

Formally Verified Compositional Algorithms for Factored Transition Systems

Mohammad Ahmad Abdulaziz Ali Mansour

**A thesis submitted for the degree of
Doctor of Philosophy of
The Australian National University**

June 2018

Declaration

Except where otherwise indicated, this thesis is my own original work.

Mohammad Ahmad Abdulaziz Ali Mansour
27 June 2018

Acknowledgement

Firstly, I would like to thank god who blessed me with the health, will, and ability to do the work in this thesis.

Secondly, I would like to thank Data61 (formerly NICTA) and the Australian National University for the PhD scholarship as well as the multiple travel grants, without which I would have not been able to conduct my research.

Thirdly, I would like to thank my supervisors, Dr. Charles Gretton and Dr. Michael Norrish, who were mentors, as well as friends to me. Our meetings and style of interaction were intellectually very stimulating and enjoyable, and in addition to their rich technical content, our discussions spanned different topics like politics, history, and culture. Dr. Michael Norrish introduced me to the basics of interactive theorem proving and spent countless hours with me in meetings helping me through the different hurdles I faced during my research. Also, without his encouragement and leveraging of his connections, it would have not been possible for me to proceed with my internships. Dr. Charles Gretton's help and guidance made it much easier to get habituated to the literature and the different concepts of artificial intelligence planning. I would also like to thank him for his relentless effort in reading and editing all the drafts I wrote in the course of my PhD studies. I would also like to thank my advisor Prof. Sylvie Thiébaux who helped gauge my bearings research-wise and career-wise, at many points during my PhD program.

In addition to my supervisors, I would like thank different researchers with whom I had very useful interactions related to the work in this thesis. I would like to thank Prof. Brendan McKay who helped me understand many concepts in group theory and graph automorphism, and formulate and prove theorems that are crucial for significant parts of this thesis. I would also like to thank Dr. Alban Grastien for our many engaging discussions and for his very useful feedback on many parts of my research, as well as for reading many drafts that I wrote through my research. I would also like to thank Dr. Patrik Haslum for his useful feedback and suggestions for my research and on some of my drafts, and for being available to answer many of my questions on the fundamentals and practicalities of AI planning. I would like to thank Prof. Daniel Jackson for his suggestion of the hotel key protocol as an example to try the upper-bounding algorithms on. I would also like to thank Dr. Miquel Ramírez for reading one of the drafts I wrote.

I would also like to thank the researchers who hosted me during my different internships and research visits. Firstly I would like to thank Prof. Lawrence Paulson for hosting me in the University of Cambridge Computer Laboratory, for his patience and generosity with his time and effort, and for our consequent fruitful and joyful collaboration. Secondly, I would like to thank Dr. Bruno Dutertre, Dr. Dejan Jovanovic, and Dr. Natrajan Shankar for hosting me in the Computer Science Laboratory in SRI International and for the very helpful discussions I had with them. From the Computer Science and Artificial Intelligence Laboratory in Massachusetts Institute of Technology, I would like to thank Prof. Brian Williams and Dr. David Wang who hosted me during my visit there.

Last, but not least, I would like to thank my family and my wife for their unlimited help and support in general, and especially during my PhD program. I would also like to thank my comrade PhD candidates: Josh, Jan, Fazlul, Karsten, Mazen, and Zuhair, for the many chats we had, technical or otherwise.

Abstract

Artificial Intelligence (AI) planning and model checking are two disciplines that found wide practical applications. It is often the case that a problem in those two fields concerns a transition system whose behaviour can be encoded in a digraph that models the system's state space. However, due to the very large size of state spaces of realistic systems, they are compactly represented as propositionally factored transition systems. These representations have the advantage of being exponentially smaller than the state space of the represented system.

Many problems in AI planning and model checking involve questions about state spaces, which correspond to graph theoretic questions on digraphs modelling the state spaces. However, existing techniques to answer those graph theoretic questions effectively require, in the worst case, constructing the digraph that models the state space, by expanding the propositionally factored representation of the system. This is not practical, if not impossible, in many cases because of the state space size compared to the factored representation.

