meet all of ANU's needs?

Stuart: No package would provide all ANU needs, but with its development framework it allowed us to adjust and expand to better meet requirements.

Dave: I'm aware of enormously expensive University IT disasters, such as RMIT. It seems as though ANU avoided that. How come?

Stuart: A good but very small group of developers (8 at most) and very strong business experts.

10.3.2 ANU's Administrative Gods

David Berriman told me the story of the godly names of online ANU administrative systems. The story began with ANU's Grants Office which created a database with the very natural acronym GODS (Grants Office Database System). It was a standalone system written in MS Access and utilised downloads from Entire Connection to maintain HR and Student data. GODS started a sequence:

ANUBIS Egyptian god of the underworld. ANU Business Information System.

APOLLO Greek god of archery, music and dance, truth and prophecy, healing and diseases, the sun and light, poetry, and more. Anu POLLing Online.

ARIES Greek god of war (usually spelled ARES). ANU Research Information Enterprise System.

HORUS Son of ISIS. Human Online Resources User System. Not really an acronym, but sounds a bit like HR and ends with S for system.

ISIS A major Egyptian goddess. Interactive Student Information System. The interface was given an Egyptian appearance. Unfortunately, it was not linked from ANU's homepage because of the incredible competition for space on the homepage – At the time it was restricted to 640 x 480 pixels! MISD resorted to promoting the system by other means, notably ISIS branded squeeze balls.

All these systems were developed at ANU by David Berriman's team. ANUBIS provided single-login access to all administrative systems through a uniform Administrative Framework. ANUBIS software was sold to the ACT Government, overseas companies and seven Australian universities.

The first attempt at ISIS (before it had that name and the Egyptian theme) had an interface that looked like 3270 screens. It was implemented in SmallTalk and some new APIs into the Fujitsu mainframe. According to David Berriman:

> It was not well received and resulted in the development of ISIS (with all of it's associated design and hype). When PeopleSoft came about ISIS was rewritten in MS ASP and embedded into ANUBIS. ANUBIS also did away with the 'themed' look-and-feel in favour of a more business oriented interface (designed by Martin Hedenstroem). The APIs into PeopleSoft were better than those into the mainframe so functionality also improved.

APOLLO was a very generalised response to a considerable demand from ANU people for support to conduct specific surveys. It won a Vice-Chancellor's award in 2004. It provided an interface to allow users to create their own customised servers. Because it was linked to ANU's PeopleSoft implementation, survey creators could precisely select the audience to be surveyed. For example, "All full-time undergraduate students", or "All casual staff." It was also secure to ANU.

APOLLO could potentially run online exams, though there was no real way to stop students getting answers from a second screen.

ANU
THE AUSTRALIAN NATIONAL UNIVERSITY
Prize for Innovative IT Application on Campus
Apollo
Corporate Information Services
Division of Information
Mr David Berriman
Mr James Blanden
Mr Martin Hedenstroem
Mr Adrian Thomson

The MISD team who developed APOLLO receiving an ANU Innovative IT award. From left: Adrian Thomson, David Berriman, James Blanden, and Martin Hedenstroem. ***Photo supplied by David Berriman.***

APOLLO was also extended to support electronic voting, motivated by a desire to reduce the time taken to count elections for ANU Council. In previous Council elections, paper votes were sent to Elections ACT and counted manually, with the poll declared a couple of weeks later.

While implementing voting in APOLLO, ANU's James Blanden found two logical errors in the statute controlling the Council election. After they were clarified, he implemented the ANU's voting system as well as first-past-the-post, Hare-Clark, etc. Elections ACT certified Apollo's counting, by running on its test data sets and comparing with the correct results. APOLLO is the only electronic system which supports ANU-specific counting.

The first ANU Council election run through APOLLO found David Berriman as Returning Officer in the Mills Room of the Chancelry. The Vice-Chancellor Ian Chubb came in and said to him something along the lines of, "This better bloody work." The button was pressed and 15 seconds later the final counts were printed, and then signed off. Ian Chubb said, "Good", and left.

David Berriman says that ARIES (University Office) was the replacement for GODS and, although built and owned by a private company, utilised ANUBIS interface components. The ANU was looking for a web-based replacement for GODS, but there were no suitable options at the time. University Office[6] had most of the required functionality but was a standalone client/server system – the company owner was contracted to redevelop his system for the web and spent many months at ANU learning from our Web Applications Team and DBAs how to go about this. (He was given access to our code and graphics, which he then adapted.)

[6] `https://universityoffice.com.au/`

10.4 2021 Situation

Information Technology Services is a single business unit responsible for audio visual services, network and cloud infrastructure, information commons, timetabling, standard desktop environments and all-ANU service desk as well as ANU's business systems.

Helen Duke (currently Director, IT Services) asked Emelyn Pecar to answer a set of questions I asked about the current situation. Release of the answers was approved by Garry Whatley (Interim CIO).

Emelyn Pecar: Current situation in IT Services

Dave: Back in the 1970s and 1980s most Research Schools and the Faculties (today's Colleges) operated their own computer systems (other than desktop services) and had their own computer support staff. E.g. the Joint Schools operated a DEC-10 computer and had a School Computer Unit with about half a dozen staff. Do colleges still have their own IT staff, or is support mostly provided centrally?

Emelyn: Support is mainly provided through Information Technology Services (ITS). The Service Desk provides support via phone, the online Knowledge Base, in person, and through the online ticketing system ServiceNow. The Acton campus is divided into four service precincts (North, South, East and West), each of which has a Service Hub team that provides IT support to colleges within the precinct area. There are a small number of schools who still have their own IT staff, however this would only amount to approximately 20 or so people.

Dave: Roughly how many staff are there in IT services?

Emelyn: Approximately 170 employed by ITS. This does not include the office of the Chief Information Security Officer which reports to the Vice-Chancellor.

Dave: Roughly how many IT staff would be employed at ANU outside ITS?

Emelyn: There are many other administrative staff employed across ANU who perform IT functions, however not necessarily in a support role. Approximately 140 who identify as performing an IT function for more than 50% of their role (UniForum 2019).

Dave: What proportion of ITS staff would have university qualifications?

Emelyn: We don't have any numbers on this, but would hazard a guess at 90-95% as having a degree/qualification is generally a requirement for ITS roles.

Dave: At a typical point in time, how many ITS staff would be working on software development projects involving more than say a person-year of effort?

Emelyn: ITS doesn't generally develop software applications — we use Software as a Service and customise and configure off the shelf applications. We do have people involved in developing data integrations, website maintenance, and PeopleSoft development activities. This is covered by the Application Services portfolio which is approximately 35–40 people.

Dave: How much use does ANU make of external contractors and consultants?

Emelyn: Contractors or non-ongoing staff are often used for project delivery to supplement existing staff and this number varies with the number and complexity of projects being undertaken. Consultants are used occasionally where specific expertise is required.

10.5 Prominent People in Administrative Computing

Pat Wilson was the first head of DPU and MSG. See Page 237.

Skaidrite Darius spent decades in charge of DPU. Her story is told on Page 274.

Ian MacNicol was head of the group which supported administrative computing at least from the late 1970s. Initially it was known as *Management Services Group* (MSG), but on 01 January 1986 separated from the CSC and became the *Information Systems Division*. The CSC provided facilities management for the new Division and a new portable building was erected to house it. **Ken Vine** was appointed as Head of the Division but did not take up the post until March 1987.

Don Hardman was in charge of ISD in the interim, later moving to Buildings & Grounds where, in the popular mind, he became associated with the profusion of concrete spheres which appeared on campus, serving as traffic bollards.

In April 1987, the Division was renamed as the *Management Information Systems Division* to differentiate it from *University Information*.

Ken Vine took over as Head, MISD in 1987 and left ANU in 1990. Prior to this role, he had been a catholic priest, and later a senior executive in the Public Service Board. In 1983 he submitted an ANU PhD thesis in Psychology titled, *Development of Linear Measurement*.

George McLaughlin, Business and Technical Manager at the Research School of Chemistry, had been involved in the evolution of management information systems under Don Hardman and Ken Vine, and, when Don left, was seconded to replace him. George's interesting career at ANU, AVCC and AARNet is documented in the panel on Page 249.

Fay Gibbons. ***Photo: David Berriman***

Fay Gibbons became Head, MISD on George's departure in 1994, and continued in the role for many years until her unfortunate illness and death. Barbara Bankovsky rated her very highly. She was a trained programmer with real-world experience, and had a good head for money. Barbara says that the PeopleSoft project grew out of a visit to Sydney by Fay and Robyn Mulholland, who was Director of Student Services. When Fay was admitted to hospital with what turned out to be her final illness, Robin Stanton visited her there, and discussed interim arrangements with her. To Barbara's surprise, Fay nominated her as Acting Head, looking for someone with a broad view of all of MISD's activities.

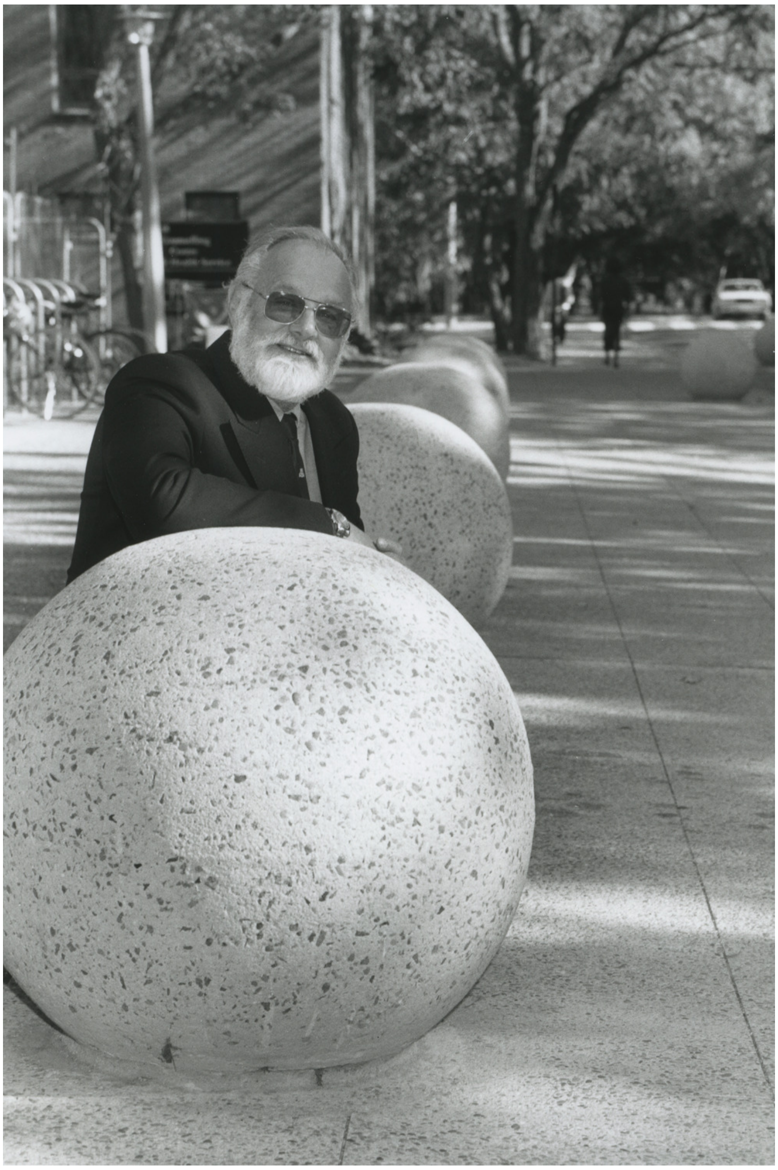

1996: Director of Facilities and Services, Don Hardman, with the newly installed 'traffic calming' balls near the ANU Sports Union. *Photo: ANU Archives.*
https://openresearch-repository.anu.edu.au/handle/1885/144908

Robin Erskine notes that Fay Gibbons initially reported to the PVC (Admin) – Chris Burgess. Chris had overall charge of Y2K and software/hardware procurements for admin computing. "I found him very good – one had to argue the case strongly but when he was convinced he supported you well – both Fay and myself."

In 1995, **Rick Van Haeften** joined ANU after the completion of a huge PC and networking project in Defence. He played a series of important roles in the evolution of administrative computing at ANU up to 2010, when he moved to Brisbane. He tells the story in the panel on Page 250.

After Rick left ANU, the ANU CIO position has been held by Peter Nikoletatos (2012 – 2014), Karen Hill (2014 – 2019) and Garry Whatley (2019 –).

A CIO Magazine article[7] anouncing Peter Nikoletatos's departure from ANU reports him as saying, "I knew it wasn't quite right and agreed to separate [from ANU]", and declining to comment further on the circumstances. According to the report, ANU was then looking to reduce its total IT budget (of up to $70M p.a.), had replaced 32 different email solutions with Microsoft Office 365, had replaced ten different helpdesks with one, and was looking to rationalise networks, storage, and private cloud.

[7] https://www.cio.com/article/3499765/anu-cio-peter-nikoletatos-departs.html

George McLaughlin recalls his career at ANU, AVCC and AARNet

Glen Robertson, from RSC, ANU was on sabbatical at the University of Sussex with my then boss Ronald Mason. As well as supporting Ron's group, I spent time at the Science Research Council's ATLAS computing lab as most of our tasks involved overnight computing runs, and early exposure to "networks."

When Glen returned to ANU, he wrote and asked if I'd like to come to Australia for a couple of years to help get the crystallography service up and running. My initial appointment in 1974 was as a Senior Technical Officer, but this morphed to a Research Officer.

Back then, XRay crystallography (like astronomy) was one of the computing intensive disciplines, and I quickly became involved with the ANU computer centre. There may be some references to my user code GMM500 in the archives, since that account held most of the crystallography software. (500 was the code for Inorganic Chemistry, RSC, and GMM was me.) I served on the Computer Centre users group and chaired it for a while.

In 1983, I was seconded as Business and Technical Manager, RSC, and with my background in computing and systems, became involved in the evolving Management Information systems, working with Don Hardman, Ken Vine and others. When Ken Vine left ANU, I was seconded as Acting Head, MISD in January 1991. I was later confirmed in the role.

During the next few years, the Australian Vice Chancellors Committee (AVCC) initiated the Core Australian Standard for Management and Administrative Computing (CASMAC), and I became heavily involved with this. In 1994, I was seconded to the AVCC to manage the CASMAC project. Fay Gibbons took over the role of Head, MISD.

CASMAC eventually failed because trying to get a consensus from 37 different universities on how to manage their admin systems was a tall ask. CASMAC split into several groups. ANU, Sydney and Macquarie formed a Natural/ADABAS subset.

In 1995, the AVCC "sold" AARNet to Telstra, and Geoff Huston and David Woodgate moved across to Telstra. AARNet bacame "the Telstra Internet." Peter Elford had already left to join Cisco. Andy Linton and Brenda Aynsley left to do other things.

The initial expectation was that the universities would simply buy services from Telstra Internet, but much to Geoff's frustration, Telstra refused to believe the exponential growth of the Internet and a torrid period followed, essentially rebuilding AARNet, and involving Optus. (I'm glossing over a heap of issues here.)

The AVCC asked me to help manage the new arrangements, and put new flexible approaches in place, establishing AARNet as an incorporated company limited by shares, with its own carrier licence (coincidentally licence number 61, the same as Australia's international telephone dialling code), its own international links, and Australian representation in the global NREN (National Research and Education Network) forums. Issues discussed in these forums included: direct participation in evolving global science projects, supercomputing initiatives, high energy physics, astronomy, and then into telehealth, new modes of educational delivery, disaster mitigation, etc.

When AARNet became an incorporated entity in 1999, I was appointed to lead the organisation, and did so until around 2004. My health then deteriorated and I reduced my involvement to the international collaborations. I parted company from AARNet around 2006.

Rick Van Haeften's involvement in administrative computing at ANU

At the completion of my RAAF project, in 1995, I joined ANU as a senior desktop support officer in the then Management Information Systems Division or MISD. My reporting line was to Fay Gibbons, the Director, Corporate Information Services (CIS). Fay reported to Robin Erskine, the Director Information Technology Services, who in turn reported to the Deputy Vice Chancellor. Later the reporting line went to Robin Stanton, as Pro Vice-Chancellor (Academic).

I later advanced to the position of Manager, Network and Desktop Support. The position carried the responsibility for management of the Network and Desktop Support Services Team; charged with providing first, second and third tier helpdesk support services to over 600 users (and network devices) within CIS and the wider University. The Helpdesk also provided support to the whole of campus for all IT matters related to Administrative systems.

In 1997 ANU embarked on the Enterprise Solution Project (ESP), which involved the replacement of the Financials, Facilities Management, Human Resources and Student Administration systems with the PeopleSoft Enterprise Resource Planning (ERP) suite of applications. My role expanded to include the initial design and implementation of the information infrastructure to support the multi-year project, including ongoing maintenance, upgrades and associated improvements to these and other major administrative and academic enterprise systems.

Some specific tasks relating to the project were:

- Design, documentation and implementation of the necessary IT infrastructure to support the development and production phases of the project.
- The implementation of Thin Client technology to provide access to the ERP applications across campus in a multi-platform desktop environment, Microsoft Windows, Apple Macintosh and Unix/Linux.
- The innovative use of web based technology to streamline processes and provide easy access to the helpdesk and a variety of software tools and applications in support of system developers and end users.

During the same period I was also responsible for coordinating Y2K activities to maintain unfettered access and operation of core University information resources.

At the formation of the Division of Information (DoI), I was appointed to implement and lead structural change within the various elements of the DoI. The role focused on the development of an integrated information support structure to manage the overall planning, development and maintenance of comprehensive information infrastructure services in support of the corporate information portfolio. Formation of the structure entailed merging the staff, equipment and activities of five disparate providers of information services.

In 2005, I was appointed to the role of Director, Corporate Information Services. The role carried the responsibility for leading the development of the University's corporate information services. At the divisional level the role was active in the formulation and support of divisional programs, through resource and management contributions, with associated management responsibility as required. These programs involved managing cross-divisional and campus-wide resources.

In 2009, I was appointed to the role of Director, Information Services. The position carried University-wide responsibilities which included leadership and management of an integrated information infrastructure support group consisting of more than 200 professional staff at various levels of management. The group was engaged in the planning, development, maintenance and delivery of a comprehensive suite of information infrastructure support services to the entire University community (staff, students and visitors). This included the provision of Service Desk, Teaching Space Support and Audio and Video Conference development and support services.

At the strategic level the focus of my role was on leadership and direction of information and communication technology and services in support of University strategic goals. I was responsible for the provision of advice to executive management and university advisory committees and for formulating IT strategy and policy.

Rick Van Haeften's reflections on former ANU colleagues

Robin Stanton. I found Robin (RBS) to be a very affable and approachable member of the senior executive. Robin had a tough job. Not only did he have to deal with the other members of the Senior Executive, he had to also navigate the many thorny issues that arose from the College Deans and their demanding Business Managers. How he managed to keep it all together, I will never know.

Robin was a guiding light for me and opened my mind in so many different ways. My career at ANU particularly, in the early 2000's, was, I felt, strongly supported by Robin. I learned much from him and owe him a lot for giving me the opportunities to advance. Robin and I remain friends and colleagues to this day.

Robin Erskine. My first meeting with Robin was with Fay Gibbons in Robin's office in the Leonard Huxley building. I remember thinking of him as someone who was confident and relaxed yet sharp and attentive. He didn't miss much. I remember writing a number of memos and letters that required Robin's approval. It wasn't often that they came back from him without changes and corrections. I learnt very quickly to triple check everything, content, spelling, grammar and punctuation before sending anything to Robin. This stood me in good stead during the course of my career. Robin was very supportive of change and new ideas. He gave me the opportunity to form the Systems and Desktop Support group within the Division of Information (DoI). This group comprised all the IT support elements within the DoI and include the Local IT Support Staff (LITSS). It was hard going as there were a number of 'ornery' characters that took some convincing to come on board. Robin was there throughout it all and I could not have made it happen without his advice and support. Like RBS, Robin Erskine and I remain friends to this day.

Fay Gibbons was Director of Corporate Information Services. She passed away in July 2005 after a period of illness. I worked closely with Fay and found her to be a very dedicated, conscientious and caring person. Fay was a vocal proponent for change in a number of key areas. During my time at ANU, Fay's greatest achievement was without doubt the implementation of the Enterprise Solution Project (ESP).

Fay was also responsible for the creation of the Business Support Groups or BSG model which has stood the test of time. Without this group in place the implementation of the ESP would have floundered.

Fay, like both the Robins, gave me a tremendous amount of support and sage advice. Her business acumen, coupled with her no nonsense yet empathetic approach made her a great boss to work with.

Geoff Barlow. Now here's a guy who reminded me of the 'energiser bunny'. He just did not stop. I had and still have tremendous respect for Geoff and his knowledge and abilities. He was able to adapt very quickly to the many major changes that happened in my time at ANU. "Just tell me what you want and leave me to get on with it." That's how I always think of Geoff. He did not settle for 'near enough'. His attention to detail was something to aspire to.

Geoff was given responsibility for running and maintaining the machine room. The room housed the main networking equipment nodes, the Super Computer and other clusters including, all the equipment for the student and administrative computing. Geoff, set up some very strict processes and procedures that controlled access and activity within the room.

Brok Glenn. As Director, Information Services, Brok Glenn was my line manager in his role as Chief Operating Officer. I had a lot of respect for Brok's role and he was very supportive during my time reporting to him.

Allan Williams. I first met Allan when he arrived at ANU about the same time as I did. Allan was part of the Web (WWW) team. I knew little about such things at the time and although Allan did not know it, I was very interested in the activities of the team. It was many years later at the formation of the DoI that Allan and I began working together. Allan always seemed a very serious guy, I rarely saw him smile and you could never tell what he was thinking. Allan could always be relied on to get things fixed under pressure. I can't remember him every having a 'bad' thing to say about anyone. A real gentleman. Allan's technical knowledge together with his capacity to design, implement and manage complex and effective end-to-end services enabled the DoI to significantly move forward from a fledgling start. Allan's depth of his knowledge and experience in IT security, ensured that the services hosted and provided by the DoI were as secure and robust as possible.

Allan was my '2IC' and I had total trust and respect for his technical knowledge, his strong people skills, his capacity for leadership and management, and his determination and honesty.

10.6 Security

Although ANU computers and individual user accounts have been subject to frequent attacks over the years,[8] it is clear that in 2021 cyber attacks are now more frequent and far more sophisticated than ever before. Protection against damage to infrastructure, fraud, and theft or, possibly worse still, alteration of ANU data is more important than ever before. Cyber security is also related to the Australian Government's concern with foreign interference.

In May 2018, there was an intrusion into ANU's systems by unknown persons. ANU came under significant pressure from the government to take decisive action, but quickly realised that its small security team within IT Services was not capable of responding the way the government expected. Brian Schmidt decided to appoint a Chief Information Security Officer (CISO). Suthagar Seevaratnam took up the post.

His first job at ANU was to carry out a threat analysis on ANU's network. This was difficult because of lack of knowledge about the setup. He brought in two external firms to carry out the work. In November they found a new and very serious breach. Suthagar called an emergency meeting with Brian Schmidt. At that stage they feared that the attacker was targeting vulnerable students. Suthagar recalls walking past the Chifley Library on their way to mid-year exams and felt sick at the scope of potential harm to them. Fortunately, it later became clear that vulnerable students were not the intended target.

It had been envisaged that ANU's CISO might prepare a five-year cyber security plan but the presence of a serious breach focused the mind. In June 2019 ANU Vice-Chancellor Brian Schmidt released a comprehensive public report into a major cyber-attack on ANU's systems which occurred in November 2018.[9] Suthagar reports Brian Schmidt saying, "It's easy to be hacked when the hackers spend more on IT resources than you do."

ANU is very different from most other large organisations. It's spread over a large campus, welcomes students and visitors on its very complex network, deploys a very wide range of hardware and software, and expects a considerable degree of freedom for its academics. Any attempt to defend against cyber attacks must allow members of the university to do what they need to do without imposing unnecessary barriers or impediments.

Suthagar wants to make things better at the same time as improving safety. In early meetings with senior staff, he received endless complaints about passwords. In response, he has removed ineffective rules about the characters comprising a password, and significantly increased the interval between forced password changes. He is also working with Thycotic and an Australian company Forticode, towards a passwordless access mechanism based on smartphone biometrics. Authentication tokens passed around the network will be one-time only – of no lasting use to hackers. Suthagar says that this system could be a world-first.

During 2020, Suthagar and his team focused on securing email, blocking a lot more spam. Around seven million messages a day are now blocked.

ANU has invested in Network Packet Brokers and is now able to map the entire scope of its network, including private networks within individual schools. They can detect obsolete and unpatched operating systems and machines infected with malware. Suthagar says ANU is building a new network, starting now. When fully implemented, each time a new machine is physically connected to the network, its owner will be required to declare whether the machine is a centrally managed machine (in which case updates will be automatically applied), or a BYOD (Bring Your Own Device) or a travel laptop, and identify themselves before any network access is allowed. The scope of access to be allowed will depend upon the type of machine. Virtual desktop technology may allow BYOD users to safely access more privileged services.