One common approach that is used to avoid constructing the state space is the compositional approach, where only smaller abstractions of the system at hand are processed and the given problem (e.g. reachability) is solved for them. Then, a solution for the problem on the concrete system is derived from the solutions of the problem on the abstract systems. The motivation of this approach is that, in the worst case, one need only construct the state spaces of the abstractions which can be exponentially smaller than the state space of the concrete system.

We study the application of the compositional approach to two fundamental problems on transition systems: upper-bounding the topological properties (e.g. the largest distance between any two states, i.e. the diameter) of the state space, and computing reachability between states. We provide new compositional algorithms to solve both problems by exploiting different structures of the given system. In addition to the use of an existing abstraction (usually referred to as projection) based on removing state space variables, we develop two new abstractions for use within our compositional algorithms. One of the new abstractions is also based on state variables, while the other is based on assignments to state variables. We theoretically and experimentally show that our new compositional algorithms improve the state-of-the-art in solving both problems, upper-bounding state space topological parameters and reachability. We designed the algorithms as well as formally verified them with the aid of an interactive theorem prover. This is the first application that we are aware of, for such a theorem prover based methodology to the design of new algorithms in either AI planning or model checking.

Contents

Declaration	iii
Acknowledgement	v
Abstract	vii
1 Introduction	1
1.1 Contributions	3
1.1.1 Publications	4
1.2 Thesis Structure	5
2 Basic Concepts and Notations	7
3 Compositional Upper-Bounding of Topological Properties	9
3.1 Related Work	9
3.2 Results	11
3.3 Compositional Upper-Bounding: Negative Results	12
3.3.1 Projection and Variable Dependency	13
3.3.2 Diameter Cannot be Compositionally Upper-Bounded	15
3.3.3 Recurrence Diameter Cannot be Compositionally Bounded	17
3.3.4 Discussion	19
3.4 The Traversal Diameter	20
3.4.1 Tightness of the Traversal Diameter	21
3.4.2 Computing the Traversal Diameter	22
3.5 Using Structural Knowledge for Better Bounds: Acyclic Dependency	23
3.5.1 Upper-Bounding the Diameter using Abstractions' Recurrence Diameters	24
3.6 The Sublist Diameter	29
3.6.1 Compositionally Bounding the Sublist Diameter Under Acyclic Dependency	29
3.6.2 The Sublist Diameter as a Compositional Upper Bound on the Diameter .	30
3.7 Exploiting State Space Acyclicity	32
3.7.1 Hotel Key Protocol	33
3.7.2 State Space Acyclicity Compositional Bounding Constructs	33
3.8 A Practical Algorithm for Upper-Bounding	35
3.8.1 Hybrid Algorithm	36
3.9 Empirical Evaluations	37
3.9.1 Quality of HYB Bounds	38
3.9.2 Comparison of HYB and N_{sum}	38
3.9.3 Planning with HYB	39
3.10 Conclusion and Open Questions	39

4	Compositional Computation of Reachability in the Presence of Repetitive Symmetry	49
4.1	Related Work	50
4.2	Results	51
4.3	Planning Problems and Additional Notation	53
4.4	Computing Problem Symmetries	54
4.5	Computing the Set of Instantiations	55
4.5.1	Finding Instantiations: Practice	56
4.5.2	Finding Instantiations: Theory	57
4.6	Concrete Plan from Quotient Plan	59
4.7	Experimental Results	61
4.8	Conclusions and Future Work	63
5	Formalisation	65
5.1	Related Work	67
5.2	Factored Transition Systems in HOL4	67
5.2.1	Topological Properties	68
5.2.2	Abstraction	71
5.3	Formalising Compositional Upper-Bounding Algorithms	72
5.3.1	Compositional Bounding in the General Case	72
5.3.2	Compositional Bounds in the Presence of Acyclicity	73
5.4	Formalising the Compositional Reachability Algorithm	81
5.4.1	Formalising Soundness of Sub-solution Concatenation (Theorem 17)	82
5.4.2	Formalising Instantiation	85
5.4.3	Formalising Theorem 18 and the Validity of Goal Augmentation	87
5.5	Concluding Remarks	88
	Bibliography	91

Introduction

The state spaces of problems in fields such as model checking and Artificial Intelligence (AI) planning can be modelled as digraphs, where vertices and edges represent states and transitions, respectively. The digraph modelling the state space is represented as a *propositionally factored transition system* in languages such as STRIPS by Fikes and Nilsson (1971) and SMV by McMillan (1993). This representation has the advantage of being exponentially smaller than an explicit representation of the digraph. That compactness is necessary due to the enormous size of state spaces of realistic systems.

Two well studied problems on digraphs are upper-bounding the digraph's topological properties (e.g. the *diameter*) and testing reachability between vertices. Solving those two problems for digraphs modelling state spaces is crucial for solving problems in AI planning and model checking. However, mainstream algorithms to solve those two problems either assume the digraph is available in an explicit form (i.e. not in a factored form) or they effectively expand the factored representation and construct the digraph modelling its state space. This is impossible in many practical cases, where factored representations of state spaces with 2^{1000} states are not uncommon. That issue is referred to as the state space explosion problem (Clarke et al. (2001)). Furthermore, assuming that constructing the digraph was possible, worst-case run-times of state-of-the-art algorithms to solve those problems are worse than linear in the size of the digraph, which is infeasible.

Different techniques were invented to mitigate the state space explosion. Computing reachability is one notable problem for which many techniques were developed to circumvent the state space explosion. For example, different modifications to path search algorithms were developed. Those techniques can exponentially reduce the path search run-time in many practical examples. Although those methods were developed multiple decades ago, they are still active research topics. Examples of such methods are symmetry breaking (Emerson and Sistla, 1996; Clarke et al., 1996, 1998; Fox and Long, 1999, 2002), partial order reduction (Godefroid, 1990; Gerth et al., 1995; Alur et al., 1997; Chen and Yao, 2009; Xu et al., 2011), and heuristic search (McDermott, 1996; Blum and Furst, 1997; Bonet and Geffner, 2001; Hoffmann and Nebel, 2001). A central drawback with many of these techniques is that they are tailored to specific solution algorithms. For instance, many heuristic search techniques are designed for explicit state space search, but it is not obvious how to combine them with SAT-based planning methods to get a significant performance boost.

Another more general approach to mitigate state space explosion is the *compositional* approach: a solution to the original problem instance is found or approximated via composing solutions to one or many (possibly significantly) smaller derived sub-problems, or “abstractions” (e.g. Clarke et al. (1994); Knoblock (1994); Williams and Nayak (1997); Berezin et al. (1998); Amir and Engelhardt (2003); Kroening (2006); Case et al. (2009); Filiot et al. (2011)). An advantage of this approach is that it is not specific to a certain solution method, i.e. it does not matter how the abstractions' solutions were computed. In this thesis we study the application of this approach to

solve diameter upper bounding and reachability. In particular, we consider computing the topological property (resp. solving the reachability problem) for a set of *minors* of the digraph modelling the state space of the system under consideration. Minors are digraphs derived by vertex contractions as well as edge and vertex deletions. These minors of the state space model the state spaces of abstractions of the given system. Minors’ topological properties (resp. reachability witness paths) are then composed into an upper bound on the topological property (resp. a witness reachability path) for the digraph modelling the given system’s state space.

In many practical cases the compositional approach is one of the few known feasible approaches to solve problems on the state space of the given factored system. This is because it avoids constructing and performing computations on the large digraph modelling the state space, and only constructs and processes (significantly smaller) minors of it. On the other hand, a drawback of the compositional approach is that it is not always the case that it is usefully applicable to the given transition system. Nonetheless, we identify structures whose presence in the transition system make it possible to compositionally solve the problems. These structures are acyclicity (in “state variable dependencies” and the state space) for compositionally upper-bounding the diameter, and symmetry (between state variables) for compositionally solving the reachability problem. We show that those structures are abundant in systems coming from standard AI planning benchmarks and standard model checking problems. We thus show that our algorithms improve the state-of-the-art in both scenarios of AI planning: proving plan existence and plan non-existence.