The new network will be software defined – infrastructure as code, allowing rapid change in configuration. No device will be trusted until it and its owner have attested their identity.

Suthagar's team of around 30 people works on many aspects of keeping ANU people and sys-

[8] See Pages 223 and 231 for examples.

[9] `https://imagedepot.anu.edu.au/scapa/Website/SCAPA190209_Public_report_web_2.pdf`

tems safe: e.g. Education, Cyber Engineering, Threat Intelligence, and Governance. They also scan the Dark Web for stolen credentials. You can access a lot of resources including humorous videos featuring Brian Schmidt, Julie Bishop and Canberra comedian, Chris Ryan at `https://cybersense.anu.edu.au/`.

10.7 Responding to Natural Disasters

David Berriman recalls that, after the 2003 bushfires which devastated Mt Stromlo Observatory (MSO), ITS staff were among the first people allowed back in. They helped recover data from damaged computers, particularly admin data.

The 2003 bushfire hit MSO on Saturday, 18 January. A wall of flame rolled over the mountain, destroying telescopes, workshops and equipment. Multi-million dollar instruments being built at MSO under the guidance of Peter McGregor for the multi-national Gemini telescopes were destroyed, along with the machines which would be needed to reconstruct them.

John McGee is proud that the then recently introduced “common IP fabric”[10] project with Voice Over IP phones, allowed restoration of phone and data service to MSO employees who had been relocated to the ANU campus, *by the Monday immediately following the fire.* Microwave connection to MSO from the accelerator tower in Nuclear Physics was restored within two weeks, and upgraded to 1Gbit/sec. Soon after a fibre-optic link to MSO was established using the Intra-government Communications Network (ICON) as far as Woden and new trenching to the mountain top.

ANU has been hit by two devastating hailstorms. The first was on 27 February 2007. It covered the ground to a depth of 20cm and created metre-deep haildrifts in Civic. It killed birds and 300 bats in Commonwealth Park. It destroyed glasshouses, skylights and roofs across ANU. Llewellyn Hall was out of action for months. Rick Van Haeften recalls:

> I remember sitting on my balcony at home, sipping a glass of red while watching the lightning around the city. Little did I know that a massive supercell storm had centred itself over ANU and was wreaking havoc across the campus. Shortly after the storm broke, I began getting a few phone calls about the loss of many IT services. I was aware that people like Geoff Barlow and others would be doing their best to get things under control and restore services as quickly as possible.
>
> I was in early the next morning and it gradually dawned on me that we were not just dealing with a storm but rather a catastrophe. The hailstorm was so severe, that ice had solidified in the guttering of many buildings including the Leonard Huxley building which housed almost all of the Campus network and computing equipment. Needless to say there was considerable damage and many services were offline for extended periods.

If anything the hailstorm on 20 January, 2020 was more severe. It occurred soon after ANU had to close due to smoke from bushfires. It has been reported that 44,000 ACT-registered cars were damaged by hail, many of them written off. In November 2020 repairs were still under way to many buildings, including University House, whose Fellows Garden and Boffins remain closed.[11] In January 2021 glasshouses had still to be repaired.

[10] See Page 266

[11] As at 15 May 2021.

Chapter 11

ANU: Networks

ANU Computer Centre Technical Report No. 38 of August, 1971, by D.E. Lawrence, was titled *A Proposed Computer Network for the Australian National University*. Driven by then-current government funding policy, it proposed a Joint Computer Centre shared between ANU and the CSIRO Division of Computing Research, and links of at least one megabit per second between the main computer and other computers on ANU and CSIRO campuses. TR38 includes the following quote:

> One of the major fallacies perpetrated by computer personnel during the 1960s was a firm belief in "Grosch's Law". This is an axiom which states that "the power of a computer increases with the square of its purchase price." The corollary of this law is that it is therefore more economical to buy a larger computer.

Technical Report No. 44 of January, 1974, by Ted Hempstead and Rex Jones reported the construction and deployment of TEX-11 (TEX = Ted & Rex), a hardware device connecting to the Unibus of a PDP-11. Ted & Rex connected the Computer Suite (Menzies) with the Computer Centre (Cockroft) and achieved 2.4 KBits/sec., severely limited by the PMG (Post Master General's Department, later Telecom, later Telstra) modems they were legally obliged to use. Plans, never realised, were in hand to upgrade the line to 10 MBits/sec, using a fibre-optic pipe.

The PMG had absolute legal control over all telecommunications. Legally, your child needed to obtain PMG authorisation to connect a tin can by a string to a tin can across the neighbour's fence! By such restrictive rules and later by seemingly exorbitant pricing of data services, the PMG and its successors applied an effective brake on data communication in Australia for many years.

11.1 Interactive Access Across Campus

As noted on Page 49, from 1973 on the CreasyLink provided terminal multiplexing over a phone line with 2.4KBits/sec. modems. You may think that a small share of 2.4KBits/sec. would be very slow, but almost ten years later, as chair of the KA-10 replacement exercise, Bob Watts declared that no student required more than a 300 Bits/sec. connection and that the University would not pay to provide higher speeds. Geoff Huston says that, after the change to VAXes, a DEC Field Service engineer (Blair Phillips) was persuaded to write a note saying that 9600 Bits/sec. was a minimum!

I believe that the CreasyLink connected directly to the Communication Terminal Module Control (CTMC) of the Univac 1108. The Computer Centre Annual Report for 1972 reveals the extent to which Computer Centre staff were involved in operating system software debugging and hardware design and fabrication:

```
... the Centre's PDP-11/20 was upgraded by DEC to an 11/45 in return for work done
by the Centre in isolating and correcting faults in the system's software.  This
machine will play an increasingly important role in handling communication
with the U1108, and in controlling a range of peripheral devices.
...
The PDP-11 is connected to the U1108 by means of the communications controller.
```

```
...
The U1108 [16-bit] link is at the stage where the printed circuit boards are
being assembled ...
```

I believe that Bruce Butterfield was working on the software (called COM I believe) to link the Univac and the PDP-11/45. I remember Erin Brent telling a story that Bruce took a sudden trip "on a slow boat to China", and didn't return, when every backup copy of the source code (kept on removable disks) was erased by a DEC engineer attempting to diagnose a hardware fault on the PDP-11. The engineer kept loading disk packs trying to get the machine to boot, but the fault was in the memory controller and it caused disks to be overwritten.

Recent discussions with people such as Richard Brent, Les Landau, Bruce Millar, and Julie Bakalor tend to confirm that an incident like that did occur, but precise details are lacking. It is the case that Bruce Butterfield left the Centre in about 1973 and returned several years later.

However, the 1973 Annual Report documents pleasing progress:

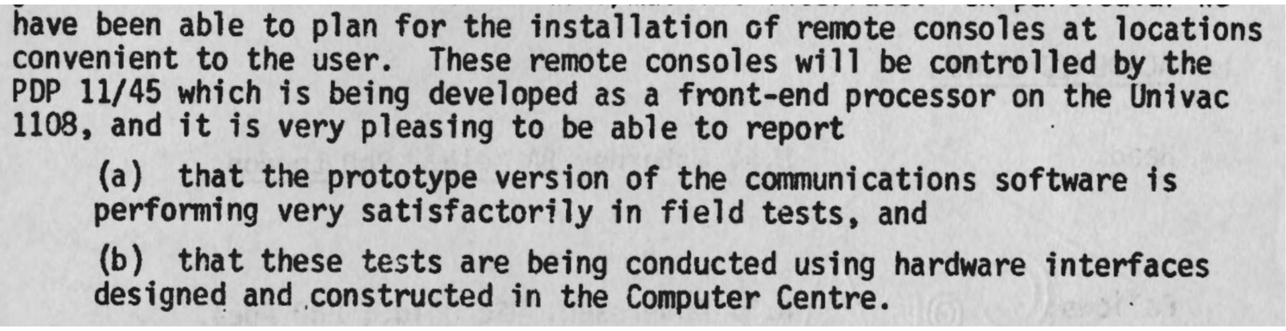

have been able to plan for the installation of remote consoles at locations convenient to the user. These remote consoles will be controlled by the PDP 11/45 which is being developed as a front-end processor on the Univac 1108, and it is very pleasing to be able to report

(a) that the prototype version of the communications software is performing very satisfactorily in field tests, and

(b) that these tests are being conducted using hardware interfaces designed and constructed in the Computer Centre.

At some stage, possibly immediately after the disaster, John Chmura took over development of COM.

Communication between the PDP-11/45 and other machines on campus relied on cables laid in a network of trenches around the campus. The cables were twisted-pair telephone cables protected by lightweight plastic 'ag-pipe'. Problems arose because the trenches were back-filled using the rear wheel of a tractor to compact the soil. That crushed the the pipe and prevented any further cables from being pulled.

John Chmura's software reached its limits after a while. He said that the addition of a single instruction to the main interrupt handler caused the system to fail.

COM was replaced by ANUNet which was also based on PDP-11s. The project, under Garth Wolfendale, started in October 1976 and was based on DECNet Phase II. According to a Wolfendale paper in volume 4 of *Computer Networks* journal, 1980,[1] the project was essentially complete by December 1977.

ANUNet continued for several years with Adrian Mortimer in charge. Greg Preston, famous for nearly always wearing orange clothing, worked on the project. ANUnet grew to include four PDP-11/34 nodes around the campus plus a central 11/50 in the CSC and an 11/45 development machine. The 11/50 was eventually connected to the Univac 1100/82 by a 50 Mbits/sec. NSA HyperChannel. Campus cabling was re-ducted and replaced.

Three MICOM PACX switches connected by 1.54 Mbit/sec links connected hundreds of terminals in the Faculties to a range of different computing services.

> **Geoff Huston:** I remember doing the evaluation for that purchase in around 1983 – we looked at the Micom and the Gandalf switches. We settled on the Micom because of the HUGE speed it achieved over twisted pair copper wires!

[1] https://www.sciencedirect.com/science/article/abs/pii/037650758090029X

11.2 Networking in DCS

In DCS, serial lines (remember RS-232?) were run from the computing facilities (CreasyLink and Nova) in Copland G10 and from the Burroughs (which I think was located on the lower ground floor of the Crisp building) to a patch panel among our offices. Each office was wired to the patch panel and had an outlet in the form of an old-style (6.35mm) phone jack, and eventually we all had Volker-Craig VC404 or DEC VT100 terminals.

Departmental offices and the terminal rooms in Copland G5 were wired up by Bill Crystal, a retired PMG technician. Bill was able to accurately terminate a 100 pair cable in as little as 20 minutes. This was remarkable! Each pair consisted of a signal wire which was colour-coded and slowly twisted around a corresponding white earth wire. He had to match the pair of colours (e.g. slate and pink) on the signal wire to the number, say 47, of the pair and then pick out the corresponding earth wire out of a hundred of them. Then he had to neatly and reliably solder that pair of wires to the correctly numbered pair of terminals on the terminal block. Oh, and I almost forgot, the pairs had to be cut to different lengths to avoid a mess of wires in the box.

To start a session on say the Univac, you would go to the patch panel, look for a spare Univac line, and plug in the cord corresponding to your office. Because lines were scarce, you were expected to unplug as soon as you finished your session. Raids down the corridor were sometimes necessary to determine whether a line was actually in use.

This system was dramatically improved when the department acquired a North Star Horizon microcomputer and used it to replace the manual plugboard. Terminals in offices were permanently connected to the NorthStar and special characters were used to connect and disconnect sessions. I think that the multiplexing software, in Z80 assembler, was written by Rob Ewin.

> **Geoff Huston:** Sean Batt, who was a high school student at the time, was bought in by Rob to help out as I recall.

A DCSnet Status Report from 9 April, 1984, written by Rob Ewin, reports that the Horizon was about to be upgraded to 24 incoming computer lines and to 32 lines for offices and workshop. The facility allowed connection of staff terminals to Suns, departmental VAX and to CSC machines through the MICOM. The Statistics department had also acquired a North Star (STATSnet) which ran the same software. Rob was investigating the possibility of increasing the speed of some lines above 1200 baud.

> **Geoff Huston:** I was at the CSC in 1982/1983 in the ANUnet group and I recall trying to convince DCS to drop the local switch and jump onto the Micom! I don't recall why they didn't want to make the move.

The same report implied that some staff were using MicroBee computers as terminals. A Wikipedia article[2] states that the original MicroBee was designed in Australia by a team including Owen Hill and Matthew Starr, was based on Z80 chips, and used DGOS, the David Griffiths Operating System. A Sydney Morning Herald article from 13 September, 1987[3] gives more detail.

Rob proposed that the old plugboard (which had been retained as a backup and a bypass mechanism) be phased out. He also gently complained about: time required to debug problems with STATSnet and Statistics department wiring; the lack of a technician to do wiring; the lack of space for the network equipment; and lack of facilities for developing the DCSnet software. He was investigating use of cross-compilers from the Department's Discovery Multiprocessor, VAX, or DEC-10.

In the early 1980s, Ethernet came to DCS. Peter Wishart tells the story:

[2] https://en.wikipedia.org/wiki/MicroBee

[3] https://www.smh.com.au/business/small-business/secret-of-project-granny-smith-20090619-cp67.html

Peter Wishart: This new fangled Ethernet thing

1983?: The department acquired a few more Sun workstations. This time it was Sun-2 workstations and we had more than one, so we needed a LAN to network them. The workstations were shipped with original Ethernet transceivers, which were ‘in-line’, meaning you had to break the Ethernet cable and connect the cable to both ends of the transceiver. But we had not ordered any Ethernet cable with the workstations, so we had to acquire that separately. And we could not quickly and easily find a source for the cable at that time. Robin Stanton and I went to local Tandy store and bought some cheap coax and some appropriate connectors. We soldered the connectors onto the coax and screwed in the transceivers, and then we had a small LAN Ethernet network connecting our new Sun Workstations. Robin was in his element using a soldering iron, it took him back to his engineering roots. For me, as a programmer and software guy, it was a new experience. It all felt so ad hoc, but it worked, probably because our network was so small. We did eventually get some proper Ethernet cable and not long after those original in-line transceivers were replaced with ‘vampire taps’ which connected directly to the Ethernet cable without breaking it.

Sun originally supplied a 3Mbits/sec Ethernet board implementing the Xerox-PARC standard with Sun-1s, but later upgraded the specification to the 10Mbits/sec 3COM version.[4]

11.3 Dial-up Service

For a considerable period, the CSC operated a dial-up service with a bank of Telstra modems, which Telstra took pride in saying were rated only to a speed of 2400 baud (bits/sec). The service was highly unreliable and generated many complaints. Of course the problems only occurred out of hours when the users went home and tried to use dial-up. Robin Erskine still recalls one complaint received from a senior academic who had been unable to order flowers for his mother’s birthday! Eventually the CSC went to tender to outsource ANU’s dial-up service. Robin Erskine and Rob Ewin ran the tender process, with Ozemail winning the business in a close race with TPG. For several years, Malcolm Turnbull was Chairman and a significant investor in Ozemail. According to Wikipedia[5], Malcolm made a 11,300% profit on his investment in Ozemail when it was sold to WorldCom in 1999. However, neither he nor Ozemail “invented the internet in Australia.”

Extract from Tony Abbott speech, 28 June 2013

In a speech to his Liberal and National parliamentary colleagues, Opposition Leader Tony Abbott has claimed that Shadow Communications Minister Malcolm Turnbull “virtually invented the internet in Australia.”

Josh Taylor article, ZDNet, 28 June 2013: `https://www.zdnet.com/article/turnbull-virtually-invented-the-internet-in-australia-abbott/`

DCS also provided a primitive dial-up service. Ray Jarvis built a telephone answering system out of FischerTechnik. When it heard the phone ring it would activate a stepper motor to raise the handset off its cradle. There was a departmental Silent 700 terminal (made by Texas Instruments) which operated by acoustic coupling. It had a couple of padded receptacles into which the ‘talk’ and ‘listen’ ends of the telephone handset could be accommodated. The accoustic coupler listened to and interpreted whistles and chirps from the other end of the line and transmitted back its own whistles and chirps.

11.4 ACSNet

According to Roger Clarke’s *A Brief History of the Internet in Australia*[6] the Australian Computer Science Network (ACSNet) was established by Robert Elz at the University of Melbourne and Bob Kum-

[4] `https://en.wikipedia.org/wiki/Sun-1`

[5] `https://en.wikipedia.org/wiki/OzEmail`

[6] `http://www.rogerclarke.com/II/OzIHist.html`

merfeld and Piers Lauder at the University of Sydney "in the mid 1970s." It used store-and-forward techniques to provide mail and file transfer services, using dial-up and, for interstate links, CSIRO's X.25 network called CSIRONET.

> **Geoff Huston:** The reason why it used CSIRONET was that CSIRONET was the 'front end' for CSIRO's mainframe system that it was trying to position as the computer bureau for the entire public sector. They used to offer a cards-in, printout-back service using Australia Post and couriers (I used it in my brief stint in the public service in 1979) and then slowly switched to a nationwide remote job entry model where CSIRONET provided the connectivity. They did not charge to use CSIRONET – they charged to use the mainframe in Canberra and these charges 'covered' CSIRONET costs. So when the USyd and UniMelb CS folk looked at the CSIRONET nodes (PDP-11/45's as I recall) they realised that all nodes could talk directly to any other node for free! As soon as they figured this out they jumped on with UUCP. UUCP was a single threaded message system – only one message at a time could be transmitted, and if that message was a (gasp) 10MB file then everything else backed up. ACSNet was a refinement of UUCP that allowed 3 virtual circuits in each direction. Small messages were put on the 'fast' queue and large messages were scheduled on the soaker queue. So ACSNet built up the core of their network riding on a free CSIRO platform.

Peter Wishart played an important role in establishing an email service based on ACSNet for DCS and ANU as a whole.

Peter Wishart: Email – Connecting ANU to the World

1983–5?: The department got one of the very first Sun workstations in the world (serial number <100). That workstation provided the department with a small machine to run a number of Unix applications. One of the first uses was to connect to email around the world. I set up a connection to ACSNe. It was network of computer science departments around Australia using Unix based software developed by the University of Sydney (the Basser Department of Computer Science). Typical of our make do approach, we had to find a cheap way to connect by dial-up to either Sydney or Melbourne to access ACSNet. I set up a connection over CSIRONET, which was an Australia wide X.25 based network accessible from the ANU through a local phone call. This provided the transport mechanism to get to Sydney (`basser.oz`) and Melbourne (`munnari.oz`). We ran the ACSNet protocol as a dial-up store and forward network over that connection. That connected our department (`anucsd.oz`) to ACSNet for email and file transfer, which in turn connected us to UUNET, ARPANET and sites in UK and Europe.

It was the first connection of email at ANU, giving access from people's desks to email around the world. I was the postmaster running the connection out of ANU. Academics outside the department discovered we had this access and were lining up to use it. But most other departments at ANU were using VAX/VMS as the time. I discovered some undocumented features in VMS which allowed me to write an interface between VMSMail and our departmental Sun (Unix). This allowed people on VMS using the ANU DECNet connections to access email around the world via ACSNet. This was a real plus for people at places like Siding Spring Observatory (SSO) and Mount Stromlo Observatory (MSO). This was before the Internet, so email had to be explicitly routed through key nodes, meaning that the email addresses could be long and complex. Addresses like `john%uk.edu.cambridge.cs!siesmo!ucbvax%otherpace.eu@munnari.oz` were common.

The department had a regular seminar series where academics and PhD students presented on various topics. They would typically get 10-15 people to these seminars. I ran a few seminars on how email addressing worked, so people could better route their emails around the world. My seminars were standing room only with 50–75 people regularly attending. I think the popularity of these seminars bemused a few of the academics in the department at the time.

When I returned to DCS in 1988, ACSNet was still in service. Every hour, a machine at the CCAE would dial up a machine in DCS and transfer email in both directions. (This was CCAE's only email service.) Our machine was permanently connected to a machine at ADFA, which was permanently connected to a machine at UNSW. The chain of such connections ended at a Melbourne University machine, whose name was `munnari.oz`, which since the early 1980s had had an email connection to ARPANET in the USA.

> **Geoff Huston:** It used to call Rick Adam's system at the United States Geological Survey (seismo.usgs.gov), Piet Bersema's system at Amsterdam (mcvax) and one other (Berkeley?). DEC picked up the phone bill for a few years, but this stopped in around 198?, and KRE (Robert Elz) wrote some accounting software and starting billing folk for the volume of outgoing mail to pay for the phone calls. One of the early ideas for AARNet was to just take over the phone bills and keep ACSNet running – Ken McKinnon, the chair of the AARNet committee, scoffed at this and said he wanted a "real network, just like the NSFNET!"
>
> For a brief period a form of email address was used, in which the uucp email address was a route statement in which nodes were separated by exclamation marks. For example, in the following:
>
> ```
> {uunet,ubc-vision,nttlab, mcvax, ukc}!munnari!csc.anu.oz!gih900
> ```
>
> if your machine knew how to get to any of the bracketed set of servers, they could pass the message to munnari, to ANUCC, and to gih900's mailbox.

In August, 1983 while at a DECUS (Digital Equipment Corporation User Society) meeting in Auckland, NZ, I was able to string together a series of network links to the DEC Field Service modem resulting in an anonymous message being printed on the operator's console of the Faculties KA-10. It caused great consternation, and the operator (Inta Skriveris) instantly logged me off.

11.5 Australian Academic Research Network (AARNet)

Roger Clarke reports that the top level domain .au was delegated to Robert Elz in 1984 and that in 1989 Geoff Huston was transferred to the AVCC (Australian Vice Chancellor's Committee) as Technical Manager of AARNet (the Australian Academic Research Network). ANU's Robin Erskine chaired a Working Committee which did a lot of the lead-up technical work.

> **Geoff Huston:** I think Roger has got that wrong – it was around 1988 or so as we had informally used 'oz' for the preceding years. It was not until 1988 that we learned of the ISO3166 committee and their two letter coding convention and the assignment of 'au' to Australia. Without a 'live' network and a thing like the DNS it did not matter what we did!

Connection of AARNet to the Internet via a 56Kbit/sec satellite link to Hawaii from the University of Melbourne Computer Centre which at the time shared the Richard Berry Building (now the Peter Hall Building) with the Department of Computer Science. Geoff Huston brought back the Proteon Router from the University of Hawaii (UH), and it was connected via an OTC (Overseas Telecommunications Commission) leased line to its counterpart at UH. Robert Elz and Chris Chaundy took care of that connection. In August 1989 ANU purchased a CISCO Hybridge and connected it to a CISCO at the Universty of Melbourne via a 48Kbit/sec link. All academic institutions were connected to AARnet by May 1990.

> **Geoff Huston:** Peter Elford joined AARNet at the start of 1990, and in May 1990 we completed the effort to join all Universities and CCAEs to AARnet with leased lines. At the start AARnet was a multi-protocol network – it supported TCP/IP, DECnet and X.25. This was modelled on the NORDUnet [a collaboration between the research and educational networks in the Nordic countries] multi-protocol network. The multi-protocol approach was intended to provide a pragmatic answer to the protocol wars at the time.

Robin Erskine recalls that in the mid-to-late 1980s there was a debate among Australian universities as to the relative merits of the British Coloured Book protocols[7] versus TCP/IP. The Carrs Report from UQ argued in favour of Coloured Book but the AVCC were aware of dissenting views from

[7] https://en.wikipedia.org/wiki/Coloured_Book_protocols

their computer centre directors and wanted not to make a mistake. Robin Erskine was asked to advise and recommended going multi-protocol. Ken McKinnon, Vice-Chancellor at the University of Wollongong played a critical role in taking this decision and getting AARnet off the ground. Once the decision was taken, AARnet purchased 37 routers and 37 communications lines, and Geoff Huston and Peter Elford criss-crossed the country getting them all installed and operational in a single month – May, 1990.

> **Geoff Huston:** DEC was a part of this. They proposed the donation / deep discount of a VAX 11/750 to each campus that signed up for SPEARnet (South Pacific Education and Research Network). That used Telstra's X.25 network (AUSTPAC). Each institution would have had to pay its Telstra bill and that was a real problem. The idea was to piggyback along DEC's efforts to transition via Coloured Book to the OSI (Open Systems Interconnection) model.
>
> The Queensland universities (led by Alan Coulter) were keen to be the focal point of a national network à la NSFNET in the US (which had started in 1988) and, as their efforts were largely rebuffed by the computer centre directors, they successfully proposed that the Vice Chancellors fund a 'requirements study'. When this was released in 1988, and recommended using Coloured Book, the computer centre directors wrote a dissenting commentary.
>
> There was a networking workshop at the University of Sydney in December 1988 that included many of the players at the time and we worked on a different model that cost a whole lot less and used new fangled 'routers'. Ken McKinnon's committee decided that their direct step was to hire a technical manager to work through this mess and I took the job, initially on a two year contract seconded from ANU, at the end of January 1989, encouraged by Robin Erskine.

George McLaughlin was seconded to the Australian Vice Chancellors Committee (AVCC) in 1994 to oversee the CASMAC (Core Australian Standard for Management and Administrative Computing) project and, after AARNet 1 was sold to Telstra in 1995, the AVCC asked him to manage the subsequent networking arrangements. The panel on Page 249 tells the story.

Interested readers are referred to the AARNet history written by Glenda Korporaal which is available online at `http://mirror.aarnet.edu.au/pub/aarnet/AARNet_20YearBook_Full.pdf`

The story of ACTEIN

The Australian Capital Territory Education Information Network (ACTEIN) program is a local university initiative to introduce the Internet to primary and secondary schools in the Australian Capital Territory. The physical connectivity is based on low cost accessible technologies, and most schools use IP dial-up as their access to the Internet. The program's main direction is not the provision of Internet access itself, but in attempting to address the issue of how the Internet can be put to work in the classroom, consequently the ACTEIN Program has a strong emphasis on technical and training support to accompany the Internet access.

From: Michele Huston, *Introducing the Internet into Australian Schools. A Case Study : The ACTEIN Program.* Paper presented at INET 95. `https://web.archive.org/web/20160103082849/http://www.isoc.org/inet95/proceedings/PAPER/202/html/paper.html`

11.6 The Internet Comes to ANU

The ANU played a major role in AARNet – Deane Terrell, Robin Erskine, George McLaughlin, Geoff Huston, and Peter Elford all played major roles, and ANU was generous in its support of their activities. AARNet was located on the ANU campus at the CSC from inception until its sale in 1995. The ANU supported a ground-breaking program to extend the Internet to local schools as part of a larger ANU community outreach program. The network was called ACTEIN (the ACT Education Internal Network) and it was run by Michele Huston, Geoff's wife. Geoff set up a network domain in which each school had a sub-domain, `schoolname.act.edu.au`.

To the great sorrow of all who knew her, Michele died from melanoma, in 2010, aged only 53.[8]

[8] `https://www.potaroo.net/michele-service/`

As noted elsewhere, DCS operated a local Ethernet from the very early 1980s. In November, 1987 an order was placed for the installation of a fibre optic Ethernet network to link the main research and teaching computer facilities on campus.[9] This was extended over the next few years to cover the whole campus, including halls of residence. Geoff Huston says that he and Mark Corbould arranged to put fibre into a point of entry for almost every building on campus and used a MAC bridge to contruct a campus-wide LAN.

The combination of the campus-wide Ethernet with the AARNet gateway afforded full internet access to the vast majority of staff and students. It was early days however, as evidenced by the following three anecdotes.

Software tools (specifically `named`) supporting the Domain Name Service, by which symbolic names such as `cs.anu.edu.au` are converted into IP addresses such as `150.203.161.98`, obtained data from a root name server in the United States and cached it. Unfortunately, each entry in the cache was associated with a time-to-live after which the entry was flushed, *even if* `named` had been unable to contact the root server. This had very bad consequences when the fibre-optic cable linking Australia to the US via Hawaii was snagged and cut by a fishing trawler. A cable laying ship was despatched to locate the break and fix it, but Australia disappeared from the Internet for almost a fortnight, by which time local `named`s had flushed their caches. For a few days it was impossible to use DNS names even for servers within ANU.

Another critical component of the Internet is the software (`routed` or `gated`) which communicates information about how to route messages from one node to another. By exchanging routing protocol messages, one computer on the internet can work out how to send messages to another computer many network hops away. For security, it is vital that a server only accepts routing information from a trusted source. A gap in such security led to a funny incident involving an academic couple. One worked at ANU and the other at the University of Canberra (UC). They shared a home computer (let's call it HC) which ran a routing service, and each of them in turn connected to their University. While connected to UC, the router on HC learned how to get to the external world via the UC gateway, and somehow assigned a very low cost to that route. It then connected to ANU and told ANU's machines that the best way to the external world was via HC and the UC gateway, which of course couldn't be reached because the dial-up connection couldn't connect to both universities at once. ANU dropped off the internet again, until the problem was diagnosed and fixed.

> **Geoff Huston:** We used RIP [Routing Information Protocol] in AARnet and the local problem was replicated nationally when an aptly named system called 'madvax' at the University of Western Australia asserted a default in RIP and pulled the network in Australia down for a couple of hours.

The third anecdote relates to AARnet charging. Each university was charged for traffic but only for data packets coming in, not for data going out. The very competent Gaby Hoffman (who very sadly died in a road accident in February 2021) carefully monitored traffic for the CSC and maintained a 'Top Talkers' website.

It seemed harmless to locate Andrew Tridgell's `samba.org` server (hardware supplied by a grateful vendor) on the ANU network. Unfortunately, although the vast bulk of the traffic was outgoing (external people downloading copies of SAMBA), every outgoing packet resulted in a TCP acknowledge (ACK) packet coming in. The volume of ACKs was so great that `samba.org` became one of the heaviest users on campus! Because the AARnet charges were not paid centrally but sheeted home to user departments, `samba.org` was quickly relocated.

11.7 Electronic Information Access Prior to the World Wide Web

The CD-ROM standard was created in 1983. Later that decade CD-ROM became very popular for distributing documents, data, and software. The Library at ANU subscribed to many CD-ROM information services before these were supplanted by online access to journals.

[9] Geoff Huston, then Head, FCU, worked on this project through 1987, doing the technical studies about fibre, repeaters and bridges. DEC project-managed the field work.

Much earlier than the World Wide Web, people at ANU were able to access *Usenet News*. Usenet is a discussion system, allowing users to post articles and read them with a newsreader program which kept track of which articles you had read. Content is hierarchically organised into groups, e.g. `rec.arts.movies`. There is no central server. Instead, distributed servers contact each other at intervals and exchange articles. In a faint foreshadowing of the huge volumes of vitriol distributed by social media in 2021, early Usenet users began to *flame* each other's posts, leading to so-called *flame wars*.

The File Transfer Protocol (*FTP*) was first specified in 1971, and the first TCP/IP version in 1980. Once ANU was connected to the Internet it became possible to download files from elsewhere using FTP over TCP/IP. A search engine for FTP archives (*Archie*) was developed by Alan Emtage, at McGill University. Once a month or so, Archie would request a file listing from each of the FTP servers it knew about. It kept these listings in local files. Search requests were handled by `grep`ping the files.

Although Tim Berners-Lee wrote the specification for the World Wide Web in 1989, several years passed before the propagation of the Mosaic browser and the CERN `httpd` daemon led to the development of the first websites.

In the meantime an alternative system called *Gopher* was released by the University of Minnesota in mid 1991. Gopher servers managed a hierarchy of documents, accessed by menu, and could link to other Gopher servers.

Veronica was a search engine for documents held on Gopher servers. It was released in November 1992 by Steven Foster and Fred Barrie at the University of Nevado at Reno. It maintained a database, constantly updated, of the menu items on thousands of servers.

Tony Barry, ANU's Deputy Librarian, played a pivotal role in providing network access to ANU information. To quote from the ANU People-Timeline[10]

> Tony Barry ... was integral to the ANU Library becoming the first Australian Library catalogue available on AARNet in the early 1990s, and to the development of ELISA (Electronic Library Information Service at ANU) in 1995. He developed the first Australian University Library gopher server, ..., and helped establish the University's first Campus Wide Information Service using gopher technology. This included the first interactive directory of web accessible Australian libraries and Australian electronic journals.
>
> ANU staff member, Lisa Bradley, remembers Tony Barry being spoken of as the 'crown jewels of library technology' at a major library conference closing.

11.8 The World Wide Web at ANU

ANU has an odd connection to the World Wide Web which predates the Web itself. At WWW8 in Toronto, when I was lunching with Tim Berners-Lee,[11] he told me that his mother, Mary Lee Berners-Lee,[12] had spent several happy years at ANU's Mt Stromlo Observatory. He then asked me what he should talk about in his keynote the next day. As a web-newcomer, I was no help at all.

ANU was home to several web pioneers. To quote Roger Clarke's *Morning Dew on the Web in Australia: 1992-95*,[13] "The evidence is clear that the first web-server in Australia was David Green's Bioinformatics site, at `life.anu.edu.au`. The server-software was installed by Gaby Hoffmann ..., and the site-content was established by David Green in mid-1992. It was therefore among the first 20 sites in the world."

At the time, people were very excited by the potential of the Web, but very little of that potential had been realised. In 1993, web browsers started to support images, and in November of that year the Cambridge University Computer Laboratory set up the first webcam, showing their coffee ma-

[10] `https://itservices.anu.edu.au/_resources/_temporary/people-timeline-print.pdf`

[11] Stop name-dropping, David! (He came and sat next to me, and I didn't know who he was until others at the table asked him what it was like to be among Time magazine's 100 most influential people of the 20th century!)

[12] `https://blogs.bl.uk/sound-and-vision/2018/01/mary-lee-berners-lee-the-joy-of-programming-and-equal-pay.html`

[13] `http://www.rogerclarke.com/II/OzWH.html`

chine.[14] This site became very popular, and the Lab was petitioned to keep the lights on at night, so that it could be viewed by people around the world. When the webcam was switched off in 2001, it made the front page of *The Times* and *The Washington Post*.

Bob Gingold recalls that in 1993 David Green (an undergraduate mate of his from Monash) was located both at ANUSF and RSBS (Research School of Biological Sciences). He built ANUSF's web. ANUSF distributed NCSA (National Center for Supercomputer Applications) software Mosaic, Telnet vis tools etc on floppies around the country until internet access made that redundant.

Other early Web pioneers at ANU listed by Roger Clarke include Tony Barry (Deputy Librarian), Michael Greenhalgh (Art History)[15], Paul Thistlewaite (see PASTIME Project on Page 178), and Chris Johnson (online coursenotes for a DCS unit).

11.8.1 Mental Health Interventions Via the Web

Helen Christensen and Kathy Griffiths (full disclosure, she's my wife) of ANU's Centre for Mental Health Research (CMHR) used the power of the Web to deliver online mental health interventions and to conduct research studies at a scale which would have been totally impractical using conventional methods. Web-delivered services offered by the Centre over the years include: *MoodGYM*, *BluePages*, *MHGuru*, and *e-Couch*. The operation of these services for the Australian population has been funded by the Australian Government since 2007. In 2015, the Centre for Mental Health Research was rolled into the Research School of Population Health. The operation and further development of MoodGYM and MHGuru was transferred to the Research School of Psychology and later spun off into the *e-Hub Health* company.[16].

MoodGYM, publicly launched in 2001, was the world's first web-delivered mental health intervention. A very large scale randomised controlled trial reported in the British Medical Journal[17] demonstrated significant mental health benefits and global take-up was huge. In 2019, the Group of Eight universities featured MoodGYM in a publication highlighting the high-impact research done by its members. ANU also reported it in the Research Impact section of its 2018 ERA return. To quote from that return:

> MoodGYM is an interactive, automated online program designed to prevent or reduce symptoms of anxiety and depression. It is available free of charge to all Australians, providing a low cost alternative to conventional face-to-face therapy. It is a vital service to people in rural and remote areas and others such as young males who are unlikely to access conventional therapy. It has been translated into five languages: Chinese, German, Norwegian, Dutch and Finnish. In 2016 in Australia, 25% of MoodGYM users have been referred by their GP. MoodGYM has users from every country in the world with more than a million people having used MoodGYM as part of their recovery from depression.
>
> ...
>
> New registrants for each of the years from 2011-16 respectively were as follows: 99,778; 105,478; 107,147; 97,073; 138,999; 138,313.

Behind the Scenes at MoodGYM

The very first version of MoodGYM was developed by an external company who modified a shopping cart system. A second company made some improvements but problems persisted.

CMHR asked Robin Stanton for help, and he passed the request onto MISD, where Fay Gibbons asked David Berriman to look into the problems. David installed MoodGYM locally and looked at what would be needed to get it working. He decided that the delivered product was a lost cause and that it should be rewritten from scratch. The result was written in Microsoft ASP (not ASP.Net, but its predecessor) – this used VBScript (Visual Basic) for the programming language and ran on IIS

[14] See `https://www.cl.cam.ac.uk/coffee/qsf/cacm200107.html`.

[15] Artserve: `http://rubens.anu.edu.au` – still online as at 01 Nov 2020, after more than 27 years.

[16] `https://ehubhealth.com/`

[17] Christensen, H., Griffiths, K. and Jorm, A. F. (2004). Delivering depression interventions using the internet: randomised controlled trial. *British Medical Journal*, **328**, pp. 1–5.

(Microsoft's web server), with a Microsoft SQL Server as the backend. This is the same technology that MISD was using at the time for its admin systems (ISIS, HORUS, APOLLO, etc.).

John Richards (DVC) granted funding for the rewrite and Robin Stanton offered David Berriman's team to do the work. A programmer named Richard Pass was recruited to do most of the ground work, with David acting as project leader and liaison with CMHR.

This rewritten version of MoodGym ran on the admin infrastructure. It initially had all the functionality to display and capture all the mental health content, surveys, navigation and login details (no trial identification, subject selection or recruiting) ... and also had a very basic reporting feature (CVS export for Excel and export files for SPSS). Unfortunately, this was very slow to run since the database had been designed to streamline data capture – this was an oversight that was only partially fixed by additional indexes and tuning.

In 2003, Anthony Bennett was recruited by CMHR to develop the framework for a new online intervention called *e-Couch*. Says Anthony:

> The e-Couch framework was intended to act as a flexible system for delivering structured content and,"it should be able to run Moodgym", was one of our thought processes to make sure the design was on the right track, although implementation of Moodgym on this framework wasn't actually something that was planned or intended from the outset.
>
> I first started designing this framework when I joined CMHR in 2003, however due to the enormous scope of the content that needed to be developed, for a time I worked on other projects and it was not until 2005 that we really started working in earnest on e-Couch. It was at this point that David Tulloh joined the team, and he was also responsible for elements of the system design, as well as implementing a number of core libraries on the framework. He departed in 2006, and I continued the development. The first version of e-Couch went live in 2007, and it was only after this that we began working on implementing Moodgym on the framework.

The new framework was written in PHP with a PostgreSQL database.

A planning workshop held at Lake Crackenback Resort made it clear that CMHR online systems were going to be used in multiple trials run by multiple universities in Australia and overseas. A key requirement was thus to build a framework for running experiments which was tightly integrated with the content delivery.

At some point, CMHR took over hosting of its online interventions from MISD (which I think was by then called CIS). Both David Berriman and Richard Pass had left, none of the CIS staff had the expertise to make changes or adjustments, and stability issues had begun to manifest themselves.

Anthony again:

> The e-couch framework was extended in a number of ways to support research trials and program translation. We designed and implemented an interface to allow program content to be translated directly, a tool for generation of replicable randomisation schedules and various other libraries that aided in the delivery of automated online research trials. Alellie Batterham was especially involved in much of this work.
>
> I'd also note that from the technical side, there were actually a lot of people involved in the implementation of the framework and the two sites. Beyond the straight up coding by those already mentioned, there was of a huge body of work (far more than the coding I would guess) implementing all the content and program flow. People who made a significant contribution to these elements include: Ada Tam, Jolanta Samoc, Kylie Bennett, Lisa Olive and Sandra Lauer.

David Berriman joined CMHR in January 2008 (after returning from extended sick leave due to bowel cancer). At this time the framework was essentially working and many trials were in progress, but reporting and data analysis was still problematic – researchers had to request data be extracted and Anthony did this manually and passed it on to them. David was tasked with building an admin interface that would allow researchers to log in and maintain and report on the trials they were responsible for. Before joining CMHR he had been with the Australian Partnership for Sustainable Repositories (APSR) – now the Australian Research Data Commons (ARDC). While there, he and James Blanden had developed a web application framework in PHP and PostgreSQL (Communications Synthesis Infrastructure (COSI) framework) – so he used this for the MoodGYM admin interface. This technology was used up till the most recent rewrite of MoodGYM, which now uses the open source framework Symfony (still PHP with PostgreSQL database).

11.9 CNIP: Centre for Networked Information and Publishing

CNIP was set up in 1993,[18] jointly by the Computer Services Centre and the Library. Mark Corbould and Tony Barry were co-heads. Other staff members over the years included Allan Williams, Monica Berko, Sam Hinton, Michele Huston, Melanie Rooney, and Gavin Mercer.

The IT Services Annual Report for 1996 provides insight into the operations of CNIP.

1996 IT Services Annual Report: CNIP activities during the year.

The Centre for Networked Information and Publishing (CNIP) worked with other areas of the University to provide publicity and communications services as part of the University's support of major community activities including IFIP 96, the Australian Science Festival and the Australian University Games.

CNIP, in collaboration with the ACT Department of Education and Training, provided teacher training, support and internet services to ACT schools under the banner, Canberra Schools on the Net (ACT Education Information Network – ACTEIN). By year end CNIP was supporting 98 schools and colleges in the ACT. Network Services continued to provide support for a bank of 48 modems providing dialin connectivity for this project.

The Australian SunSITE (Software Information and Technology Exchange), which is located within CNIP, continued to develop as an ftp archive and mirror site and community publishing web site focussing on Australia and the Asia Pacific Region. The most popular web sites on the SunSITE were Questacon (the National Science and Technology Centre) and the South Pacific Information Network. At the end of 1996 there was 550MB of information published on the SunSITE web server and on average 80,000 documents were being delivered each week. Notable sites established on SunSITE in 1996 include:

- Archeology in Australia
- RSPCA
- International Council for Science Education
- Federation of Engineering Institutions of South-East Asia
- Pacific Network for Sustainable Development
- Institute of Public Administration
- National Shelter.

A SunSITE objective is to provide the Australian community with a better level of service than is often available when accessing offshore servers. As 95% of accesses were from Australia it appears to have met this objective. The SunSITE ftp archive holds 40Gb of software and information and is accessed on average 35,000 times (involving the transfer of some 23Gb of data) per week.

CNIP also coordinated the University's web publicity for the 1996 open day and provided a home site for the combined ACT universities open day publicity. An IT Services poster was developed to hand out at Open Day and subsequently distributed throughout ANU.

Number of accesses to campus networked information servers

CNIP operated two major web publishing servers in 1996 – www.anu.edu.au and SunSITE (see above). The server www.anu acts as a general web publishing service for the University community. It currently hosts 6.8Gb of data and delivers approximately 2.8Gb of data per week (internally, nationally and internationally). This service saw significant growth in 1996. From delivering some 137,000 documents per week in January demand grew to 350,000 per week by year's end.

[18] To the best of my knowledge – I was unable to find the relevant annual reports.

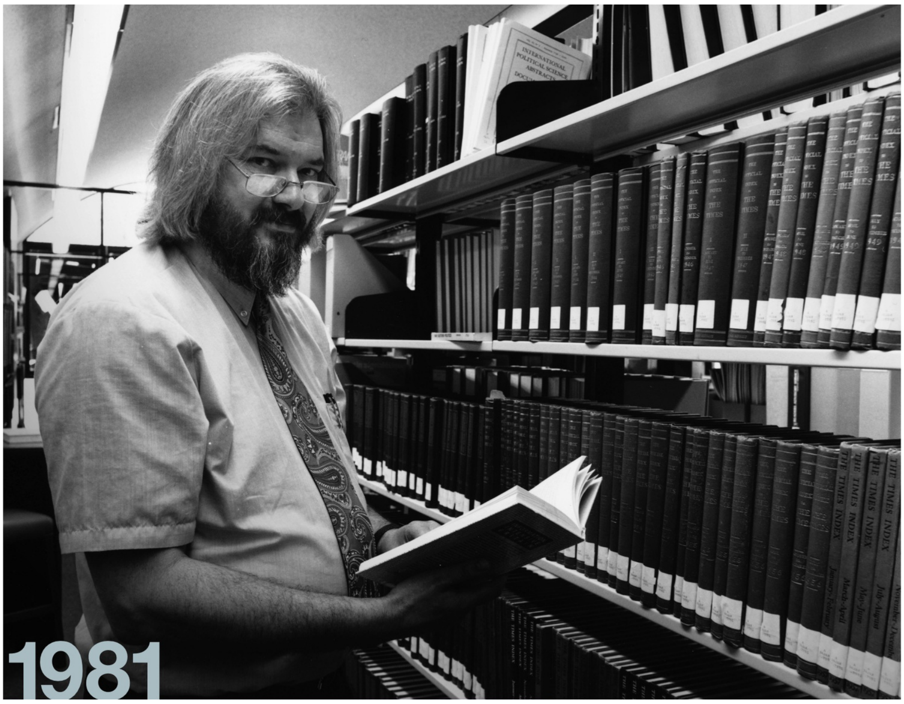

Tony Barry, ANU Deputy Librarian. *Photo extracted from* **`https://itservices.anu.edu.au/_resources/_temporary/people-timeline-print.pdf`**

After returning to Canberra after more than five years away, Melanie Rooney worked part-time for ANU on a range of web projects:

- CNIP,
- ANU Graduate School,
- Chifley Library Outreach Program,
- Research School of Chemistry.

CNIP was abolished in 1997, along with the IT User Support group, in the context of a substantial cut in recurrent funding to IT Services. Many staff opted to take a redundancy package.

11.10 Eduroam

Eduroam is a mechanism which allows members of an educational institution to access the internet using the `EDUROAM` wireless network at another participating institution. Access requests are authenticated using an authentication server at the person's home institution. The person attempts to connect using a fully qualified login, such as `u199999@anu.edu.au`. From the domain name, the Eduroam service works out how to connect to the authentication service for that domain and forwards the username and password to it.

Institutions operating an `EDUROAM` wireless network typically allow authenticated users to set up a virtual private network (VPN) to their home institution. Eduroam may be of tremendous benefit to people in circumstances where mobile phone service is poor, unavailable, or expensive. ANU is an Eduroam participant and has been for many years.

11.11 Network Consolidation

In the late 1990s, Robin Stanton was PVC (Academic) with responsibility for the Division of Information, including the Library, MISD and CSC. At the time, there were four independent networks on campus:

1. Ethernet, by now servicing every area of campus;
2. The administrative FNA coax network, linking administrators around the campus to the Admin. FACOM;
3. The telephone network, comprising hundreds of pairs of twisted copper cable; and
4. The television network, allowing distribution of lectures and other teaching materials around the campus.

The telephone network was controlled by the Head of Facilities and Services (F&S), Warwick Williams, who had previously been University Secretary, and was administered by Murray Napier. F&S had also started to build a security monitoring network.

By the early 2000s it became clear that significant expenditure would be needed on ANU's network infrastructure. In particular, ANU's main PABX (telephone exchange) was lacking in capacity and needed replacement, the campus backbone network based on ATM (Asynchronous Transfer Mode) switches needed upgrading, and campus telephones were supported by thousands of copper wire pairs individually connected back to the PABX in University House.[19]

At this time, network "convergence" was an industry trend, and Robin Stanton championed it in the form of a "common IP fabric". John McGee was recruited as Head, Networks and Communications in around 2001.[20] John came from the Snowy Mountains Authority (SMA, later known as Snowy Hydro) and had set up an integrated communications network, "from Cooma onward", including within 330 KV switchyards.

John found that ANU actually had more than ten separate PABXes in addition to the main one at University house, each of them locally managed. They were in halls of residence, as well as at MSO and Siding Spring observatory. After F&S were persuaded of the merits of convergence, John wrote a Request For Quotation (RFQ) for a Voice-Over-IP (VOIP) system capable of supporting 10,000 extensions. At the time CISCO was the only established VOIP player, but Lucent (itself spun off from AT&T in 1995) had spun off Avaya and Avaya had a VOIP product. Contrary to the, "no-one got fired for buying CISCO", wisdom, ANU opted for Avaya as the supplier, even though support for 10,000 extensions was only provided in a then unreleased version of its software.

Then there was the question of ANU's network fabric. At the time there were 130 separate locally controlled network entities, e.g. Ethernet switches, and there were 100 separate email domains. John and his team set about upgrading the network, trying to engineer an end to the local fiefdoms, by providing a faster, more reliable service. Of course DCS (with specialised needs) was a hold-out and John tells me that Bob Edwards kept his team honest, and was a great help in testing out campus-wide support for IP version 6.

ANU's new network had a CISCO core in a redundant mesh configuration with 1 Gbit/sec interlinks. At SMA, John had been asked to achieve no more than 30 seconds network downtime in a year, and his team aimed for a similar reliability level at ANU. The challenge with achieving reliability is not only creating redundant links and equipment but in creating robust mechanisms to switch-over in the event of failure of a node or link.

Consolidation was eventually successful, though several administrators were terribly concerned that student and admin data would share the same network.

A benefit of the new consolidated model is described in the Natural Disasters section on Page 253.

Between 2003 and 2005, another upgrade was undertaken. The CISCO switch/routers remained as the Core Network – with their capacity extended to support multi-1GigE capacity and 10Gbit/sec interfacing to AARNet via new and additional circuit boards. However, the tender for the Access Network replacement was won by Enterasys (in 2021 owned by Extreme Networks after a number of fraud cases in the USA). Their switches met all requirements and offered the lowest life cycle price.

Redundant cross-campus links were now 10Gbit/sec, and by mid-decade, ANU was connected

[19] I was amazed that when the CS&IT building was built in 1994/5, three hundred copper wire pairs were laid across campus to support its telephones.

[20] Originally Head, Communications until a marketing person was appointed with the same title.

to AARNet by three 10Gbit/sec links. Each building was provided with fibre verticals and structured copper wiring to invidual ports. Every port supported 1Gbit/sec, including all the rooms in halls of residence. Note that this is 40 times faster than the 2021 standard NBN connection! All told ANU had 12,000 gigabit network ports!

In addition to this, 700 wireless access points had been rolled out by the time of John's departure from ANU in 2009. When I say his departure from ANU, I mean the end of his employment contract with ANU. He physically left ANU in 2005, relocating to Tasmania and working remotely. The Avaya phone on his desk was, of course, reachable via an ANU extension.

John led a team of 12 to 14 people, covering networks, security, telephony, physical infrastructure, and support. His office was right next door to ANUSF in the Leonard Huxley Building and his team became responsible for supporting ANUSF infrastructure, down to the level of modeling airflow within the computer room. He is fulsome in his praise of the people who worked for him: Geoff Barlow supporting ANUSF hardware, Darren Coleman as network engineer, Gaby Hoffmann's brilliance in spotting suspicious activity on the network (scanning 100 times per second!), Markus Buchhorn on digital futures, and Andrew Howard, who came from AARNet.

Darren Coleman was called to evaluate network switches from a non-CISCO vendor and, to his surprise, found CISCO copyright notices in their code.

John McGee thinks that the setting up of ANU's Division of Information, including libraries as well as IT infrastructure, enabled good things, and cites as an example the benefits of turning library ground floors into fully networked student hubs.

11.12 Research Networks

During the 1990s and beyond, researchers around the country could see enormous potential research gains through collaboration and sharing of data and computational networks over high speed networks. Unfortunately, the speeds available through available from available networks were too low and latencies too high.

This led to substantial investment of research funds into advanced networks to overcome the limitations. The ACSys CRC and the Australian Partnership for Advanced Computing (APAC) were heavily involved. Markus Buchhorn (a recent doctoral graduate from MSSSO) was a key player, as were the ACSys Directors: Michael McRobbie, Robin Stanton, John O'Callaghan, and Darrell Williamson.

Slowly but inexorably, other organisations, particularly AARNet Pty Ltd, were working hard to upgrade the performance of more generally available networks. At the time of writing, even the much-maligned National Broadband Network (NBN) is capable of supporting simultaneous High Definition streaming video to a large proportion of the Australian population.

From today's standpoint, the performance of the advanced networks described in the following subsections seems unremarkable, but at the time, they were significant advances.

11.12.1 The Research Data Networks (RDN) CRC

ANU was heavily involved in the RDN CRC – Two of the four research areas were joint projects with the ACSys CRC. The RDN CRC ran from 1993–1999. See Page 182 for more detail.

11.12.2 The Australia-Japan Network Link (AJNL)

The AJNL was a direct frame relay service between Australaia and Japan, established in January 1998, by Optus and KDD. Initially the connection was rated at 768kB/sec.

The aim of AJNL was to improve support for collaborative research between the two countries by overcoming the congestion and high latency of routes via the USA. The Australian end of the connection was managed by the ACSys CRC. (See Page 178.)

11.12.3 The Asia-Pacific Advanced Network

With enthusiastic support of Australian researchers and from the Japanese Ministry of International Trade and Industry (MITI). The AJNL was very quickly (May 1998) connected to APAN.

The 1999/2000 Annual Report of AJNL lists ten separate projects making use of AJNL / APAN. One of them is the *Remote Control of Robots Project* involving Alex Zelinsky's robotics group within RSISE and the Japanese Real World Computing Program. Each group was able to demonstrate the ability to control a robot at the remote site.

11.12.4 GrangeNet

The high speed GrangeNet research network was launched in 2002 by a consortium comprising AARNet, the Australian Partnership for Advanced Computing, Cisco and PowerTel.

2002: SMH article anouncing launch of GrangeNet.

> Canberra-based Grangenet's capacity of one gigabit per second makes it Australia's fastest public network.
>
> It has nodes in Canberra, Sydney, Melbourne and Brisbane. Tasmania, South Australia, the Northern Territory and Western Australia are connected via AARNet, the network that belongs to the nation's universities. Each node has access to a supercomputer-level research facility.
>
> Grangenet itself employs the advanced IPv6 network protocol; special routers convert conventional IPv4 traffic.

Sydney Morning Herald, 22 October 2002. `https://www.smh.com.au/technology/network-paves-fast-track-for-research-20021022-gdfqqe.html`

Chapter 12

Women in ANU Computing

Around the time ANU's first computers were arriving, Australian society was characterised by overt gender discrimination in employment opportunity and salaries. In the scientific world, there is a long history of undervaluing the contributions of female scientists. More recently, there has been general recognition of the need to end discrimination and to appropriately recognize female achievement.

This chapter starts with an attempt to recognize the contributions of women in the early days of ANU computing, and the employment discrimination they faced. It then looks at the gender balance among students and staff in DCS across the years, and among IT staff at ANU. It concludes with a summary of commitments made by CECS to achieve demographic representational parity.

12.1 The First Women in ANU Computing

Despite discrimination and gender stereotyping documented in this section you will see some evidence that ANU was less discriminatory than other employers.

Claire Wehner working at Mt Stromlo in the 1940s. *Photo via* **`https://services.anu.edu.au/campus-environment/trails-tours/mt-stromlo-heritage-trail/a-home-on-the-hill`**

12.1.1 Claire Wehner and the IBM 610

There is no doubt in my mind that the first woman to use an ANU computer was **Claire Wehner** (née Beech), in 1960. She may even have been the first ANU person to use an ANU computer. Claire had been working at Mount Stromlo Observatory (MSO) as a Scientific Assistant since December 1947. On arrival she had been given a desk, a slide rule, a book of mathematical tables, and astronomical calculations to make. Around 1949 she graduated to using a simple hand-punch calculator. Some time after that she took a year's leave in England, where she worked on Maurice Wilkes's Electronic

Delay Storage Automatic Calculator (EDSAC). Quoting from the written record of an ANU Heritage Office Oral History Recording of an Interview with Claire Wehner:[1]

> Worked in a maths laboratory at Cambridge University as one of a small group of scientific assistants, assisting the students who were building EDSAC, the first big computer built in England;[2]
>
> The four women scientific assistants doing this work would manually check the computer's calculations and outputs after it had been programmed, to make sure that the computer program was working correctly; the four women would all sit together around a big table in what had been the University's medical faculty to do their manual calculations.

On return to Stromlo,[3] Claire started working on the Commonwealth Time Service, taking comparative observations on stars, and performing calculations. When she married in January 1955 she was forced to resign her permanent job with the Department of the Interior (which ran MSO at that stage).

Claire Wehner adjusting for the Time Service in the 1950s. *Photo courtesy Mt Stromlo Archives. Via* ***https://xnatmap.org/adnm/ops/prog/Orroral%20Obs/astro_geod%20orroral.htm***

She was immediately rehired as a temporary employee on the same salary but lost her job security and her superannuation entitlements. When MSO, by then part of ANU, acquired an IBM 610 computer, Claire used it for the Time Service calculations she had previously undertaken manually. (See the photo of Claire and the IBM 610 on Page 27.) Quoting again from the oral history record:

> ... the Observatory acquired its first computer; it was pretty elementary, about the size of a piano, and was programmed by and produced outputs on punched paper tape; an electrical keyboard created the punched paper tape inputs;
>
> This computer did the calculations that CW had previously done manually; decoding the punched paper tape was done by looking at the pattern of noughts and ones;
>
> Eventually three people were able to operate the computer – CW and two other male scientific assistants;
>
> The computer was made by IBM and CW remembers many emergency late night/early morning calls to the IBM technician to handle problems;
>
> To meet deadlines for papers and material for Stromlo scientists for the International Geophysical Year,[4] CW had to run the Stromlo computer non-stop for 48 hours; she then had to go to Sydney to the big IBM computer there (a machine the size of a warehouse) and had to use the computer at night after its normal work was over; had to work every night for a week just to get through the required work;

[1] `https://openresearch-repository.anu.edu.au/handle/1885/117118`.

[2] The use of the Colossus computers at Bletchley Park from 1944 – 1945 was kept secret until the 1970s.

[3] Claire had returned to MSO by 1952.

[4] It must have been for some other big occasion because the International Geophysical Year was 1957/8 and the computer didn't arrive until 1960.

> The ANU's Department of [Theoretical] Physics also acquired a computer, which was larger than the Observatory computer, so she would sometimes go down to the campus to use it, driving down in the old Chevrolet even when pregnant;
>
> CW continued work until she had her first child, Martin, on December 7, 1963. She stayed at work until her waters broke and she had to be taken to hospital: there was no provision for maternity leave for her;
>
> This was more or less the end of her work for the Stromlo Observatory: although the Director, Bart Bok, asked her to do some work writing a computer program at home, it proved too hard to manage work and the baby at the same time;

Claire Wehner returned to IT work at the ANU Library in the mid-1970s and worked there until at least 1993. I understand she died in April 2020.

12.1.2 Women and the IBM 1620s

As noted in Chapter 2, when ANU acquired an IBM 1620 in 1962 (located on the ground floor of the Oliphant Building), **Elizabeth Reid** was the first operator/programmer and **Margaret Campbell** was one of the first operators. Elizabeth's subsequent storied career is well documented.[5] Unfortunately I have not managed to discover anything about Margaret's subsequent career. **Lorrel Sherar**, later known as Ann Apthorp, joined the facility as an operator on 12 August, 1964, on a salary of £1037 p.a, after doing well on an aptitude test. She arrived just after the upgrade of the 1620 from a Model I to a Model II.[6]

When **Barbara Davidson** joined the facility in early 1965, she says that Lorrel was the only other female staff member there. The two of them attended a course together, and Lorrel was soon promoted to Programmer. Lorrel (now Ann) says:

> I was offered a move to ANU's Management area and to be classified as a Programmer. Prof Le Couteur told me to accept it! My salary jumped up to £1848. ANU was using Unit Record Equipment in the Management area. Data was on punched cards, but sorting and collating and subsequent reports were done on various machines that had large panels that were actually 'wired' to perform what was required. Wish I had taken some photos of them![7]
>
> I have nothing but wonderful memories of my time working at ANU. We had the opportunity to be introduced to some amazing people and can say we worked in the same building as people like Prof Le Couteur, Prof Titterton, Sir Mark Oliphant and many budding scientists eg John O'Callaghan and Ken Muirhead. Early learning, prior to the pending arrival of the IBM 360, was via computer room books (Daniel McCracken etc.), talking to the engineers and mostly gleaning programming information from anyone else who was using the computer! Then along with the 360 introductions in Sydney, ANU organised short courses with IBM in RPG[8] and PL/1 for some of us. I never looked back....and ended up introducing PL/1 into Hong Kong!
>
> I do consider that I was just lucky to have been in the right place at the right time ... and ANU gave me the 'opening'.

[5] `https://genderinstitute.anu.edu.au/elizabeth-reid-transforming-adversity-impact`

[6] `https://openresearch-repository.anu.edu.au/bitstream/1885/149339/2/2.1.8.3(3)-17-1964.pdf`

[7] ANU's Psychology Department subsequently made use of card sorters and collators, which may have been inherited from this area, until the 1980s. They later passed briefly through DCS and no doubt were eventually dumped.

[8] Report Generator

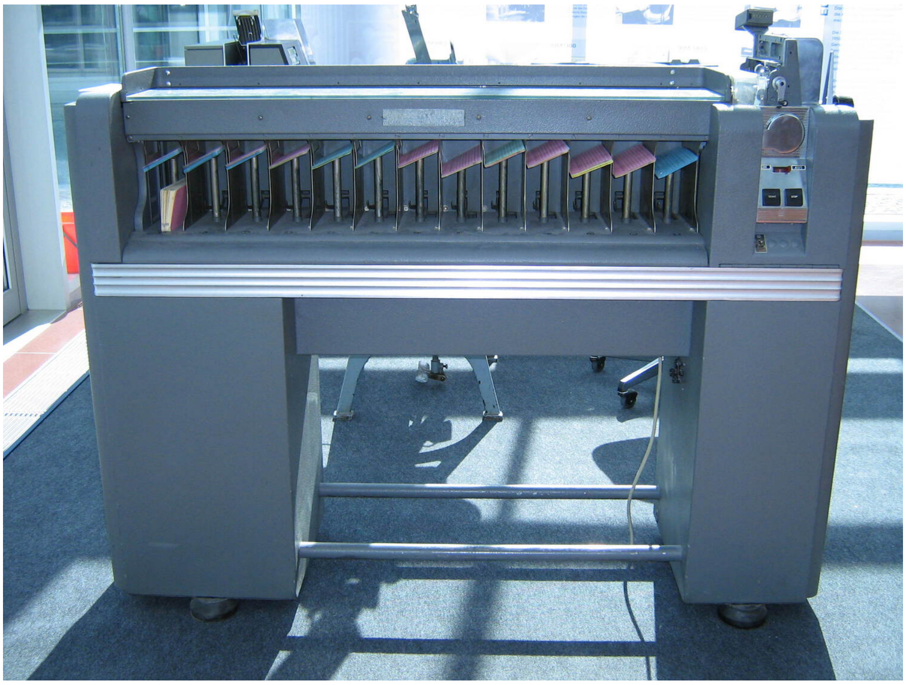

An IBM 082 card sorter, similar to one used at ANU. *Photo from* `https://en.wikipedia.org/wiki/Unit_record_equipment` *under Creative Commons 2.5* `https://de.wikipedia.org/wiki/Benutzer:waelder`

Lorrel (Ann) thinks that ANU's IBM 360/50 may have been one of the first in Australia and has a memory that the PL/1 compiler produced some error messages in German, testimony to the international nature of IBM's software development.

When Barbara Davidson arrived in Canberra from New York with her husband in late 1964, she found significant discrimination against women:

> Early in 1965, I began looking for a job in Canberra. I had been a Programmer/Analyst for General Electric for 3 years and had experience on the IBM 1620 and IBM 704. It surprised me to see the job ads in the Canberra Times: "Computer Programmer: Salary for men £4000, women £2000."[9]
>
> I considered teaching Mathematics, in which I also held a certificate. That turned out to be even less promising. I interviewed at CSIRO and they offered me a job. They apologised that I was a woman and they were "required to pay me half salary."
>
> Then the ANU computer facility began planning for their conversion from the IBM 1620 to the IBM 360. I was given a job at **full** salary. I worked for them until late 1966. I helped with the conversion. I worked with various staff and faculty to set up their jobs for the computer. I wrote documents for others to learn various processes.
>
> I was treated with respect and closer to equal than most of the women around me. There were several professors and researchers learning to use the computer, not many others actually working in the Computer Centre. I was reluctant to rock the boat and tended to follow rather than to be aggressive. I do remember a couple of instances where I made bold decisions, but most of the time I let others lead.
>
> I only remember one other woman, an Australian, Lorrel Sherar[10], working with the 1620 when I started my job. She was an operator. She was intelligent and strong minded. We travelled together to Sydney and to Melbourne for several training courses. She went on to be very successful in several international businesses.

Barbara returned to the US, and taught part time for 18 years in the then newly founded Computer Science Department at Western Washington University.

[9] Approximate amounts.

[10] Later Ann Apthorp.

12.1.3 Women and the IBM 360/50

Brian Robson recalls that at the time of the transition from the 1620 to the 360/50 there were four programmers, and that all were female. We know that Barbara Davidson was one of them, but unfortunately the annual reports for Theoretical Physics don't list general staff.

The ANU Computer Centre (ANUCC) Annual Report for 1966[11] shows that all four ANUCC programmers (BF Butterfield, PN Creasy, JF Temperly and P Tindale) were male, and all three operators (TR Martin, A Sanders and P Thompson) were female. It is likely that programmers continued to be employed by Theoretical Physics for some period after the Computer Centre commenced operation. That explanation is strengthened by the fact that the list of Academic Publications for ANUCC in 1966 include: Davidson, Barbara A. and Temperly, J. *Introduction to Digital Computer Programming using PL/1 for the IBM/360*. ANUCC Technical Report No. 3.

The ANUCC Annual Report for 1966 records that, "The computer is also used extensively by the Administration's Data Processing Unit. A typical application is the calculation of the University Pay." By 1967, the number of ANUCC operators had doubled to six: three were male , and the females were Martin, Brinkley and McKinley. Julie Rohl (later Julie Bakalor) joined as the only female ANUCC Programmer in June 1968.

The 360/50 was located on the top floor of the Cockroft Building. My understanding (partly based on personal experience) is that the 360/50 was operated on a schedule which allocated time to: academic processing, during which jobs were run by ANUCC operators; administrative processing, during which jobs were run by Data Processing Unit (DPU) operators; maintenance, during which the machine was handed over to IBM engineers; and out-of-hours, during which 'licensed' academics or research students could book the machine in blocks of two hours.

Unfortunately, I have been unable to locate any annual reports covering the operation of DPU or the Management Services Group (MSG), prior to their incorporation into the Computer Services Centre (CSC) on 26 Oct 1978. However, the ANU Reporter article *Gender No Barrier*, celebrating the achievements of Skaidrite Darius[12] sheds light on the operations of the 360/50 and on the hurdles faced by female employees.

Skaidrite Darius around the time of her retirement in 1988. ***Photo: ANU Archives and Darren Boyd.***

[11] https://openresearch-repository.anu.edu.au/bitstream/1885/149339/4/2.1.8.4%283%29-13-1966.pdf

[12] https://reporter.anu.edu.au/gender-no-barrier

Skaidrite Darius (then Skriveris) first joined ANU in 1955 as an accounting machinist in the John Curtin School of Medical Research.

> ... her boss suggested she apply to join the University's new IBM data processing team.
>
> "In those days, when they put 'male' in an advertisement it meant only males could apply — the job wasn't meant for females and I thought, I'll never get it," she recounts.
>
> Her boss continued to encourage her and, after she was told 35 men had failed the test, she was curious. She wanted to know why the test was so hard and why they had failed.

She got the job, and in her new role, she travelled to attend IBM courses in Sydney and Melbourne.[13]

> "I spent the whole year going from one place to another," she says. "The first trip was to Sydney. Downstairs was the course for the punch operators and upstairs was for the programmers."
>
> "I walked in but I was a bit late. I went upstairs and opened the door and the room was full of males. The lecturer didn't even ask my name, he just said, 'no, no, you're in the wrong place – it's downstairs, that's where the punch operators, the girls, are'."
>
> However, Skaidrite was registered to attend so when she told him her name, he was shocked she was female. This became a common scenario as the years went on.
>
> ...
>
> The first computer at ANU was on the top floor of the Physics building.[14] The entire floor was reserved exclusively for academics during the day. Administrative functions could be processed only between midnight and four in the morning.
>
> During the day, Skaidrite and her team would carry out the preparation, such as writing or checking programs and assembling the data to be processed. An operator would go in during the night to do the processing.
>
> "However, if you did anything wrong during the day, like if you misspelt one word, then it wouldn't work," she says.
>
> More often than not, Skaidrite was called in when issues arose during the processing.

I encourage you to read the full Reporter article for an account of all the discriminations and inconveniences endured by Skaidrite. A Canberra Times article, *In a little apartment in O'Connor lives a true trailblazer for women*[15] outlines her frightening history as a refugee from Latvia, and reports that ANU's payroll recorded her as Male in order to overcome barriers to female participation. However I believe there is a decade error in the sentence, "In the mid-1950s, IBM introduced computing processing to the ANU to deal with everything from payroll to stores to student results." It should read, "mid-1960s".[16]

The 1979 CSC Annual Report shows Skaidrite as DPU Supervisor with 24 staff. She retired in 1988 when her daughter became pregnant.

12.2 Gender Balance Among DCS Staff

Vicki Peterson was the only female academic in DCS for the first four years, and the most junior appointment (Tutor). Between her departure in late 1974 and her return more than five years later, the department included no female academics.

[13] I believe she transferred to DPU late in 1962. According to internal ANU telephone directories held by ANU Archives, Skaidrite's telephone number changes to the DPU one between the July 1962 list and the January 1963 one.

[14] As documented in this history, the IBM 360/50, located on the top floor of the Cockroft (Physics) Building was actually ANU's fourth computer.

[15] `https://www.canberratimes.com.au/story/6665428/in-a-little-apartment-in-oconnor-lives-a-true-trailblazer-for-women/`

[16] Brian Robson doesn't remember DPU using the IBM 1620s for administrative computing. In the Computer Unit Annual Report for 1964, "Administration" accounted for 0.46% of the useful work of the 1620, but he thinks this may have been administration of the computer. Lorrel Sherar's recollection is that MSG/DPU were using unit record equipment, not computers, until the arrival of the 360/50.

Four Professors in the School of Computing, 2021: Top to bottom: Sylvie Thiébaux, *Photo: Bob Williamson;* **Amanda Barnard,** *Photo: Amanda;* **Kerry Taylor.** *Photo Kerry;* **Lexing Xie** *Photo Lexing.*

By 2019, things seemed to have improved. Sylvie Thiébaux had become, in January 2011, the first female professor in Computer Science and the first in the College of Engineering and Computer Science.[17] Unfortunately for DCS, she has decided to take a voluntary redundancy in 2021. See her entry on Page 132.

Since Sylvie's promotion to professor several others have followed. In September 2017, Genevieve Bell became a Distinguished Professor at ANU, founding the Autonomy, Agency and Assurance (3A) Institute. See the panel on Page 185. In 2020, Amanda Barnard was appointed as a Senior Professor of Computational Science. See her entry on Page 105. Lexing Xie was promoted to Professor level in 2020, having been in the department since December 2010. See her entry on Page 139. Finally, Kerry Taylor was also promoted to Professor level in 2020. Kerry's entry is on Page 132.

These senior women in DCS, and of course Genevieve Bell in the School of Cybernetics, serve as role models and provide encouragement for talented junior females to pursue careers in computer science.

In 2007, Lynette Johns-Boast received the *WICked Woman of the Year* award, from the Canberra Women in Information and Communication (WIC) for her contribution to developing, encouraging and mentoring young women in industry and at university. See her entry on Page 118. Lynette shared some disquieting comments on gender issues in DCS.

> As a woman, it was very interesting working at the various incarnations of DCS. I had always worked in a male dominated and hierarchical environment, but the university was something completely different, that took some years to understand and to learn to cope with. DFAT was male dominated, but it was hierarchical rather than sexist. Once I moved to the IT area, it was even more male dominated than elsewhere – at one point I was the single female in a team of about 50 – I think we ended up with 4 women in the team, but I never felt there was a gender bias, I never felt disrespected because I was female, there was never any more sexism than you found elsewhere in society.
>
> Moving to the university was where this came to the fore for me and really surprised me. My overwhelming experience at ANU was that I was disadvantaged/disrespected because I didn't have a PhD, I was female and because I came from an industry/applied background. Added to that, I was concerned about education and believed that I could make the biggest difference through providing high quality learning opportunities, so largely I was irrelevant to many of my colleagues. My experience counted for nought. From a student perspective, I was rarely aware of any negative bias because I was female (until we began taking in large numbers of international students, and then there were a few from certain backgrounds who had real difficulty dealing with women). Students seemed to really appreciate that I could provide links for their learning to real-world experiences and uses.
>
> The School has and has had some fabulous individuals, and like everywhere else has had its fair share of dead wood for a variety of reasons. I'm biased because I think the whole system of higher education is broken and that brokenness is reflected in recruitment into the School. Again, I could go on for ages about this.

12.2.1 Current staff gender balance

The ANU's Athena SWAN institution application of March 2019,[18] shows that, in CECS in 2017, women made up 53% of professional staff but only 21% of academic staff – 32% overall. The same report shows that across the whole university the gender balance was substantially better. Women made up 60% of professional staff and 40% of academic staff – 51% in total. These figures include full-time, part-time, and casual staff.

Unfortunately, the most recent data on the CECS People database (as at 27 Feb 2021) has not been updated to reflect the 2021 situation for DCS. The data presented in the following panel shows the percentage of female staff at about 22%. There is a somewhat higher percentage at Level E[19], but the numbers are small enough that arrival or departure of even one female would change the percentage substantially. Sylvie Thiébaux has been kind enough to analyse the School of Computing staff list from late March, 2021.

[17] Australia's first female computer science Professor was Jennie Seberry, appointed as a full professor at ADFA in 1987.

[18] `https://www.sciencegenderequity.org.au/wp-content/uploads/2019/09/Australian-National-University.pdf` Note that there appears to be a missing "Research Only Academics" line in the CECS part of the table on Page 15.

[19] After the departure of Sylvie Thiébaux in early 2021, three out of ten professors in the new School of Computing will be female.

Gender statistics for DCS staff

- *Counts derived by Sylvie Thiébaux from official School of Computing staff list as at 30 March 2021.*
- *Includes some people who are in the process of leaving.*

Level	Count	Female Count	Percentage Female
E	11	3	27
D	9	3	33
C	14	2	14
B	33	9	27
A	2	0	0
All	69	17	25

12.3 Gender Balance Among Computer Science Students

The low representation of females in DCS over the years has been a concern. In 1976 it was reported (see panel on Page 58) that only 22% of student enrolments were female. Statistics for 2018, passed on to me by Alistair Rendell, show that, although total student load increased by a factor of 14 (90 to 1273 EFTSL) between 1976 and 2018, the female proportion barely changed – 22% to 24%. This is despite strong efforts to encourage female enrolments.

Two panels on Page 279 use Commonwealth Higher Education data to compare the recent situation with that 20 years ago. The analyses break down the data by completions v. enrolments, and domestic v. overseas students. The growth in the proportion of overseas students is dramatic, while the drop in domestic student completions seems concerning as does the drop in female proportion of domestic student completions.

Randy Olsen presents a nice graph of percentage female graduations[20] versus year, for a number of different disciplines in the US. For Computer Science it shows a rise from 1970 to 36% in about 1983 followed by a decline to below 20% in 2012. I don't know whether that rise and fall has been evident at ANU.

Thena Kyprianou explains the origin of the Paul Thistlewaite Memorial Scholarship.

Paul was very conscious of the lack of women in his field. It was very much his view that gender balance was a desirable thing for the advancement of the discipline and the world as a whole. He addressed this by going out of his way to tailor his teaching to make the few women in his third year course feel included and encouraged to go on to honours year and post grad work. I recall him discussing with me social work-scenarios which might arise in a typical office which computational logic analysis/ programs could address. He would then set those as assignments for students as he thought they would interest the women students much more than pure mathematics or engineering questions. He was always looking for ways to encourage female students to remain in the discipline.

It was for this reason I thought an appropriate scholarship in my husband's memory would be one which funded a female student to undertake an honours year.

12.3.1 A Gender Balance Highlight: Data Mining and Analysis

The course COMP8410 Data Mining in Autumn 2018 achieved gender parity. It was probably the first CS course ever to do so. Following on from this, the two programs Master of Applied Data Analytics and Grad Dip of Applied Analytics reached and exceeded gender parity in 2019. I asked Kerry Taylor, who is heavily involved in teaching these courses, to speculate about reasons:

1. That COMP8410 course instance is specially for students of the Grad Dip and Masters of Applied Data Analytics, so all the things about the programs (below) also apply to that (compulsory) course. Except for maybe one thing – I am female and I teach it. I also teach it more broadly to postgrad and undergrad students – but there is no gender parity there.
2. The programs are convened by a female (me) who has done a lot of personal outreach around the public service – visiting offices and promoting the programs.

[20] `http://www.randalolson.com/2014/06/14/percentage-of-bachelors-degrees-conferred-to-women-by-major-1970-2012/`

3. The marketing budget (and these programs might well have had more than usual, although I do not really know) has often been spent on Women-focused venues e.g *Code like a Girl* and *Women in Big Data*
4. The programs target intake at postgraduates who have been in the workforce for a time – it is possible that girls who never thought of computing at school then discovered something in the workplace.
5. The programs accept a very broad intake of undergraduate degrees (so encourage direction-changing, which might be specifically attractive to women, or at least it creates an opportunity for women that discover an interest later in life as above)
6. The programs themselves are interdisciplinary and make a big feature of that– COMP, STAT and SOCR (social science) and overall quite rounded – across all three areas but with specialisation at the advanced level in one of the three. Maybe interdisciplinary is attractive to women.
7. A lot of the students are in the public service – and sometimes sponsored by the public service, and the public service has been working for a long time now to increase female participation. It might be (speculating) that the population from which we draw students actually has more women.

Ironically none of this based on data!

Analysis of ANU IT Completions data in the Higher Education uCube

- *From* `http://highereducationstatistics.education.gov.au/`
- *Accessed: 28 Feb 2021*

Total cohort: The overall completions in IT at ANU more than tripled between 2001 (125) and 2019 (406).
Domestic students: However, the total number of domestic students decreased in absolute terms (2001: 109; 2019: 84). Comparing the cumulative total of domestic student completions for the 5-year period from 2001-2005 (543) with that for 5-year period from 2015-2019 (326) to reduce the effects of single year variations, domestic completions had dropped by 40% between early and most recent periods. In 2001, domestic students accounted for 87.2% of the total completions in the field of IT at ANU but only 20.7% in 2019.
Overseas students: There was a 20-fold increase in the total number of overseas student completions between 2001 and 2019 (2001: 16; 2019: 322).
Women accounted for 20% of completions in 2001, a figure which remained the same in 2019 (20.2%). However, there was a drop in the percentage of female completers among domestic students over this period (18.3% in 2001 compared with 9.5% in 2019). Comparing the cumulative total of domestic student completions by women for the 5-year period from 2001-2005 (107) with that for 5-year period from 2015-2019 (52) to reduce the effects of single year variations, the enrolment of domestic female students had dropped by 51.4% between the early and most recent periods.

Analysis of ANU IT Enrolments data in the Higher Education uCube

- *From* `http://highereducationstatistics.education.gov.au/`
- *Accessed: 28 Feb 2021*

Total cohort: The overall enrolments in IT at ANU almost tripled between 2001 (670) and 2019 (1974).
Domestic students: However, the total number of domestic students increased by only 10% during this period (2001: 562; 2019: 616). Domestic students accounted for 83.9% of the total enrolments in the field of IT at ANU in 2001 but only 31.2% in 2019.
Overseas students: There was a greater than 12-fold increase in the total number of overseas student enriolments between 2001 and 2019 (2001: 108; 2019: 1358).
Women accounted for 22% of enrolments in 2001, a figure which remained the same in 2019 (22.2%) although the number of female enrolees among domestic students dropped somewhat over this period (115 versus 102) as did the percentage of domestic students who were women (20.5% in 2001; 16.6% in 2019). However, this effect was not as substantial as that observed for completions. (Are women enrolling at a higher rate than they are completing? And is this because they proceed to other fields or drop out all together?)

12.4 Gender Balance Among ANU IT Staff

In the 1970s and 80s, women were much better represented in the Computer Services Centre and the service section of the Computer Centre, than they were in academic computer science. For several early years, Julie Bakalor (née Rohl) was Chief Programmer. Other female programmers included Heather McKinley, Sue Groom (née Leckie), Dale Minchenton, Melanie Rooney (née Bleeze), Debra Hinton, Erin Brent, and Kathy Handel.

Erin Brent left Computer Services in the mid-80s to lead the IT team at ANU Library. Kathy Handel developed a strong reputation for her ability and willingness to support users. She went on to lead the User Services Group. See Page 193 for some stories and photos about some of these influential women.

Female CSC operators included Margaret Brinkley, Val Airey, Dorota Janiszewska, Inta Skriveris, and Cher Chung.

The Academic section of the Computer Centre was entirely male.

The Data Processing Unit, headed by Skaidrite Darius, had a largely female staff, but they were likely employed at relatively low levels. In the 80s, the Faculties Computer Unit was headed in turn by Monica Berko, Harriet Michell and Melanie Bleeze (Rooney). The staff of the CSC's Microcomputer Information Unit was substantially female. It included: Gloria Robbins, Rosemary Coggins, Virginia Woodland, and Melissa Waterford.

For many years, Mary Rose ran the Computing Section in Joint Schools and her programming and operational staff often had a female majority. Among them were Anne Sandilands, and Monica Berko. Yvonne Pittelkow was the Statistical Consultant and Margie Wood was the IT support person in Economics, RSSS.

The administrative computing service at ANU also included many female members and was headed for some years in the early 2000s by Fay Gibbons. Rick Van Haeften has written about Fay on Page 251. In more recent years, the leadership group of IT Services has had very strong female representation. Although a departing female CIO (Karen Hill) was replaced by a male (Garry Whatley), four out of five of his direct reports are female. Emelyn Pecar (IT Services Governance Officer) tells me that IT Services have been reporting gender statistics to CAUDIT (Council of Australian University Directors of Information Technology) since 2015. She says that the proportion of female direct CIO reports has not dropped below 50% in that time, and that ANU is well above the average of other universities. Across the entire IT Services organisation ANU female employment rates are and always have always been higher than the average for the Australian IT industry. The current proportion is 30%.

Another influential woman during the 1980s was Jennifer Barreda, who ran the ANU Apple Consortium arrangements through ANUTech, selling millions of dollars worth of computers and accessories to departments, staff and students at heavily discounted prices.

12.5 Current Gender Policy in CECS

There have been various initiatives over the decades to improve female participation but success has been patchy. The College of Engineering and Computer Science (CECS) seems determined to do better.

In the *Reimagine Engineering and Computer Science: Strategic Intent 2019 – 2025* document[21] we find, "Our benchmark is demographic representational parity. Initially we will prioritize our imbalance in women and Indigenous students, who are conspicuously underrepresented." On International Women's Day 2019, CECS anounced a target of 50% female participation in the college by 2030:[22]

> We know that we are coming off a low base, but we are committed to working inclusively and in an evidence-led way to improve the experience and diversity of the CECS student and staff cohort. The College is not shying away from its commitments to address ingrained diversity challenges and is approaching them head on. That's why we have created the very first Associate Dean, Diversity and Inclusion position at the University, and established a Diversity and Inclusion Committee. We are also committed to developing an integrated strategy founded on research. These are the first steps in improving the experience of women in the engineering and computing disciplines, striving for 50:50 participation by 2030.

Let's hope these goals are achieved.

[21] `https://cecs.anu.edu.au/sites/default/files/strategic_intent_2019.pdf`

[22] `https://cecs.anu.edu.au/news/cecs-reaffirms-commitment-gender-equity`

Chapter 13

ANU Computer Science: How Are We Doing?

This concluding chapter looks at ANU's achievements in the Computer Science discipline over the past half century. After fifty years of computer science research and teaching at ANU there are many impressive achievements.

On the research side, ANU academics and PhD students have made significant contributions to: machine learning; robotics; computer vision; fundamental algorithms; combinatorics, persistent programming and programming languages in general; garbage collection; parallel and high performance computing; logic programming; immersive environments; data analysis, matching and mining; information retrieval and web technologies; VLSI design; and artificial intelligence. Appendix D is inevitably incomplete but lists significant research outputs: Spinoff companies; Software which has been influential or commercially successful; and Research papers with more than a thousand citations.

On the teaching side, an important measure of success relates to the success of graduates from ANU's graduate and undergraduate programs. Appendix C lists a subset of computer science alumni who have gone on to make a distinguished name for themselves in their subsequent careers.

This chapter first highlights recognitions given to individual ANU computer scientists, then looks at external assessments of ANU computer science as a whole. There follows a section of "testimonials" from external people who have interacted with ANU over the years. The chapter, and the book, concludes with some final remarks about the path travelled, the devastating effects of the COVID pandemic in 2020, and best wishes for the next 50 years.

13.1 Individual recognitions and achievements

13.1.1 Honours and Distinctions

Richard Brent has Fellow status in ACM, IEEE, SIAM and the Australian Academy of Science. He has been awarded the Hannan Medal by the Australian Academy of Science and the Moyal Medal by Macquarie University. Steve Blackburn is also an ACM Fellow, while Tony Hosking is an ACM Distinguished Scientist.

Bob Williamson is a Fellow of the Australian Academy of Science and the Australian Mathematical Society. Brendan McKay is also a Fellow of both those organisations. Another Fellow of the Australian Academy of Science is Richard Hartley. Robin Stanton is a Fellow of the Australian Academy of Science, Technology and Engineering.

Amanda Barnard (RMIT University, 2020), Andrew Tridgell (ANU, 2018), and David Hawking (Université de Neuchâtel, Switzerland, 2003) have been awarded honorary doctorates.

13.1.2 Vice-Chancellor's Awards

In 2018, Ben Swift received a Vice-Chancellor's Citation for Outstanding Contribution to Student Learning.

13.1.3 Knuth Cheques

If you are a computer scientist you will surely know of Donald Knuth's long-standing offer of small payments to the first person to report a verified error in volumes of his *Art of Computer Programming*, in *TEX and MetaFont* or in TeX itself. Of course the honour of finding an error in the great man's work is worth far more than the value of the cheque. One presumes that very few of them are cashed. I am aware of several such uncashed cheques in the possession of ANU researchers:

Reporter	Date	Error in
Richard Brent	1975	Vol. 3 *Art of Computer Programming*
Brendan McKay	1988	*TeX and MetaFont*
Brendan McKay	1996	Vol. 2 *Art of Computer Programming*
Richard Brent	1997	Vol. 2 *Art of Computer Programming*
Henry Gardner	1997	A Knuth FORTRAN program

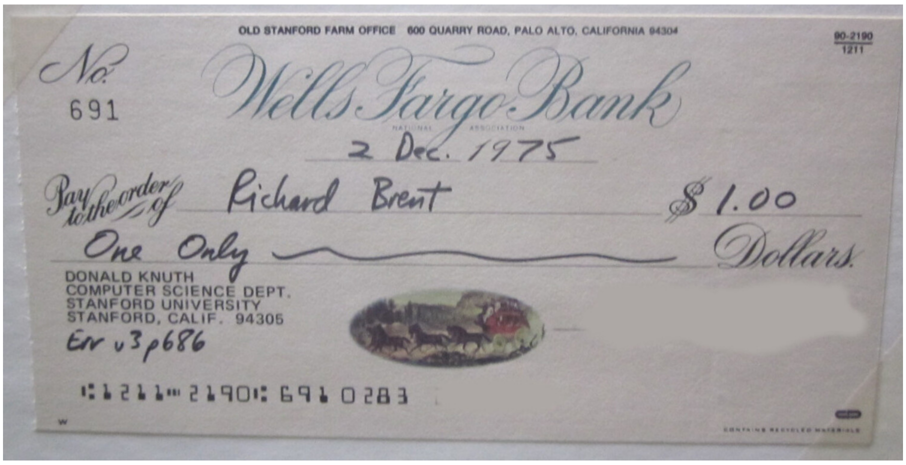

OLD STANFORD FARM OFFICE 600 QUARRY ROAD, PALO ALTO, CALIFORNIA 94304

90-2190
1211

No.
691

Wells Fargo Bank
NATIONAL ASSOCIATION

2 Dec. 1975

Pay to the order of Richard Brent $1.00

One Only Dollars

DONALD KNUTH
COMPUTER SCIENCE DEPT.
STANFORD UNIVERSITY
STANFORD, CALIF. 94305

Err v3p686

[redacted]

⑆1211⑆2190⑆ 691 0283

Mike Robson also received a Knuth cheque. He used it for many years as a bookmark in Vol 1 of *The Art of Computer Programming*.

13.1.4 Erdős Numbers

Brendan McKay has an Erdős number of 1, having published with him in 1984.

Trevor Vickers has calculated a similar distance metric based on Alan Turing and the "knows" relation rather than "has published with". On this basis, he claims a Turing number of 3 based on the fact that he once had afternoon tea with Donald Michie's wife, Donald having worked with Alan Turing at Bletchley Park. If only Donald had been home when Trevor called, Trevor could claim a Turing number of 2!

13.2 External Assessments of ANU Computer Science

13.2.1 ERA: Excellence in Research for Australia

Successive governments have been keen to promote science and innovation, while ensuring that government research funding is spent effectively.

The Howard government developed the Research Quality Framework (RQF) as an objective means of assessing the quality and impact of publicly funded research, with a view to maximising the benefits to the community. It followed the introduction of the Research Assessment Exercise (RAE) in the UK. In 2006, Julie Bishop (the relevant minister) anounced funding to enable the first cycle of RQF to take place in 2008.

This did not happen, because the Rudd Government replaced RQF with ERA (Excellence in Research for Australia.) The ERA process has run in 2010, 2012, 2015, and 2018. Outcomes[1] are reported by Field of Research (FoR) codes. ANU outcomes for the broad "Information and Computing Sciences" code (08) have rated "5 – Well above world standard" in each of the four exercises. This is an amazingly good result. As far as I can see, only the University of Melbourne comes close, with a 4 in 2010, and 5s after that.

Let's hope that the departure of so many senior staff in 2020 will not harm ANU's rating in the next exercise.

13.2.2 Ranking of World Universities

There are many organisations publishing rankings of universities around the world. Some of them provide rankings for individual disciplines. For example Shanghai Ranking's *2020 Academic Ranking of World Uni-*

[1] `https://dataportal.arc.gov.au/ERA/Web/Outcomes`

versities for Computer Science & Engineering[2] lists ANU in the same 51–75 group as UNSW, behind UTS (13), University of Adelaide (40) and University of Sydney (49).

The 2021 Times Higher Education rankings for Computer Science[3] ranks ANU at 65 out of 827, behind University of Melbourne (64), and UNSW (54).

13.2.3 Good Universities Guide Australia

The Good Universities Guide rates Australian universities from the perspective of new students. For Undergraduate students in Computing and Information Technology, here are the ratings for ANU and two comparison universities. The reader is referred to the Good Universities Guide[4] for an explanation of the dimensions.

The Good Universities Guide says:

> These performance ratings and rankings are created using a combination of resources provided by the Department of Education, and graduate and student experience surveys in Australia.

Dimension	ANU	UniMelb	UNSW
Graduate Salary	$65,000	$65,000	$65,900
Full-Time Employment	86%	79.7%	86.6%
Learner Engagement	51.9%		55.3%
Learning Resources	80.6%		82.2%
Overall Experience	69.7%		66%
Skills Development	69.4%		73.8%
Student Support	66.5%		61.8%
Teaching Quality	70.9%		75.3%

Only two dimensions were shown for the University of Melbourne when I visited the site on 22 March 2021.

13.3 Testimonials From People Who Have Worked With ANU CS

Another way to assess the success of ANU computer science is by asking people for their view of it.

Bill Gibson, one of the first Honours cohort, recently retired as CIO of the Australian Taxation Office.

"When I was doing computer science at ANU, I was just enjoying the exploration of new ways of thinking and applying the emerging technology. I have fond memories of Robin Stanton trying to (sometimes futilely) educate the student masses and of spending many hours working on a Data General mini-computer constantly re-booting it, as I had got something wrong! I wasn't made for Academia but I was able to help corporates and government agencies benefit by exploiting computing technology – I wouldn't have achieved this without the curiosity and excitement from hands-on computing science at the Department."

Nola Whitecross, Departmental Administrator during Molinari / Johnson headships

"I loved working at CS, it was my lucky day when I landed that position and I have only positive memories of my time in the Department."

[2] http://www.shanghairanking.com/Shanghairanking-Subject-Rankings/computer-science-engineering.html

[3] https://www.timeshighereducation.com/world-university-rankings/2021/subject-ranking/computer-science#!/page/1/length/25/sort_by/rank/sort_order/asc/cols/stats

[4] https://www.gooduniversitiesguide.com.au/course-provider/australian-national-university-anu/ratings-rankings/undergraduate/computing-and-information-technology

Mitsuo Ishii, Former Head, Fujitsu Laboratories, Kawasaki, Japan

"Many visits to ANU to run the collaborative research with ANU, having enjoyed sightseeing, hiking and cycling near Canberra. It's more of an impressive memory than research meeting.

I attach a photo of Symposium held at that time. The researchers in this photo played a pivotal role in parallel computer at Universities and Research Institutes in Japan. The joint research with ANU was such an influential activity."

"Regarding Fujitsu People involved in the development of supercomputers, members who had participated in the joint research between ANU and Fujitsu Labs are now working, supporting and taking the lead in the computer division in Fujitsu Ltd."

Mark Reid, ANU Research Fellow, now working for Apple, in Canberra

"Yes I worked with Stephen Gould in 2015 to put together the initial version of Craft [of Computing].

I was a research fellow at the ANU school of computer science from 2007-2016. I was part of the ML group there and worked closely with Bob Williamson on the theoretical side of machine learning.

One of the things that drew me to ANU after my PhD at UNSW was the calibre of the people there and at NICTA. Bob, Alex Smola, Marcus Hutter, and others were researchers at the cutting edge of ML theory at the time.

Throughout my time there I was very impressed by the work the school put out. People like Lexing, Stephen, Edwin Bonilla, Scott Sanner, and a slew of fantastic PhD students made it a very ambitious and fun research environment.

I moved to Apple in the Bay Area at the beginning of 2016 after being contacted by someone there. I'm still working on the same project but moved back to Canberra in mid-2017 where I continued work remotely. The project has not been released so I can't go into details about it. "

Michael Hind, Distinguished Research Staff Member, IBM Research

"ANU was a key creator and collaborator on the hugely successful Jikes RVM open source project, which received the ACM SIGPLAN Programming Languages Software Award in 2012. They contributed significant innovations, architectures, and community support that was key to the impact of the system."

Alex Zelinsky, former ANU professor, now Vice-Chancellor, University of Newcastle

"Overall comments on ANU? Great people and a fantastic environment to innovate. Lots of good people that were funded for success – pursuing the frontiers. John Moore was a legend. He knew how to spot talent early and then backed the people 100 percent. Sadly, he has passed away. A great man – who is dearly missed."

John O'Callaghan, former Chief, CSIRO Division of Information Technology

"I followed the developments in computer science at ANU ever since I graduated from ANU with a PhD in pattern recognition in 1969 – well before the Department started. Since that time, I have witnessed the emergence of ANU DCS as a significant research capability with world-class researchers and as a leader in national research programs and collaborative joint ventures.

I would like to mention the outstanding leadership of Robin Stanton who was instrumental to the initiatives that I was involved in. Within DCS, Robin upheld the ANU culture of world-class expertise and leadership in national programs.

One of the strategies to build DIT was to build relationships with other research organisations at its locations in Canberra, Sydney and Melbourne. In Canberra, Robin Stanton and I developed the proposal for the Division to move on to the ANU Campus and share a new building with DCS. This proximity was key to strengthening proposals to build collaboration and form joint ventures with support from the Federal Government. The main highlight was the establishment of the Cooperative Research Centre (CRC) for Advanced Computational Systems (ACSys) in 1993."

Dave Grove, IBM Research, ACM Fellow

"ANU obviously played an important role in the growth of the Jikes RVM open source community. It was the academic home of early community leaders like Steve Blackburn and Daniel Frampton. The ANU also hosted all of the project's publicly accessible testing infrastructure that allowed us to make continuous functional and performance regression testing possible. It's hard to remember how important (and unusual) that was for even a moderately active open source project like Jikes RVM in the days long before the emergence of cloud computing, TravisCI, GitHub, etc. That infrastructure, and the people at ANU who kept it running, were essential to the long term success of Jikes RVM as a widely-used research platform.

For me personally, ANU is the university with which I've had the longest and most fruitful research collaborations. Between the Jikes RVM and X10 projects, I was able to work with multiple ANU faculty members over more than a decade, host their students for summer internships, and help the students develop their own research careers via many phone calls at odd hours of the day. My only regret is that I managed to visit ANU just once during that time."

Eliot Moss, UMass, ACM Fellow, IEEE Fellow

"Over the years, Computer Science at ANU has developed strong ties with Computer Science at the University of Massachusetts Amherst, now a college of its own. Professor and School Director Tony Hosking was the first PhD student of Eliot Moss, recently emeritus at UMass, and Professor Steve Blackburn was Moss's first post-doctoral researcher. Both have maintained strong ties with UMass and with colleagues they met there, and students on both sides have benefited from continuing research collaborations, that have been enormously successful in terms of not only publications and grants, but also graduated students now out, building our reputations around the globe. I myself prize these collaborations and the inspiration and energy that our visits back and forth bring. ANU Computer Science continues on a tremendous upward trajectory, which will only continue! Congratulations!"

Andy MacFarlane, Reader, Centre for HCI Design, City University, London

"In 1998, I visited ANU for three months as part of the PhD in Information Science I was completing at City University. My PhD investigated the best way to distribute inverted file indexes in parallel machines to speed up tasks such as searching, indexing, maintaining, passage retrieval and query optimisation. City did not have the infrastructure to support this, but fortunately David Hawking came to the rescue and I was able to access the DEC AlphaFarm and the Fujitsu AP1000 and AP3000.

I spend three very happy months in chilly Canberra, building friendships with many members of the department including Dave Walsh, Dave Sitsky, Steve Blackburn and Nick Craswell. I stayed in Bruce Hall, and I was amazed by the variety of wildlife in Australia, but the cockatoos outside my room making a racket when I was trying to sleep was sometime a little wearing!

My time at ANU was extremely fruitful and I was able to develop my ideas further and the software I was writing for my PhD. I was able to conduct all the experiments I needed to finish my thesis and make a contribution. I graduated a couple of years later, and have been back at City as an academic for nearly 20 years. My visit to ANU proved to be a strong foundation for my career, and I look back fondly to my time there."

David Sitsky, Chief Scientist, Nuix Pty Ltd

"I had a brief stint working at Cisco Systems in embedded operating system work, before joining a software startup called Nuix, which I am still working at 20 years later! Many ANU people have worked at Nuix, including Peter Bailey, Nigel Snoad, Raj Nagappan, plus two others that at the time of writing still work there: Luke Quinane and Daniel Noll. All of these individuals clearly had quality education from ANU which has enriched Nuix. I have no doubt the quality education and work experience I obtained at ANU helped me enormously in my professional career, which I am always grateful for."

13.4 Last words

2020 was a disastrous year for many universities around the country, due to the COVID-19 pandemic and its consequences. ANU has had to curtail major developments and find hundreds of millions of dollars in savings. CECS has been particularly badly affected. Unfortunately, this means that many valued staff members have left or will leave the Computer Science areas of the college. Despite this, many strong performers remain, and we can be optimistic for the future.

Over the decades, computer scientists from ANU have made major contributions to the discipline of computer science and the industries it feeds, through research, teaching, and outreach. As we have seen, they have been highly visible on the world stage and participated in numerous valuable collaborations, both domestic and international. Some ANU Computer Science research has been commercialised, but so far we have seen no unicorn let alone dekacorns or hectocorns.[5] Could Quantum Brilliance,[6] builders of room temperature quantum computers, and founded by ANU physicists Andrew Horsley and Marcus Doherty, be the first? Which of our academics or PhD students will in the future emulate Larry Page or Sergei Brin?

ANU is now equipped with a very high level of computational and communications infrastructure to support computer science and computing generally. Network connectivity from just about every room in the ANU is far superior to most people's home connections via the NBN. Thanks to the location of the National Computational Infrastructure on the ANU campus, all ANU researchers have access to fantastic computational power. I estimate that Gadi, with its close-to 10^{16} Flops, is at least thirteen orders of magnitude more powerful than the IBM 610 which, from the beginning of 1960 to the end of 1961, was ANU's only computer.

Over the decades, computer science activities have been carried out in at least a dozen different administrative units. In my opinion there was often an unfortunate lack of communication and collaboration between groups in existence at the same time. Currently, it seems that most computer science researchers are located in the new School of Computing, or at least within the College of Engineering and Computer Science. While the scale of these organisations makes it difficult to create the marvellous esprit de corps of DCS in its first decades, hopefully large scale co-location will lead to more and better mentoring, exciting synergies, and to greater and better outcomes.

[5] Unicorn – Company valued at more than one billion dollars. Work the others out for yourself!

[6] `https://quantumbrilliance.com/`

Appendix A

ANU Computing Timelines

Timelines showing highlights from ANU Computing history. ANU computers are shown in this colour. Items relating to DCS or its successors are shown in this colour.

1960 – 1974

1960 IBM 610 arrives at MSO. Claire Wehner uses it for Time Service, Don Faulkner and others for astronomical calculations. Claire had worked on EDSAC project.

1961 Pat Wilson becomes Manager of DPU, at the time using Unit Record Equipment rather than computers.

1962 Ken Le Couteur of Theoretical Physics arrranges for IBM 1620 Model 1 to be installed in Oliphant Bldg, RSPhysS. Brian Robson is in charge. Margaret Campbell and Peter Tindale are Operators, Elizabeth Reid is Operator/Programmer. Ken had worked with Colossus, Brian with CSIRAC. Skaidrite Darius transfers to DPU.

1964 1620 Model II replaces Model I which is moved to MSO. Lorrel Sherar becomes an Operator of the 1620.

1965 Barbara Davidson employed as Programmer on the 1620. Lorrel Sherar transfers to DPU as Programmer.

1966 Mike Osborne appointed Director, ANUCC. IBM 360/50 installed in Cockroft Bldg and used for administrative as well as academic computing. Peter Creasy employed as an ANUCC Programmer.

1968 Julie Rohl (Bakalor) joins ANUCC as programmer, later Chief Programmer. John O'Callaghan completes first ANU CS PhD.

1970 Computer Science sub-department created within Department of Statistics, Faculty of Economics. Ray Jarvis appointed as Senior Lecturer, Peter Creasy and Tony Cantoni as Lecturers.

1971 Computer Science I commences 02 March as a second-year subject. Vicki Peterson appointed Tutor. David Hawking is first-year undergraduate. DCS acquires Data General Supernova.

1972 Univac 1108 installed under Menzies Library. MSG acquires Inforex 5000. ANU introduces semesters. CS B01/B02 commence. B01 teaches ANSI FORTRAN and numerical methods. Computer Science II taught to third years. Robin Stanton, Brian Molinari join DCS as Senior Lecturers. Tony Cantoni leaves. Richard Brent joins ANUCC.

1973 CS C01/C02/C03 taught. Peter Kelo does joint 4th year honours with Pure Maths.

1974 First 4th year honours cohort: Graeme Williams, David Hawking, Bill Gibson, Liz Barker, Les Landau (part-time). John Hurst arrives as a Lecturer.

1975 – 1994

1975 Univac 1100/42 in Leonard Huxley Bldg replaces Univac 1108. Vicki Peterson leaves. Rob Ewin and David Hawking join as Tutors. John Hurst takes over B01 and teaches Structured Programming and Algol W. Malcolm Newey joins ANUCC. Pat Wilson becomes full-time Head, MSG.

1976 DCS becomes a Department with joint membership of the Faculties of Economics and Science. Bev Johnstone becomes Administrator.

1977 Mike Robson takes up Lecturing Fellowship in DCS. ANUCC split into CSC and VCCRG (academic). Bob Watts appointed interim Director, CSC. ACSC-0 (Programming Language Systems) held in Canberra.

1978 Bob Landford appointed Director, CSC. MSG and DPU incorporated into CSC.

1979 DEC KA-10 installed in Copland G4, to serve undergraduate teaching requirements. David Hawking is site manager. Geoff Huston replaces him as Tutor. Richard Brent appointed foundation Professor. First year CS teaching (A01/A02) commences. SGS is renamed “The Faculties”.

1980 Univac 1100/82 replaces Univac 1100/42. Facom M160F acquired to run admin systems – to be rewritten in Adabas/Natural.

1981 Malcolm Newey joins DCS. Robin Erskine becomes Assistant Director, CSC.

1982 Brian Anderson founds Department of Systems Engineering.

1983 Faculties DEC KA-10 replaced by DEC KL-10. Richard Brent seconded to Centre for Mathematical Analysis. Robin Stanton becomes Head, DCS.

1984 Faculties DEC KL-10 replaced by VAXCluster. ANU is a founding member of the Apple Universities Consortium (AUC).

1985 Richard Brent becomes Head, Computer Sciences Lab in RSPhysSE.

1986 DCS moves to top floor, Crawford Bldg. ARP founded.

1987 ANU acquires Fujitsu VP-50 supercomputer, to be upgraded to VP-100 in time for Australian Bicentenary. Univac 1100/82 decommissioned. Bob Landford leaves ANU. Don Hardman interim Head, MISD, pending arrival of Ken Vine.

1989 The Internet arrives in Australia through AARNet. Undergraduate engineering program founded. Darrell Williamson is its first Head.

1990 First intake of BE students. The Pub and Colour Lab Sun workstation labs set up. First ANU-Fujitsu CAP Project Workshop. Connection Machine CM2 installed in Huxley Bldg..

1991 Fujitsu AP1000 installed in Crawford Bldg. George McLaughlin becomes Head, MISD.

1992 First website in Australia: `life.anu.edu.au` Fujitsu VP-2200 acquired.

1993 FEIT formed in time for first BE graduation. Computer Science becomes part of FEIT. Robin Stanton becomes inaugural Dean. Brian Molinari becomes Head, DCS. ACSys and RDN CRCs commence in new building now known as Brian Anderson Bldg. CNIP launched.

1994 RSISE founded with Brian Anderson as Director. Launch of Australian Co-operative Supercomputer Facility. Fay Gibbons becomes Head, MISD.

1995 – 2021

1995 DCS moves into north wing of CS&IT Bldg, CSIRO CMIS into south wing.

1996 Fujitsu VPP300 acquired.

1998 Launch of AJNL and APAN networks. Chris Johnson takes over as Head, DCS. John Lloyd becomes Head, CSL. ANU acquires PeopleSoft.

1999 S@NITY search engine (later P@NOPTIC, later Funnelback) launched at ANU. Launch of Australian Partnership for Advanced Computation with John O'Callaghan as Executive Director.

2000 ACSys CRC winds up. ANU proves to have successfully mitigated Y2K risks. Attempted sale of last vector processing supercomputer (Fujitsu VPP300) on eBay. Seeing Machines founded.

2001 APAC acquires a 1 teraFlops machine, a Compaq/HP Alphaserver SC.

2002 Launch of GrangeNet. Launch of NICTA.

2003 Tom Gedeon becomes Head, DCS, but soon becomes Acting Dean, FEIT and then Deputy Dean.

2004 Chris Johnson returns as Head, DCS.

2005 Creation of CECS. Death of Fay Gibbons. RIck Van Haeften becomes Director, Corporate Information Services. APAC acquires a 4 teraFlops SGI Altix 3700.

2006 Burial of ANU's last mainframe – Fujitsu GS8400 in front of Huxley Bldg.

2007 Launch of National Computational Infrastructure (NCI). Bob Williamson takes over from John Lloyd as Head, CSL.

2008 Henry Gardner becomes Head, DCS.

2009 NCI acquires Sun Constellation (Vayu), 140 teraFlops.

2010 ARC Special Initiative on Implantables (Bionic Eye) launched.

2011 Sylvie Thiébaux becomes first female professor in DCS.

2012 NCI acquires a 2 petaFlops Fujitsu supercomputer (Raijin).

2013 Alistair Rendell becomes Head, DCS.

2014 Australian Centre for Robotic Vision established.

2015 NICTA winds down, merges with CSIRO ICT Centre to become Data61.

2017 Establishment of 3A Institute under Genevieve Bell.

2018 Suthagar Seevaratnam appointed ANU's first Chief Information Secutity Officer after major hack of ANU systems.

2019 Tony Hosking becomes Head, DCS. New Hanna Neumann Bldg opened. ANU-ASD Co-Lab launched on its top floor. NCI Fujitsu/Lenovo supercomputer (GADI) comes into service with over 9 petaflops peak performance.

2020 Year of COVID-19 pandemic with major budgetary consequences. Significant downsizing in CECS.

2021 School of Computing (DCS) and School of Cybernetics formed within CECS.

Appendix B

Deans, Heads, and Directors

Unfortunately, there is a dearth of official records. It is likely that there are some errors in these tables.

Head, Department of Computer Science

Including time as a sub-department of Statistics.

1971 – 1979	Ray Jarvis
1979 – 1983	Richard Brent
1983 – 1993	Robin Stanton
1993 – 1997	Brian Molinari
1998 – 2002	Chris Johnson
2003	Tom Gedeon
2004 – 2008	Chris Johnson
2008 – 2009	Henry Gardner

Director, Research School of Computer Science

Called “School of Computer Science” between 2009 and 18 Feb 2011.

2008 – 2013	Henry Gardner
2013 – 2018	Alistair Rendell
2019 –	Tony Hosking

Head, Department of Systems Engineering

1982 – 1991	Brian Anderson AC
1992 – 2004	John Moore (Alex Zelinsky acted as Head while John was on leave.)

Director, Research School of Information Science and Engineering

1994 – 2002	Brian Anderson AC
2003 – 2008	John Richards

Head, Computer Sciences Lab

1985 – 1998	Richard Brent
1998 – 2007	John Lloyd
2008 – 2009	Bob Williamson

Dean, Faculty of Engineering and Information Technology

1993 – 1997	Robin Stanton
1997 – 2000	Darrell Williamson
2001 – 2003	John Baird
2004 – 2005	John Richards

Dean, College of Engineering and Computer Science

2006 – 2007	John Richards
2007 – 2008	Mick Cardew-Hall
2008 – 2011	Chris Baker
2012 – 2014	John Hosking
2014 –	Elanor Huntington

Appendix C

Distinguished Alumni

Graduates of computing related degrees from any part of ANU, who went on to a distinguished career. Apologies to those deserving people inadvertently omitted.

Name	Grad. yr	Position, Achievement, etc.
Apostoloff, Nick	2015	Machine Learning Manager, Apple Computer, Academy Award 2017, Scientific and Technical
Bailey, Peter	1997	Principal Applied Scientist, Microsoft
Blackburn, Steve	1998	Professor of CS at ANU
Cai, Jie	2011	Head, Search and Discovery, Tencent Medical
Corbould, Mark	1979	Asst. Director-General IT, National Library of Aust.
Craswell, Nick	2000	Principal Group Science Manager, Microsoft Bing
Elford, Peter	1983	Co-founder of AARnet with Geoff Huston
de Ferrari, Lisa	1990	Barrister, Senior Counsel
Fidge, Colin	1989	Professor of CS, QUT
Fifield, Davin	1992	VP, Product Development, Oracle Corporation
Gibson, Bill	1974	CIO at Australian Taxation Office
Hansen, David	1995	CEO, Australian e-Health Research Centre
He, Zhen	2003	Assoc. Professor of CS, Latrobe University
Huston, Geoff OAM	1978	Chief Scientist, APNIC, Member, Internet Hall of Fame
Keronen, Seppo	1993	Head of Advanced Development, Nokia
Le, Quoc	2006	Research Scientist, Google Brain, h-index 87
Liang, Weifa	1998	Professor in DCS
Lokan, Chris	1985	Assoc. Professor at UNSW, Canberra
Mackerras, Paul	1988	rsync (with Andrew Tridgell), Linux for PowerPC
Manning, Chris	1989	Director, Stanford Artificial Intelligence Lab.
McConnell, Paul	1980	Founder, Vocus and TransAfrica Communications
McRobbie, Michael AO	1979	President, University of Indiana
Michael, Gavin	1996	Head of Technology, CitiBank, NY, CTIO, Accenture
Molinari, Rory	1992	Principal Software Engineer, PayPal
Nadasi, Daniel	2007	Engineering Director, Google Photos, Sydney
O'Callaghan, John AM	1969	Exec. Director, Aust. Partnership for Advanced Computation
O'Kane, Mary	1981	Chief Scientist, NSW, Vice-Chancellor, University of Adelaide
Over, Andrew	2008	Principal engineer, Google Australia
Poole, David	1982	Professor of CS, University of British Columbia
Popple, James	1993	Australia's first FoI Commissioner
Roxas (née Onate), Rachel	1994	Vice President for R&D of National University of the Philippines
Sharma, Dharmendra AM	1992	Dean and Distinguished Professor, University of Canberra
Sitsky, David	1993	Chief Scientist, NUIX
Sloane, Anthony	1984	Assoc. Professor of CS, Macquarie University
Sutherland, John	1992	CIO, Cochlear, now CIO, Ramsay Health
Syme, Don	1992	Principal Researcher, MSR Cambridge, Designer of F# language
Taylor, Kerry	1995	Professor in DCS
Taylor, Sam	2001	Senior Director of Engineering at Arm, Cambridge
Tridgell, Andrew OAM	1999	DSc (Honoris Causa, 2018), Creator of Samba, rsync (w. Mackerras)
Upstill, Trystan	2004	VP of Engineering, Google
Wagner, Michael	1979	Three honorary CS professorships
Wanless, Ian	1998	Professor of Mathematics, Monash University
Whigham, Peter	1983	Director, Spatial Information Research Cent., University of Otago
Williams, Graham	1991	Director of Data Science, Microsoft Singapore
Yang, Jian	1995	Director of Research, Dept. Computing, Macquarie University

Early ANU Computer Science PhD Completions

Year	Name	Department	Supervisor	Topic
1968	John O'Callaghan	Eng. Physics	Steve Kaneff	*Pattern recognition using some principles of the organism-environment interaction.*
1969	George Watson	ANUCC	Mike Osborne	*Chebyshev approximation with applications to the numerical solution of differential equations.*
1970	Iain MacLeod	Eng. Physics	Steve Kaneff	*A study in automatic photo-interpretation.*
1971	David Ryan	ANUCC	Mike Osborne	*Transformation methods in nonlinear programming.*
1973	Les Jennings	ANUCC	Mike Osborne	*Orthogonal transformations and improperly posed problems.*
1973	Richard Weiss	ANUCC	Bob Anderssen	*Numerical procedures for Volterra integral equations.*
1973	Frank de Hoog	ANUCC	Bob Anderssen	*Product integration techniques for the numerical solution of integral equations.*
1975	Bernard Omodei	ANUCC	Bob Anderssen	*Stability of the Raleigh-Ritz-Galerkin procedure for elliptic boundary value problems.*
1977	Jim Abbott	ANUCC	Richard Brent	*Numerical continuation methods for nonlinear equations and bifurcation problems.*
1978	Krisom Jittorntrum	ANUCC	Mike Osborne	*Sequential algorithms in nonlinear programming.*
1978	Michael Wagner	ANUCC	Bruce Millar	*The application of a learning technique for the identification of speaker characteristics in continuous speech.*
1979	Graeme Chandler	ANUCC	Bob Anderssen	*Superconvergence of numerical solutions to second kind integral equations.*
1979	Paul Pritchard	DCS	Malcolm Newey	*An axiomatic semantics for expression languages.*
1979	Michael McRobbie	Philosophy	Bob Meyer	*A proof theoretical investigation of relevant and modal logics.*
1980	Andreas Griewank	DCS	Richard Brent	*Analysis and modification of Newton's method at singularities.*
1980	John Slaney	Philosophy	Bob Meyer	*Computers and relevant logic : a project in computing matrix model structures for propositional logics.*
1981	Mike Robson	DCS	Published work	*Worst case fragmentation of dynamic storage allocation algorithms.*
1982	Chris Johnson	DCS	Robin Stanton	*Data structure representation and transformation.*
1982	Stephen Edwards	DCS	John Hurst	*Definition of programming languages using transformational semantics.*
1982	Mary O'Kane	Eng. Physics	Iain MacLeod	*Acoustic-phonetic processing for continuous speech recognition.*
1982	David Poole	DCS	Robin Stanton	*The theory of CES : a complete expert system.*

Notes:

1. A couple of non-CS theses prior to John O'Callaghan's made use of computers.
2. In most cases 'year' is the year-of-publication as recorded in ANU's Open Theses collection: `https://openresearch-repository.anu.edu.au/handle/1885/3`

Crawford Prize for Outstanding PhD Thesis

Year	Name	Department	Supervisor	Topic
1996	Wee Lee	RSISE	Bob Williamson, Peter Bartlett	*Agnostic learning and single hidden layer neural networks.*
1998	Ian Wanless	DCS	Brendan McKay	*Permanents, matchings and Latin rectangles.*

University Medals: Computer Science or Information Technology (Up to 2020)

Year	Name
1983	Peter Wigham
1985	Andrew Kirk
1987	David Wanless
1991	Richard Walker
1992	Donald Syme
1993	David Sitsky
1994	Seng Wai Loke
1994	Michael Kirk
1997	Ian McIntosh
1997	Kent Skinner
1999	Matthew Gray
1999	Stephen Nees
2000	Nigel Tao
2000	Vijay Boyapati
2000	Michael Compton
2001	Shaun Tipson
2001	Kee Siong Ng
2002	Filip Radlinski
2004	Simon Thornton
2006	Linda Postniece
2009	Alexander Davies
2009	Christopher Pelling
2009	James Thompson
2010	Aaron Defazio
2010	Mayank Daswani
2010	Ramana Kumar
2010	James Barker
2012	Daniel McNamara
2013	James Bornholt
2014	Jesse Wu
2014	Timothy Sergeant
2015	Daniel Filan
2015	Kerry Olesen
2016	John Aslanides
2016	Theodore Olsauskas-Warren
2017	Andrew Wrigley
2017	Quyu Kong
2017	Samuel Toyer
2018	Dmitry Brizhinev
2019	Alexander Soen
2019	Matthew Brown
2019	Michael Cohen
2019	William Wei You Shen
2020	Brenda Wang
2020	Nuoya Wang
2020	Zixian Cai

Notes:

- DCS had no representative on the ANU Awards Committee for many years.
- The list above was supplied by the School of Computing in May 2021. ANU's list of University Medalists at `https://www.anu.edu.au/students/program-administration/prizes/university-medal` does not list subject area. The "Previous Winners" document `https://www.anu.edu.au/files/prize/Previous%20winners%20of%20the%20University%20Medal%20and%20the%20University%20Prize_1.pdf` does, but at the time of writing, only covers winners up to 2015.

Appendix D

Influential Outputs

Spinoffs and Startups

Company	Founded	Research Area	Business area
ReachIn Technologies AB	≈ 1995	Virtual environments	Surgical training simulator
Seeing Machines	2000	Computer vision	Operator monitoring for transport safety
Funnelback	2005	Information retrieval	Enterprise and website search
Saluda Medical	2010	Implantables	Closed-loop spinal stimulation for pain relief
Interferex	2011	Networks	Parallel path internet for video
Bionic Vision Technology	2017	Implantables	Bionic eye
Quantum Brilliance	2020	Quantum physics	Room temperature quantum computers
PostAc® – ANU / Data61	Soon	Machine learning	Job search for higher degree graduates

Notes:

- Initial funding for ReachIn Technologies came from a Swedish funding body.
- Mingchao Yu (2016 CECS PhD *Throughput and Delay Optimization of Linear Network Coding in Wireless Broadcast*) made significant contributions to the technology of Interferex. ANU Connect Ventures was an investor.
- As at the time of writing, PostAc® has paying customers but no company has yet been formed.
- Henry Gardner informs me that Wedges were sold to the Powerhouse Museum, ADFA and the CSIRO Discovery Centre.

Software

Author	Year	Name	Comments
McKay, Brendan	1981	nauty	C, was FORTRAN and 7044 assembler
Tridgell, Andrew	1992	Samba	C
Hawking, David	1992	PADRE	C
Slaney, John	1994	Magic	C
Tridgell, Andrew & Mackerras, Paul	1996	rsync	C
McKay, Brendan	1996	plantri	C
Blackburn, Steve	2002	MMTk / JIKES RVM	
Blackburn, Steve	2006	DaCapo benchmarks	
Christen, Peter	2008	Febrl	Python, 33,500 downloads
Sorenson, Andrew & Swift, Ben	2011	Extempore	livecoding environment
Bui, Minh	2015	iq-tree	

Publications

Author	Year	Citations	Title
Hartley, R	2000	27431	Multiple view geometry in computer vision
Smola, AJ	2002	19491	Learning with kernels: Support vector machines, ...
Smola, AJ	2004	10519	A tutorial on support vector regression
Lloyd, JW	1987	8527	Foundations of logic programming
Smola, AJ	2001	4989	Estimating the support of a high-dimensional distribution
Minh, BQ	2015	4961	IQ-Tree: A fast and effective stochastic algorithm for ...
Williamson, RC	2001	4876	Estimating the support of a high-dimensional dist.
Brent, RP	2013	4546	Algorithms for minimization without derivatives
Hartley, R	1999	4496	Bundle adjustment – a modern synthesis
Bartlett, P	1998	3330	Boosting the margin: A new explanation for the effectiveness ...
Bartlett, P	2000	3239	New support vector algorithms
Williamson, RC	2000	3024	New support vector algorithms
Hartley, R	1997	2990	In defence of the eight-point algorithm
Christen, P	2014	2557	Context aware computing for the internet of things: A survey
Sweetser, PK	2005	2520	GameFlow: a model for evaluating player enjoyment in games
Anderson, BDO AC	1971	2137	Wireless sensor network localization techniques
Norrish, M	2010	2031	seL4: Formal verification of an operating-system kernel
Minh, BQ	2017	2002	ModelFinder: fast model selection for accurate phylogenetic ...
Bartlett, P	2002	1834	Rademacher and Gaussian complexities: Risk bounds and ...
Mendelson, S	2002	1834	Rademacher and Gaussian complexities: Risk bounds and ...
Bartlett, P	1999	1702	Neural network learning: Theoretical foundations
Blackburn, S	2006	1572	The DaCapo benchmarks: Java benchmarking ...
Hosking, T	2006	1572	The DaCapo benchmarks: Java benchmarking ...
Hartley, R	1997	1557	Triangulation
Smola, AJ	2004	1555	kernlab – an S4 package for kernel methods in R
Bartlett, P	1998	1512	The sample complexity of pattern classification with neural ...
Gould, S	2018	1487	Bottum-up and top-down attention for image captioning ...
Bartlett, P	2000	1479	Regularization networks and support vector machines
Williamson, RC	2000	1479	A generalised representer theorem
Smola, AJ	1999	1448	Support vector method for novelty detection
McKay, BD	1981	1442	Practical graph isomorphism
Brent, RP	1982	1435	A regular layout for parallel adders
Williamson, RC	2000	1388	Support vector method for novelty detection
Li, HD	2015	1165	The Visual Object Tracking VOT2015 challenge results
Minh, BQ	2013	1228	Ultrafast approximation for phylogenetic bootstrap
Jarvis, RA	1983	1214	A perspective on range-finding tech.s for computer vision
Minh, BQ	2018	1188	UFBoot2: Improving the ultrafast bootstrap approximation
Williamson, RC	2004	1090	Online learning with kernels
Brent, RP	1974	1039	The parallel evaluation of general arithmetic expressions
Christen, P	2012	1026	Data matching concepts and techniques for record linkage, ...
Hutter, M	2005	1012	Universal artificial intelligence: Sequential decisions based on ...

Notes:

- Only papers with at least 1000 Google Scholar citations, as at 27 Sep 2020, are included.
- For multi-author papers, only the ANU author is listed.
- All papers for people with a current ANU connection were considered. For people no longer at ANU, only papers published while at ANU were considered.
- Four other papers by Brian Anderson, with more than 1000 citations, appear to be in the area of control theory rather than computer science.
- The citation count for the first paper on the list is of the right order of magnitude. In March 2021, Google Scholar was erroneously attributing the book to someone other than the book's authors. The correct citation count in March 2021 was above 29,000.

Appendix E

Computer Science Degrees

In the 1970s, computer science subjects could be taken as part of the traditional line-up of ANU degrees – principally BSc, BA, or BEc. Later there was a trend toward named degrees, prompted by a feeling that a degree in computer science might lead to more employment opportunities than a BSc, even if the BSc included as many computing units. This led to the introduction of named degrees such as the Bachelor of Information Technology and Bachelor of Software Engineering. By 2020 the number of degrees had expanded considerably, even though the following have come and gone.

- Bachelor of IT (Information Systems)
- Bachelor of IT (Software Engineering)
- Bachelor of IT (Commerce)
- Bachelor of IT (Economics)
- Bachelor of IT (Science)

2021 and beyond[1]

- Diploma of Computing
- Bachelor of IT (plus Honours)
- Bachelor of Advanced Computing (with Honours)
- Bachelor of Advanced Computing R&D (with Honours)
- Bachelor of Applied Data Analytics (plus Honours)
- Bachelor of Engineering (with Honours)
- Bachelor of Engineering R&D (with Honours)
- Graduate Certificate of Applied Data Analytics
- Graduate Certificate of Data Engineering
- Graduate Certificate of Machine Learning and Computer Vision
- Graduate Diploma of Applied Cybernetics
- Graduate Diploma of Applied Data Analytics
- Graduate Diploma of Computing
- Master of Applied Cybernetics
- Master of Applied Data Analytics
- Master of Computing
- Master of Advanced Computing
- Master of Engineering in Electrical Engineering
- Master of Engineering in Mechatronics
- Master of Engineering in Renewable Energy
- Master of Machine Learning and Computer Vision
- Master of Philosophy
- Doctor of Philosophy
- The Bachelor of Software Engineering (with Honours) is planned to have its last starting cohort in 2021, with teach out to follow.

[1] List provided by Tony Hosking on 07 Jan 2021

Appendix F

You might also like to read

A Vision Splendid: The History of Australian Computing Graeme Philipson, Australian Computer Society http://www.acs.org.au/content/dam/acs/acs-publications/ACS-ebook-2017_A-Vision-Splendid_The-History-of-Australian-Computing.pdf

The CIRRUS Computer – Computing & Computing Science Collection The University of Adelaide. https://www.adelaide.edu.au/uni-collections/collections/CIRRUScomputer/ – links to many fascinating documents.

Firty Years of Computing at The University of Melbourne Alistair Moffat. https://people.eng.unimelb.edu.au/ammoffat/fifty-years/mof06history.pdf

History of AARNet https://www.aarnet.edu.au/about-us/history

AARnet: 20 Years of the Internet in Australia Glenda Korporaal http://mirror.aarnet.edu.au/pub/aarnet/AARNet_20YearBook_Full.pdf

A Brief History of the Internet in Australia by Roger Clarke http://www.rogerclarke.com/II/OzIHist.html

Morning Dew on the Web in Australia: 1992-95 http://www.rogerclarke.com/II/OzWH.html

Streams in the History of Computer Education in Australia: An Overview of School and University Computer Education Arthur Tatnall and Bull Davey https://link.springer.com/content/pdf/10.1007%2F1-4020-8136-7_9.pdf

The Making of The Australian National University, 1946 – 1956 S.G. Foster & Margaret M. Varghese, Allen & Unwin, 1996. https://press.anu.edu.au/publications/making-australian-national-university

Fire in the Belly: the first 50 years of the pioneer school at the ANU Trevor Ophel and John Jenkin, Research School of Physical Sciences and Engineering, 1996. ISBN 0 85800 048 2. https://physics.anu.edu.au/about/history/fire_in_the_belly/

CSIRO Computing History Chapter 3 Robert C. Bell https://csiropedia.csiro.au/csiro-computing-history-chapter-3/

History of ANU ANU Archives http://archives.anu.edu.au/collections/university-archives/history-anu

Computer Science Department Annual Reports 1976 – 1981 Many are missing. https://openresearch-repository.anu.edu.au/handle/1885/149234. Note that details of DCS activities for 1971 – 1975 are included in the Statistics Department annual reports 1962 – 1986: https://openresearch-repository.anu.edu.au/handle/1885/149237

Computer Centre Annual Reports 1968 – 1991 Some are missing. https://openresearch-repository.anu.edu.au/handle/1885/149187

Australian Partnership for Advanced Computing https://web.archive.org/web/20090120211513/http://www.apac.edu.au/

Appendix G

IDLE THOUGHTS

A poem for those old enough to remember the Univac

Twas BRKPT and the I/O queue
Was SYMing Fastrand like the wind,
All idle was the CPU :
Production had just FINned.

‘Beware the UNIVAC, my son,
It’s Fastrand and its high-speed drums,
And FIELDATA, and listen for
The CTMC’s hum!’

He quickly found an unused ’scope,
And then keyed in his SITE-ID;
He typed ‘@RUN’ and then sat back
To wield his CRT.

‘NO ACTIVE RUN’ it answered back,
And ‘WAITING ON FACILITY’,
‘BAD STATUS WORD FROM CSF’
And then just ‘SYMB 03’.

‘I’ll fix you now’, he shouted out,
‘You’ve finally got me ired.
I’ll use the system’s terminal
1200-BAUD, hard-wired!!!’

‘I’ll write a loop in SSG
To make your ferrite holler,
Ten thousand runs, and in each one
Ten ERs to FORK\$.’

‘Each job will write ten 9-TRK tapes
And then rewind and read them;
Each tape, of course, is punched to cards
For backup if I need them!’

As fast as light his fingers write:
@SETC and @TEST, @JUMP,
@XQT and then, for spite,
A full post mortem dump.

He wiped his hands upon his shirt
And then he FINned his run
And scurried to the console
To sit and watch the fun.

‘MEMORY FAULT’ the system cried,
And ‘PARITY-07 ADG’,
‘PANIC DUMP IMPOSSIBLE’,
And ‘ERROR 53.’

‘Oh, Frabjous day, calloo, callay,
I’ve made the system stall!’
He tore it from the pagewriter
And taped it in the hall.

Twas BRKPT and the I/O queue
Was SYMing Fastrand like the wind,
All idle was the CPU :
Production had just FINned.

From: Australian National University Computer Centre
Programming Systems Bulletin No. 93
November/December 1973
(Author uncredited)

Appendix H

Lindsay Botten's History of NCI and its Predecessors

This history, written by former NCI director Lindsay Botten, was published on an old version of the NCI website but disappeared during a refresh. It has been restored from the Wayback machine (`archive.org`*) and is quoted here with Lindsay's permission. The original document was dated July 2013. I've also included Lindsay's May 2017 update. (*`https://web.archive.org/web/20170609055534/http://nci.org.au/about-nci/our-role/history/the-nci-years-2007-onwards/`*I've made a few small edits.*

NCI can trace its lineage back through three stages of the evolution of high-end computing services in Australia. These are:

- The Early Years: The initiation of high-performance computing services through the Australian National University Supercomputing Facility (ANUSF) from 1987;
- The APAC Years: Its extension to a national role under the Australian Partnership for Advanced Computing (APAC), hosted by ANU from 2000– 07, during which national HPC service was provided from ANUSF, a national partnership was formed, services were broadened to include a range of outreach activities to build uptake, and a national grid program, and nascent data services were established.
- The NCI Years: The current stage of advanced computing services that have been developed from 2007 onwards under the badge of NCI, again hosted by ANU, which are characterised by the broadening and integration of services, the evolution of a strong sustaining partnership, and the transition from high-terascale to petascale computational and data infrastructure to support Australian science.

Before describing this evolution under the relevant headings below, it is interesting to characterise this growth in capacity and capability in a single figure (below) which displays the evolution in HPC capacity over 25 years in Australia.

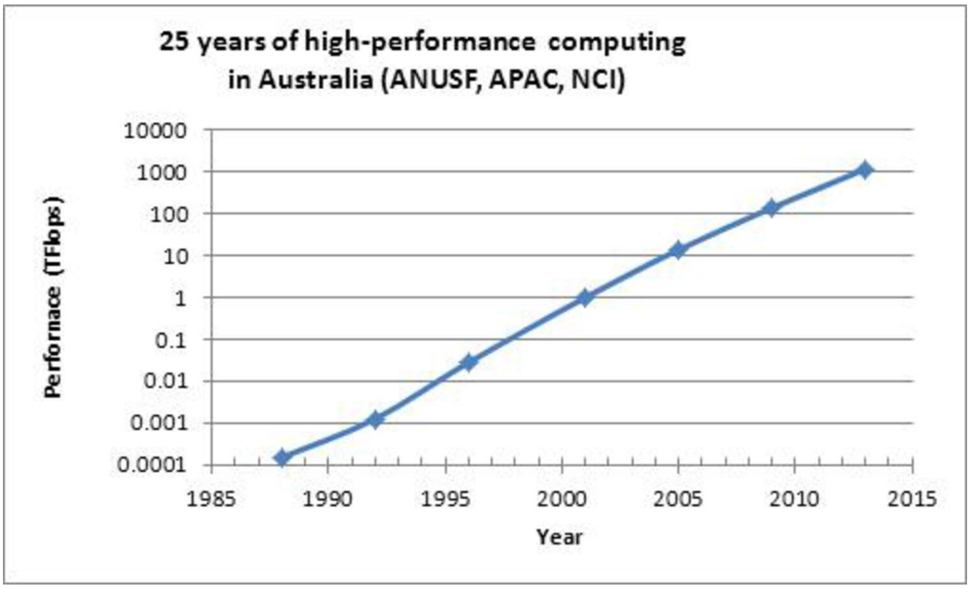

Summary of the Growth of Capability of HPC Systems at NCI (and antecedents) from 1987

Year	System (name)	Processor	Memory	Storage	Peak Perf.	Sust. Perf. (SPEC)	Initial TOP500 Rank
1987	Fujitsu VP50, 100	Vector	64 MB	12.6 GB	0.286 GF		
1992	Fujitsu VP2200	Vector	512 MB	27 GB	1.25 GF		
1995	SGI Power Challenge XL	MIPS R10000 ×20	2 GB	77 GB	6.4 GF		
1996	Fujitsu VPP 300	Vector	14 GB		28 GF		59
2001	Compaq/HP Alphaserver (sc)	DEC Alpha ×512	0.5 TB	12TB	1 TF	2,000	31
2005	SGI Altix 3700 (ac)	Intel Itanium ×1920	5.5 TB	100 TB	14 TF	21,000	26
2009	SGI Altix XE (xe)	Intel Xeon ×1248	2.5 TB	90 TB	14 TF	12,000	
2009	Sun Constellation (Vayu)	Intel Xeon ×11, 936	37 TB	800 TB	140 TF	0.24M	35
2013	Fujitsu Primergy (Raijin)	Intel Xeon ×89, 256	300 TB	12.5 PB	1.2 PF	1.6M	24
2020	Fujitsu/Lenovo (Gadi)	Intel Xeon ×142, 152	1.2 PB	80? PB	9 PF	?	25

Notes:

1. This table edited by David Hawking to add data and to reformat for LaTeX.
2. Other ANU high-performance computers such as Fujitsu AP1000, Connection Machines CM2 and CM5 are not listed.
3. Abbreviations GF, TF, and PF relate to GigaFlops, TeraFlops, and PetaFlops.

H.1 The Early Years: 1987 – 2000

ANU was the first Australian university to take the strategic decision to establish high-end computing infrastructure and services to support research. Prior to 1987, ANU researchers, like those at many other universities, undertook their research computing on general purpose mainframe systems that were shared with the administrative and enterprise computing needs of the University. During the mid-1980s, however, driven by growing internal demands for access (principally from the astronomers and chemists), and the emergence of dedicated HPC systems in peer, international universities, ANU took the decision to allocate funds from its research budget to procure its first supercomputer—a Fujitsu/Facom VP100, following an approach to the market in 1987. Elsewhere in Australia, at that time, the only other (known) supercomputers were operated by the national science agency, CSIRO (Control Data Cyber 205), and the Bureau of Meteorology.

From the inception of the supercomputer, ANU established a dedicated support unit, within its Information Technology Division, known as the ANU Supercomputing Facility, with a number of “academic consultants” (initially three in number), to work with researchers to implement and optimise their codes on the supercomputer system. These academic consultants, who were all research-trained, would today be referred to as “computational scientists”.

Over the years, ANUSF grew under its foundation Head, Dr Bob Gingold, and its Academic Director, initially, Prof. Don Faulkner and subsequently Prof. Denis Evans. The relationship with Fujitsu evolved in a number of ways through research projects undertaken jointly by Fujitsu Laboratories and ANU in areas that included computational chemistry, linear algebra, and computer vision. Of these, the first of these continues through until today as a contracted research and development activity between Fujitsu’s Technical Computing Division and ANU.

Also over the years, the Fujitsu infrastructure was progressively upgraded through a VP2200 series vector machine in 1992, and a VPP300 series vector-scalar system in 1996, with the peak performance of the installation growing from 150 MFlops through 1.25 GFlops, to 28 GFlops for the final Fujitsu system. It is interesting to compare the performance of these very early systems with what is available in the consumer market today and to surmise what the future will hold in 20 years. In this particular instance, the second of these systems, the VP2200 has a peak performance which is slightly inferior to that of an Apple iPad2, according to benchmarking undertaken in 2011 by Prof. Jack Dongarra, of Linpack fame.

From the onset of services from the ANUSF facility, ANU made available 10% of its supercomputer system to researchers nationally. From the early 1990s, however, motivated by ANU’s example, a number of other Australia universities, either individually or through consortia, sought to establish comparable supercomputing facilities. Evident at the time were activities associated with state-based consortia in NSW and Queensland, and institutions such as the University of Melbourne, the University of Adelaide, and the University of Tasmania. In keeping with the times, there was a diversity of infrastructure installed, in many cases funded by the antecedent of the Australian Research Council, and partner organisations. Amongst these were the Connection Machine CM2 and CM5 systems (including one at ANU within its Centre for Information Science Research), an Intel Paragon, and various SGI scalar systems, e.g., the Power Challenge series.

As may be gleaned from the description above, Australia's approach to the development of research computing infrastructure, during the early 1990s, was quite balkanised, with the absence of a national policy to inform developments. The practical manifestation of this was the absence of an international class supercomputer, as measured by the Top500 list. Ultimately, in 1994, the Prime Minister's Science and Engineering Council (PMSEC) recommended a substantial investment to establish a facility of national scale. While funds of approximately $30M were set aside in 1994-95 by the Keating Government, the absence of an implementation plan saw the strategy languish for two more years.

Following the change of Government at the national government election of 1996, the financial exigencies at the time saw the pool of funds allocated to the acquisition of a national supercomputer reduced significantly. This was counterbalanced, however, by instructions from government that the Department implement a program to decide upon the implementation by the end of 1997. Accordingly, In the latter half of 1997, the Department of Employment, Education, Training and Youth Affairs (DEETYA) sought competitive bids, in effect a public tender, for the establishment of a "Centre of Excellence in High-Performance Computing and Communications". By the close of the tendering process, five applications had been submitted: from Queensland, Western Australia, Tasmania, South Australia, and the largest of the bids which was the ONLY, 'nearly national' proposal led by ANU, and involving every university of NSW and South Australia, the University of Queensland, and the University of Western Australia. The absence of Victorian universities from any of the bids, however, led to the abandonment of THIS competitive approach, and the subsequent initiation of a strategic approach that led ultimately to the next phase of development, referred to as the APAC years.

H.2 The APAC Years: 1999–2006

Early in 1998, following the abandonment of the competitive bidding process, on the advice of Departmental officials, the Minister for DEETYA, Dr David Kemp, put in train a process to develop a more inclusive, large-scale organisational model to create a genuinely national infrastructure. Accordingly, from early 1998, an advisory board, referred to as the Interim Board of the National High-Performance Computing Centre (chaired by Dr Michael Sargent), with representation from every state, and each of the research funding agencies was put in place, and tasked with the development of a proposal for a suitably national model to put to the Minister for funding. What emerged from the deliberations of the Interim Board, advised by an expert national committee, together with independent advice from the Director of the then (US) National Partnership for Advanced Computational Infrastructure, Dr Sid Karin, were recommendations for a framework for the Australian Partnership for Advanced Computing (APAC), the objectives of which were to:

- Fund a peak facility at ANU, underpinned by a partnership model to facilitate high-performance computing activity in every state, along with
- Expertise and education programs that could operate, both centrally and nationally.

These recommendations were accepted by the Minister, Senator Amanda Vanstone, who allocated $19.5M to initiate the APAC program. What was less well known is that the bulk of the initial funding for APAC was obtained from the Australian Research Council, which at the time was operated as a program with the Department. The decision to fund an out-of-program activity, such as APAC, from the ARC led to questions about the independence of the ARC, and contributed ultimately to the decision to establish the ARC as an independent authority with its own budget.

The recommendation to site the APAC National Facility at ANU was based on the University's leadership of, and its considerable previous investments in, HPC infrastructure, together with its willingness to share access to its facilities. The decision to host APAC from ANU, however, was rooted in the strong support by ANU's then Vice-Chancellor, Prof. Deane Terrell, that ANU could, and should, play a leading role in such national initiatives. ANU's successful carriage of the governance and management of APAC owes much to Professor Terrell, and to Professor Robin Stanton (then Pro Vice-Chancellor of ANU), who anchored the responsibilities of the host organisation.

Early in 1999, the Board of APAC was established under the chairmanship of Professor David Beanland, and the appointment of the foundation (and sole) Executive Director, Professor John O'Callaghan followed soon thereafter. APAC was formally launched late in 1999, having initiated the process by which the:

- Partnership could be established, through proposals from organisations, and consortia of organisations;
- National Facility, i.e., the national peak system, could be established at ANU.

By 2000-01, APAC as a "partnership of partnerships" was largely in place, with the membership being:

- Australian National University, as Host;
- Commonwealth Scientific and Industrial Research Organisation;
- Australian Centre for Advanced Computing and Communications (ac3), a consortium representing the universities of NSW;
- Queensland Parallel Supercomputing Foundation (QPSF, later QCIF), a consortium representing the universities of Queensland;

- South Australian Partnership for Advanced Computing, a consortium representing the universities of South Australia;
- Tasmanian Partnership for Advanced Computing (TPAC), a consortium led by the University of Tasmania;
- Victorian Partnership for Advanced Computing (VPAC), a consortium of Victorian universities;
- iVEC, a consortium of CSIRO and the universities of Western Australia.

The National Facility was ultimately commissioned early in 2001, following some early delays in the procurement process. With the installation of a Compaq/HP Alphaserver SC system of 1 TFlop capacity, Australia, at last, had an internationally competitive system, benchmarked at rank 31 in its debut on the Top500 list.

With the establishment of services, APAC operated in two phases, respectively from 2001–04, and from 2004–07. Its first phase, funded largely by monies from the ARC, comprised the:

- Establishment of services from the National Facility – which were delivered from the ANU Supercomputing Facility;
- Building of national expertise in, and the uptake of, the use of advanced computing in research, through the partnership.

Following the commencement of services from the National Facility, ANU decommissioned its ageing Fujitsu system, and subsequently serviced all of its high-end computing requirements from the National Facility, through a substantial access share which it paid for through a combination of cash and in-kind contributions, embodied in the hosting arrangements for the NF system. Each of the APAC partners, subscribed to an access share of the NF system (although, none as large as ANU), complementing their local facilities. The largest share of the NF (approximately 50%) was allocated to the APAC Merit Allocation Scheme, an open access allocation process, rooted in ANU's earlier processes, and strongly aligned with research merit and computational suitability.

APAC also established capacity and capability programs in each of its partners, through its Centres of Expertise programs (in which each partner took leadership in a dedicated field or application area), as well as through its Computational Tools and Techniques program.

The second phase of APAC operations (2004—07) were funded through the System Infrastructure Initiatives of the Commonwealth Government. This entailed the:

- Refreshing of the National Facility peak system; and
- Nurturing of a nascent national eResearch program nationally through its Grid and Collaboration Tools program, and the beginnings of a data and storage program.

With regard to the former, the Compaq/HP system was replaced in 2005 by a 14 TFlop SGI Altix 3700 system, which ranked as 26 on its debut on the Top500 list.

Throughout the APAC years, operational services for the peak facility were provided through the ANU Supercomputing Facility, which was accountable to Executive Director of APAC through the APAC Agreement. During the early years, ANUSF was led by its foundation Head, Dr Bob Gingold, and subsequently by his successor, Dr Ben Evans, following Dr Gingold's retirement.

Towards the end of this second phase of APAC, a Department of Education, Science and Tourism (DEST) commissioned a review of APAC in advance of further infrastructure investments that were to be conceptualised under a new, collaborative framework, known eventually as NCRIS (National Collaborative Research Infrastructure Strategy). An international review noted the considerable success of the APAC program, including the high-quality of service from, and the internationally competitive ranking of, the peak system, the cultivation of a national community of researchers and support specialists, the embedding of high-level computing expertise across a bread of application areas, the nurturing of grid expertise and collaboration tools, and the fostering of engagement and commitment from government at state and federal level.

With the implementation of NCRIS from 2007 onwards, the Australian Government set in place a funding allocation to promote the uptake of eResearch techniques across the span of research, in a manner broader to that of the eScience programs in some other nations. While the APAC model had been a considerable success, the strategic framework for Platforms for Collaboration was conceptualised differently, with distinct, investment programs for:

- Interoperability and Collaboration Infrastructure—through the Australian Research Collaboration Service (ARCS),
- Data Infrastructure — through the Australian National Data Service (ANDS),
- High-performance Computing — through the National Computational Infrastructure (NCI),

in addition to other, smaller investments in authentication and authorisation, and networking. These are referred to in the following section, which covers the NCI years.

H.3 The NCI Years: 2007 –

The NCRIS investment framework for eResearch was progressively implemented from 2007 onwards, with the signing of funding agreements by the Commonwealth Government and the lead organisations for each of NCI (with ANU), ARCS (with VPAC), and ANDS (with Monash University). Indeed, it can be argued that APAC spawned the initial activities of each, with:

- NCI conceptualised as solely a supercomputing facility, built from the base of the successful APAC National Facility program;

- ARCS inheriting the grid and collaboration program activities of APAC, together with a governance structure based around the original APAC partners; and
- ANDS building data discovery, access and publication services from the early roots established during the APAC years.

For NCI, its evolution has occurred in two phases, that are not entirely independent, and in ways that are significantly different to that of its antecedent.

The First Phase

The Funding Agreement that established NCI (known as NCRIS NCI Project) was executed by ANU and the Commonwealth Government in June 2007 for an amount of $26M, with the objectives being to:

- Procure/sustain a capability computing system commensurate with international standing;
- Advise Government on further "specialised system" investments to provide for targeted application areas;
- Establish/operate a merit-based open access allocation scheme for researchers at publicly-funded research organisations;
- Provide support and expertise services; and
- Maintain a supported strategic plan for national computational needs.

With the regard to the major system acquisition, the goal at the time was to procure a capability system, the performance of which would be comparable to that of the "Track 2" systems being established in the USA, at the time, by the National Science Foundation. Quite early, it became clear that, with only the funds available, this would be a difficult goal to realise without the injection of further investment/co-investment. And so began the process of building the partnership and governance model that exists today, under the leadership of a Steering Committee chaired by Emeritus Professor Mark Wainwright, a former Vice-Chancellor of UNSW, and through the office of Professor Robin Stanton, a Pro Vice-Chancellor at ANU, and the ANU Delegate of the NCRIS NCI Contract.

With the Steering Committee (which was the antecedent of today's Board) established in 2007, and expanded in 2008 to include, as institutional members, three of the national agencies—CSIRO, the national science agency, the Bureau of Meteorology, and Geoscience Australia—alongside ANU, as the host organisation, the first steps were taken to establish the sustaining partnership that is now in place. With the sails set, ANU's delivery on the goals of the contract was initiated with the appointment of the foundation Director of NCI, Professor Lindsay Botten in 2008. As in the APAC era, all services in this first phase of NCI were provided through the ANU Supercomputing Facility, which was accountable to the NCI Director and to Steering Committee.

Early in 2008, at the time at which ANU was about to tender for a new peak system (under the NCRIS agreement), the Bureau of Meteorology was planning to approach the market for a new operational weather forecasting system. With strong synergies apparent, it was agreed that BoM and ANU would issue a joint tender, with the goal of procuring interoperable systems that would facilitate research opportunities and collaboration in climate and weather science, and which also had the potential to enhance the service robustness for operational weather forecasting. This procurement took place under the governance of a Joint Steering Committee comprising BoM, ANU/NCI, and CSIRO representatives. The evaluations were completed by October 2008, at which time the world was then in the tightening grip of the Global Financial Crisis (GFC)—which spanned the entire period of the contract negotiations with the successful tenderer, Sun Microsystems.

A contract was executed by ANU and Sun in March 2009, for a 16 rack, 12,000 core, 140 TFlop Sun Constellation known as Vayu—a distributed memory cluster based on Intel Xeon Nehalem technology with a QDR Infiniband interconnect, with the US dollar exchange rate having dropped by approximately 25 per cent during contract negotiations, thereby impacting the scale of the procurement. The first phase (one-eighth) of the Sun system at NCI was installed in September 2009, more than replacing the capability of previous SGI Altix 3700 system, with subsequent upgrades in February and April of 2010 bringing the system to its full capacity. Concurrently, BoM completed its negotiations with Sun at the same time, and its five rack system, with dual-rail Infiniband, entered production during 2010.

The emerging NCI collaboration took a most significant step forward late in 2008 with a decision by CSIRO to enter into partnership with ANU as the initial custodians of NCI—with CSIRO committing to its ongoing, strong level of co-investment in NCI by executing with ANU the Partner Service Agreement on 24 December 2008. QCIF and iVEC joined as minor partners shortly thereafter in 2009, followed in 2010 by Geoscience Australia, and Intersect (the NSW university consortium), which was able to repurpose a substantial ARC Linkage Infrastructure Grant for 2009 (led by the University of Technology, Sydney) to take services in NCI from 2010–14, rather than to procure a separate,standalone system for its consortium.

The year 2008—09 also saw the first steps taken to increase the diversity of computational resources available to Australian researchers, with the decision by the NCI Steering Committee to invest in two Specialised Facilities, in keeping with one of the goals of the NCRIS contract. Following an open call for expressions of interest from research computing facility operators, the decision was taken to invest, alongside CSIRO, in each of the:

- Specialised Facility in Bioinformatics, located at the University of Queensland—with institutional partners CSIRO, UQ, QCIF, QFAB, the State Government of Queensland and NCI, and
- Specialised Facility in Imaging and Bioinformatics, located at Monash University (and known now as MASSIVE) – with institutional partners CSIRO, Monash University, the Australian Synchrotron, the State Government of Victoria, VPAC and NCI.

Contracts were signed with the lead agents for each facility in December 2009, with preliminary services available from mid-2010, and with a full production service made available to researchers through the NCI-funded share under its Merit Allocation Scheme from 2011–13, in the first instance.

The Second Phase

The second phase of NCI's development commenced relatively early, with announcements in the Commonwealth Government Budget of May 2009 of the new Super Science initiatives that formed part of Government's economic stimulus package, overlapping the implementation of the first phase under the NCRIS agreement. The 2009–10 Budget contained four major announcements for e-infrastructure in Australia:

- An allocation of $97M to support the implementation of data, cloud and collaboration infrastructure, to be taken forward through ARCS;
- An allocation of $50M for the continuation of ANDS;
- An allocation of $80M for the establishment of high-performance computing infrastructure to support radioastronomy and Australia's bid for the SKA—from which emerged what is the Pawsey Centre through the iVEC consortium in Western Australia;
- An allocation of $50M for the Climate HPC Centre Project through ANU—which ultimately has been implemented through the NCI Collaboration.

The first of these did not proceed in the form originally planned, with ARCS being discontinued in mid-2011 at the conclusion of its NCRIS funding, and with the implementation of the project being reconceptualised into the two projects that are today known as:

- NeCTAR (National eResearch Collaboration Tools and Resources) Project, through the University of Melbourne ($47M), and
- RDSI (Research Data Storage Infrastructure) Project ($50M), through the University of Queensland.

Each of the Super Science investments, including the Climate HPC Centre Project through ANU/NCI, was funded from the Education Investment Fund, one of Australia's Nation Building Funds, with this choice bringing special and challenging constraints that required government funding be used solely for infrastructure procurements and developmental activities, precluding its use for operations or recurrent costs.

The project objectives were:

- The establishment of a petascale HPC facility prioritised to the needs of research in climate change, earth system science, national water management, which simultaneously would support meritorious and high-impact research in all fields that required access to capability-class computing;
- The identification and support of data-intensive and flagship science application aligned with other government research and infrastructure investments;
- The development and implementation of an access model that would meet the priority use requirements, as well as providing open access on research merit to researchers at publicly-funded research organisations;
- The construction of a purpose-built data centre to house the HPC facility, capable which could be upgraded to handle future usage for at least five years;

Implicit in these, and the strong constraint on the use of the funds, were that ANU, as the contract holder, had to demonstrate to the Commonwealth its capacity to meet the recurrent costs of a supercomputer (of an appropriate scale), before commencing the construction of the data centre and the initiation of the system procurement. With all the substantial recurrent costs having to be met through the co-investment of stakeholders, it took some time to form and cement the partnership, and to establish the business plan and access model, before moving ahead with the establishment of the contracted infrastructure. Slightly more than two years elapsed from the Budget announcement date, to the commencement of construction of the data centre and the release of the tender for the petascale system – with major intermediate milestones being the execution of the Funding Agreement with the Commonwealth (April 2010), and the execution of the NCI Collaboration Agreement (July 2011).

The Collaboration Agreement lies at the heart of NCI partnership and sets out the framework and objectives for the collaboration, puts in place its governance, lays out the access model, and underpins NCI's business model and planning. With the execution of the Collaboration Agreement by the initial partners, ANU, CSIRO and BoM in July 2011, the University was able to demonstrate to the Australian Government that it had secured at least $8M per annum for 2012–15 (i.e., sufficient to mount the operations of a supercomputer of an appropriate scale), at which point the green light was given for the commencement of construction of the new data centre, and the initiation of the public tender for the HPC system, both in August 2011. The process of building the Collaboration continued and by early 2012, it had expanded to include Geoscience Australia, Intersect Australia (the NSW consortium), QCIF (the Queensland consortium), and a new consortium of six research intensive universities (ANU, Adelaide, Monash, USNW, Queensland, Sydney) whose participation in NCI was facilitated by a substantial ARC Linkage Infrastructure Grant led by ANU (on behalf of NCI), with approximately $11M per annum having been secured. This partnership, which includes four national organisations, together with the high leveraging of the Commonwealth's infrastructure monies from the Government through substantial co-investment, are two of most distinctive features of NCI.

The Collaboration Agreement in setting the foundations for the implementation of this current phase of NCI also formed the backdrop for an organisational change, from the beginning of 2012, by which the University brought together the NCI Project Office and the ANU Supercomputing Facility into a single operating unit (known as NCI-ANUSF internally, and NCI externally), and operating within with a governance model by which ANU governs the operations of NCI on the advice of the NCI Board, to within the limits of its Statutes and policies.

With the Collaboration Agreement in place from July 2011, the data centre construction was initiated in August 2011 under the managing contractor, G. E. Shaw and Associates, and was completed in September 2012, with the building formally handed over to ANU in November of that year. Concurrently, the tender for the HPC system also proceeded through a public tender from August 2011, closing in late October. The multi-stage evaluation process continued through until late-April 2012, with the Board accepting the final evaluation report in May 2012, and recommending to the University that it proceed with the procurement of a 1.2 petaflop, distributed memory cluster from Fujitsu. The contract for the procurement was executed in mid-June 2012, leading to the delivery of the infrastructure from about September 2012 onwards, and ultimately to the debut of the system, now known as Raijin, on the Top500 list in November 2012 at rank 24, with a peak performance (Rmax) of 940 TFlops (for the 93 per cent of the system that had been implemented at the time). In the months that followed, the DDN filesystem was commissioned, and system software evaluated and configured, culminating in an early user service from late April/early May 2013, and with a full user service from mid-June 2013. Also associated with the Fujitsu contract is a most significant collaboration framework which is being implemented to take NCI and its partners into the future, and, in particular, to deliver optimal value for the peak system procurement, through a program of work to optimise the implementation and performance of Australian climate and earth system science modelling suite, ACCESS, on Raijin, and which will investigate the implementation of today's codes on the processor architectures that will likely be a feature of the next-generation systems.

From relatively early in the history of NCI (from around 2009 onwards), it was apparent to the Steering Committee, and subsequently the Board, that the conceptualisation of government investments at the infrastructure layer (i.e., HPC, cloud, data) was orthogonal to the realisation of research outcomes through a solutions-based approach. Thus, NCI began its path towards the comprehensive and integrated framework that underpins its delivery of high-performance solutions. The initial activities, under NCI, in cloud computing and data-intensive computing, complementing its established role in high-performance computing, ramped up from 2009–10, with substantial upgrades to, and modernisation of, the storage infrastructure occurring in 2011–12 with a SGI solution.

Building on this framework, NCI's Board took the strategic decision in 2011 that the NCI Collaboration, which had been established to support operations of a petascale supercomputer, should broaden its scope to provide the comprehensive and integrated services that had long been envisaged. By that time, the RDSI and NeCTAR projects were underway, with both activities seeking proposals from national eResearch organisations to become nodes of each of the national storage infrastructure network, and the national research cloud. NCI submitted comprehensive proposals to each of RDSI and NeCTAR, with each built on a strong research community engagement, and with the goal of realising research outcomes through complementary high-performance solutions. The build-out of the integrated infrastructure solution was undertaken during 2013, with storage procurements from SGI and DDN, and with the establishment of a 3,200 core high-performance cloud from Dell (architected with a focus on data-intensive applications within the NeCTAR OpenStack Cloud Federation).

The story since 2013 is told incrementally in NCI's annual reports, found online at `http://nci.org.au/about-nci/annual-reports/`.[1] These provide updates on Australian research outcomes and impacts that have been enabled by the NCI platforms and its expert support team, and the evolution of the infrastructure platform—with further incrementation of the storage from Netapp in 2015, a 40% upgrading of NCI's supercomputer capacity in 2016 with a 22,000 core Lenovo NeXtScale (Broadwell) system co-funded by a substantial allocation from the NCRIS Agility Fund and the NCI Collaboration, together with major storage upgrades funded from the same source.

Winding the clock forward to 2017, NCI has evolved as the national, high-end research computing service – one which is in the vanguard of international advanced computing, delivering solutions that encompass computational modelling and the needs of big data, that enable research of excellence and impact, which deliver in the national benefit, and which maintain Australian's international competitiveness in research and innovation.

The partnership that underpins NCI has grown in strength since its formation, anchored by the Australian National University as NCI's host, and three of Australia's national science agencies—the Australian Bureau of Meteorology, the national science agency, CSIRO, and the national geoscience agency, Geoscience Australia. Today, this partnership includes many of Australia's research universities, or consortia thereof, and medical research institutes, and sustains two-thirds of the annual recurrent costs ($18 million in 2017), with contributions from the Australian Government's National Collaborative Research Infrastructure Strategy (NCRIS) providing the remainder.

Through its tightly-coupled, high-performance computing and data platforms, overlaid with internationally renowned expertise (60 staff) in computational science, data science and data management, NCI provides essential services that underpin the requirements of research and industry, today and into the future.

As of May 2017, the infrastructure platform comprises:

- Supercomputer, Raijin — a hybrid Fujitsu Primergy (2012-13, Intel Xeon Sandy Bridge) and Lenovo NeXtScale (2016, Intel Xeon Broadwell) system of 84,656 cores in 4416 compute nodes (together with a number of nVIDIA Tesla K80 and P100 GPUs, and Intel Xeon Phi Knights Landing nodes), with 300 terabytes of main memory, a hybrid FDR/EDR Mellanox Infiniband full fat-tree interconnect, and 8 petabytes (real) of high-performance (150 Gbyte/sec) operational storage — aggregated peak performance 2.1 petaflops
- OpenStack Cloud — Dell, 3,200 cores (Intel Xeon Sandy Bridge), 25 TB of memory, 160 TB of SSD, 13 ceph nodes

[1] This URL was no longer active at 26 March 2021. It may be possible to find it on the Wayback Machine.

- Project and Collection Storage — totalling 22 (real) petabytes in three distinct Lustre filesystems of 50 Gbyte/sec (SGI), 70 Gbyte/sec (DDN) and 120 Gbyte/sec (Netapp) bandwidth — to be upgraded to 36 (real) petabytes in June/July 2017, with the replacement of the oldest of the filesystems with NetApp/Fujitsu and HPE storage.
- Tape storage — comprising a dual site archive of 2 x 12.3 PByte capacity in Spectra libraries using LTO-5 technology, and HSM/Lustre of 2 x 18.2 Pbyte capacity using Spectra libraries and IBM TS1140/50 technology drives.

Today, NCI plays a pivotal role in Australia's research and innovation system, supporting the work of more than 5,000 researchers across more than 500 projects being undertaken in 35 universities, 5 national science agencies (including CSIRO, Bureau of Meteorology, and Geoscience Australia), 3 medical research institutes, and industry. The breadth and depth of NCI's involvement in Australian research and innovation can be seen in case studies associated with Australia's national science and research priorities (`http://nci.org.au/research-news/nci-today-case-studies/`), and in the numerous research highlights (`http://nci.org.au/research-news/research/`).[2]

R&D which is reliant on NCI spans the full spectrum—from fundamental research, through the strategic and applied, and on to actual industrial applications. Today, NCI services are both necessary and influential in enhancing the competitiveness and impact of outcomes in every field of science and technology. An increasing number of fields are now highly dependent on the fusion of "big compute" and "big data" that NCI provides—in weather and climate science, the earth sciences, earth observation and remote sensing, medical research, and astronomy, amongst others.

Within the university sector, NCI provides the essential high-performance computing and data foundation for more than 200 research projects, ARC and NH&MRC Centres of Excellence and Industry Hubs, and fellowships. Funding for these projects from the Australian Research Council and the National Health and Medical Research Council totals around $60 million per annum, or approximately $250 million over the lifetimes of these projects.

In the domain of the national science agencies, notably BoM, CSIRO and GA, NCI provides critical program-level support in earth system science, by serving as the development platform of the Australian national weather and climate modelling suite, ACCESS. NCI is also a national hub for major national and international satellite earth observation collections (through the Australian Geoscience Data Cube), used in the earth, marine and environmental sciences, agriculture for research, the informing of policy development, and the development of important information products for primary industry.

At the time of this update in May 2017, the Australian research sector is awaiting the release of the National Research Infrastructure Roadmap (2016–17) and its subsequent funding and implementation plan. In the realm of HPC, the draft Roadmap has highlighted the urgency of funding for the renewal of national HPC, and has acknowledged that peak computing facilities must "encompass the needs of big data (processing, analysis, data mining, machine learning), in addition to its traditional role of computational modelling and simulation … compris[ing] tightly-integrated, high-performance infrastructure able to handle the computational and data-intensive workflows of today's research, together with expertise in computational science, data science and data management." – the directions in which NCI has evolved since 2012.

NCI looks forward to the next phase of its development in this environment, and the strengthening of its role in Australia's national research and innovation system.

[2] Both of these pages have also disappeared.

Appendix I

The History of Samba

```
Full Samba History
This is from a history file I wrote in 1997 for Samba

Contributor:    Andrew Tridgell and the Samba Team
Date:           June 27, 1997
Status:         Always out of date! (Would not be the same without it!)

Subject:        A bit of history and a bit of fun
============================================================================
```

This is a short history of this project. It's not supposed to be comprehensive, just enough so that new users can get a feel for where this project has come from and maybe where it's going to.

The whole thing really started in December 1991. I was (and still am) a PhD student in the Computer Sciences Laboratory at the Australian National University, in Canberra, Australia. We had just got a beta copy of eXcursion from Digital, and I was testing it on my PC. At this stage I was a MS-DOS user, dabbling in windows.

eXcursion ran (at the time) only with Dec's 'Pathworks' network for DOS. I had up till then been using PC-NFS to connect to our local sun workstations, and was reasonably happy with it. In order to run pathworks I had to stop using PC-NFS and try using pathworks to mount disk space. Unfortunately pathworks was only available for digital workstations running VMS or Ultrix so I couldn't mount from the suns anymore.

I had access to a a decstation 3100 running Ultrix that I used to administer, and I got the crazy notion that the protocol that pathworks used to talk to ultrix couldn't be that hard, and maybe I could work it out. I had never written a network program before, and certainly didn't know what a socket was.

In a few days, after looking at some example code for sockets, I discovered it was pretty easy to write a program to "spy" on the file sharing protocol. I wrote and installed this program (the sockspy.c program supplied with this package) and captured everything that the pathworks client said to the pathworks server.

I then tried writing short C programs (using Turbo C under DOS) to do simple file operations on the network drive (open, read, cd etc) and looked at the packets that the server and client exchanged. From this I worked out what some of the bytes in the packets meant, and started to write my own program to do the same thing on a sun.

After a day or so more I had my first successes and actually managed to get a connection and to read a file. From there it was all downhill, and a week later I was happily (if a little unreliably) mounting disk space from a sun to my PC running pathworks. The server code had a lot of 'magic' values in it, which seemed to be always present with the ultrix server. It was not till 2 years later that I found out what all these values meant.

Anyway, I thought other people might be interested in what I had done, so I asked a few people at uni, and noone seemed much interested. I also spoke to a person at Digital in Canberra (the person who had organised a beta test of eXcursion) and asked if I could distribute what I'd done, or was it illegal. It was then that I first heard the word "netbios" when he told me that he thought it was all covered by a spec of some sort (the netbios spec) and thus what I'd done was not only legal, but silly.

I found the netbios spec after asking around a bit (the RFC1001 and RFC1002 specs) and found they looked nothing like what I'd written, so I thought maybe the Digital person was mistaken. I didn't realise RFCs referred to the name negotiation and packet encapsulation over TCP/IP, and what I'd written was really a SMB implementation.

Anyway, he encouraged me to release it so I put out "Server 0.1" in January 1992. I got quite a good response from people wanting to use pathworks with non-digital unix workstations, and I soon fixed a few bugs, and released "Server 0.5" closely followed by "Server 1.0". All three releases came out within about a month of each other.

At this point I got an X Terminal on my desk, and I no longer needed eXcursion and I prompty forgot about the whole project, apart from a few people who e-mailed me occasionally about it.

Nearly two years then passed with just occasional e-mails asking about new versions and bugs. I even added a note to the ftp site asking for a volunteer to take over the code as I no longer used it. No one volunteered.

During this time I did hear from a couple of people who said it should be possible to use my code with Lanmanager, but I never got any definite confirmation.

One e-mail I got about the code did, however, make an impression. It was from Dan Shearer at the university of South Australia, and he said this:

> I heard a hint about a free Pathworks server for Unix in the Net channel of the Linux list. After quite a bit of chasing (and lots of interested followups from other Linux people) I got hold of a release news article from you, posted in Jan 92, from someone in the UK.
>
> Can you tell me what the latest status is? I think you might suddenly find a whole lot of interested hackers in the Linux world at least, which is a place where things tend to happen fast (and even some reliable code gets written, BION!)

I asked him what Linux was, and he told me it was a free Unix for PCs. This was in November 1992 and a few months later I was a Linux convert! I still didn't need a pathworks server though, so I didn't do the port, but I think Dan did.

At about this time I got an e-mail from Digital, from a person working on the Alpha software distribution. He asked if I would mind if they included my server with the "contributed" cd-rom. This was a bit of a shock to me as I never expected DEC to ask me if they could use my code! I wrote back saying it was OK, but never heard from him again. I don't know if it went on the cd-rom.

Anyway, the next big event was in December 1993, when Dan again sent me an e-mail saying my server had "raised its ugly head" on comp.protocols.tcpip.ibmpc. I had a quick look on the group, and was surprised to see that there were people interested in this thing.

At this time a person from our computer centre offered me a couple of cheap ethernet cards (3c505s for $15 each) and coincidentially someone announced on one of the Linux channels that he had written a 3c505 driver for Linux. I bought the cards, hacked the driver a little and setup a home network between my wife's PC and my Linux box. I then needed some way to connect the two, and I didn't own PC-NFS at home, so I thought maybe my server could be useful. On the newsgroup among the discussions of my server someone had mentioned that there was a free client that might work with my server that Microsoft had put up for ftp. I downloaded it and found to my surprise that it worked first time with my 'pathworks' server!

Well, I then did a bit of hacking, asked around a bit and found (I think from Dan) that the spec I needed was for the "SMB" protocol, and that it was available via ftp. I grabbed it and started removing all those ugly constants from the code, now that all was explained.

On December 1st 1993 I announced the start of the "Netbios for Unix" project, seeding the mailing list with all the people who had e-mailed me over the years asking about the server.

About 35 versions (and two months) later I wrote a short history of the project, which you have just read. There are now over a hundred people on the mailing list, and lots of people report that they use the code and like it. In a few days I will be announcing the release of version 1.6 to some of the more popular (and relevant) newsgroups.

Andrew Tridgell
6th February 1994

It is now May 1995 and there are about 1400 people on the mailing list. I got downloads from the main Samba ftp site from around 5000 unique hosts in a two month period. There are several mirror sites as well. The current version number is 1.9.13.

It's now March 1996 and version 1.9.16alpha1 has just been released. There have been lots of changes recently with master browser support and the ability to do domain logons etc. Samba has also been ported to OS/2, the amiga and NetWare. There are now 3000 people on the samba mailing list.

It's now June 1997 and samba-1.9.17 is due out soon. My how time passes! Please refer to the WHATSNEW.txt for an update on new features. Just when you think you understand what is happening the ground rules change – this is a real world after all. Since the heady days of March 1996 there has been a concerted effort within the SMB protocol using community to document and standardize the protocols. The CIFS initiative has helped a long way towards creating a better understood and more interoperable environment. The Samba Team has grown in number and have been very active in the standards formation and documentation process.

The net effect has been that we have had to do a lot of work to bring Samba into line with new features and capabilities in the SMB protocols.

The past year has been a productive one with the following releases: 1.9.16, 1.9.16p2, 1.9.16p6, 1.9.16p9, 1.9.16p10, 1.9.16p11

There are some who believe that 1.9.15p8 was the best release and others who would not want to be without the latest. Whatever your perception we hope that 1.9.17 will close the gap and convince you all that the long wait and the rolling changes really were worth it. Here is functionality and a level of code maturity that ..., well – you can be the judge!

Happy SMB networking!
Samba Team
ps: The bugs are ours, so please report any you find.

It's now October 1998. We just got back from the 3rd CIFS conference in SanJose. The Samba Team was the biggest contingent there.

Samba 2.0 should be shipping in the next few weeks with much better domain controller support, GUI configuration, a new user space SMB filesystem and lots of other neat stuff. I've also noticed that a search of job ads in DejaNews turned up 3900 that mention Samba. Looks like we've created a small industry.

I've been asked again where the name Samba came from. I might as well put it down here for everyone to read. The code in Samba was first called just "server", it then got renamed "smbserver" when I discovered that the protocol is called SMB. Then in April 1994 I got an email from Syntax, the makers of "TotalNet advanced Server", a commercial SMB server. They told me that they had a trademark on the name SMBserver and I would have to change the name. I ran an egrep for words containing S, M, and B on /usr/dict/words and the name Samba looked like the best choice. Strangely enough when I repeat that now I notice that Samba isn't in /usr/dict/words on my system anymore!

Appendix J

Dave’s vantage points

For the record, I viewed ANU computer science and computing at ANU from the following perspectives:

1971	Arrived at ANU, programmed FORTRAN in the Pure Maths dept.
1972 – 1973	Undergraduate computer science + psychology + physics
1974	In first computer science honours cohort
1975 – 1979	Tutor in computer science
1977	Enrolled for PhD as a Staff Candidate (Computer Animation, never completed)
1979 – 1983	Site Manager of ANU student computing facilities (SGS DEC-10, Faculties Computer Unit)
1983 – 1988	Computer Services Centre evangelist for text processing, laser printing, networking and microcomputers
1988 – 1998	Head Programmer, Computer Science Department.
1991 – ?	Manager, Fujitsu AP1000 system; Application Engineer, Connection Machine CM2
1993 – 1995	Roles in the design and construction of the CS&IT building.
1994 – 2000	Part-time secondment to ACSys Cooperative Research Centre.
1998	Completed PhD by Published Work
1998 – 2008	Adjunct position in DCS while CSIRO Senior Principal Research Scientist in CS&IT building
2008 – 2013	Adjunct / Honorary Level E while Chief Scientist at Funnelback Pty Ltd
2013 – 2018	Adjunct / Honorary Level E while Partner Architect at Microsoft Inc
2018 –	Honorary Level E post retirement.

I supervised eight PhD students in DCS and one from the Department of Engineering.

My education in computer science enabled me to meet many influential people: I was interviewed by Satya Nadella and had a public argument with Larry Page. I wrote a Canberra Times feature on Richard Stallman. I had breakfast with Bill Gates, morning tea with Edsger Dijkstra, lunch with Tim Berners-Lee, and afternoon tea with Linus Torvalds. Proposals for dinner will be considered on their merits.

Index of People

Index of Hardware & Networks

Index of Software & Services

ANU's first computer, an IBM 610 acquired by Mt Stromlo Observatory in 1960, revolutionised the astronomical calculations performed by Claire Wehner for Australia's National Time Service. In 1962, Ken Le Couteur, who had used the Colossus computers at Bletchley Park, acquired an IBM 1620 for the Theoretical Physics Department. Elizabeth Reid (yes, that Elizabeth Reid) was its first Operator/ Programmer and Brian Robson was in charge.

Undergraduate teaching of Computer Science commenced in March 1971, when Ray Jarvis led the Computer Science sub-department of the Statistics Department in the Faculty of Economics. The introduction and subsequent expansion of Computer Science was led by student demand but strongly opposed by other departments in times of static or diminishing ANU resources.

In 2021, ANU's GADI supercomputer may be 13 orders of magnitude faster than the IBM 610 and undergraduate loads in Computer Science have grown by a factor of 13.

This history tells the story of the very dramatic transition from then to now, complete with disasters and triumphs, threats of amalgamation, fruitful collaborations, heroic achievements, and revolutionary change in all the university's processes. It tells of wonderful machines and software, dramatic re-organisations and most of all fascinating people. More than 100 people have contributed to the story.

"This is quite a superb book! Congratulations! It's all that I expected, and more. I loved the scene-setting stuff in the 1970s! Every page brings a smile; it really is a disarming and affectionate epistle."

Henry Gardner, former Head, DCS