In addition to developing new compositional algorithms, we formally verify their validity in the interactive theorem prover HOL4 by Slind and Norrish (2008). To do that we develop a rich formal proof library about transition systems. Many theorems in our formalisation apply to infinite state transition systems and accordingly this library can be used for application well beyond the algorithms that we verify in this work, like algorithms for hybrid systems, for instance. A problem that we faced while developing these compositional algorithms is the sheer difficulty of manually verifying their properties and soundness. Using the proof assistant HOL4, we discovered many mistakes in our work at different stages and even mistakes in the literature, in a process we would describe as “computer-aided algorithm design”.

In addition to aiding us develop our algorithms and giving us extra assurance about our results, we believe that the formalisation of AI algorithms in general is of great utility, if those algorithms are to be deployed in safety critical applications or in autonomous exploration of space. There is a twofold justification for requiring formal correctness guarantees. First, if the output of the algorithm is hard to test—e.g. that a plan does not exist—then we require correctness to assure the sound operation of the system at hand. For example, if a chemical plant can be rendered safe in the event of a potentially dangerous subsystem failure, the planner must be able to identify that course of action which averts a catastrophic failure. Secondly, if the output is computationally easy to test—e.g. sound operation is easily guaranteed—it remains a potentially dangerous waste of time and resources to have a computationally expensive planning subsystem that produces unusable plans. This is especially the case when such resources are scarce and one cannot afford to invest in an algorithm without guarantees, like with autonomous exploration of space.

Propositionally Factored Transition Systems In this representation, every state (i.e. vertex) is defined as an assignment to a set of Boolean state variables. For example, consider a hotel that has a number of rooms, each of which opens with one of a set of keys at a given time. In this case, the set of state variables assigned indicating whether a certain room opens with a certain key at the state.

In the factored representation multiple transitions (i.e. edges) are factored into actions that identify which states are connected to which states. For instance, consider an action that, given

the variable indicating that a certain room opens with a key i is true (its precondition), negates that variable and sets another one indicating that the room currently opens with a different key j (its effect). This action factors all the transitions that go from states where the room opens with key i to states where it opens with key j . In addition to the compactness of this representation, without it, we would not be able to pursue the proofs of our results or define concepts like the *sublist diameter* (see Section 3.6), as naturally.

Minors and Abstractions Compositional approaches involve solving given problems for state spaces of “abstractions” of the given system, which are minors of the digraph modelling the given system’s state space. Abstractions that we consider are *projection* over a set of state variables (see Section 3.3.1), *snapshotting* over an assignment of a set of state variables (see Section 3.6.2), and *quotienting* (see Section 4.5) over a partition of state variables. In projection, state variables are removed from the factored system. This is equivalent to repetitive vertex contraction operations in the state space. In snapshotting, actions that violate the given assignment (i.e. if it holds, they either cannot execute or change it in their effect) are removed from the factored system. This is equivalent to consecutive edge and vertex deletions in the state space. In quotienting, given a partition of the state variables, all members of an equivalence class are removed, except for one representative from each equivalence class. This is also equivalent to consecutive edge and vertex deletions in the state space.

1.1 Contributions

We start by discussing our results on upper-bounding topological properties of digraphs modelling state spaces. Our first two results are negative results concerning compositional bounding of two prominent topological properties: the diameter (the longest distance between any two states) and the recurrence diameter (the length of the longest simple path). We show that any method that compositionally bounds the (recurrence) diameter using only “state variable dependency” analysis is unsound. We do so by demonstrating that the (recurrence) diameters of minors modelling projections’ state spaces cannot be composed, solely based on inter-projection variable dependencies, to upper-bound the (recurrence) diameter of the state space of the system. In other words, a compositional algorithm needs more “information” than just variable dependencies to compositionally upper-bound the (recurrence) diameter.

We define two new topological properties that can be compositionally bounded, which are themselves upper bounds on the diameter and/or the recurrence diameter. Thus, they can be used to compositionally upper-bound the diameter and/or the recurrence diameter. The first property, the *traversal diameter*, can always be compositionally bounded, regardless of the structure of the transition system. However, it can be exponentially looser than the diameter and the recurrence diameter. The second property, the *sublist diameter*, can be compositionally bounded in systems that can be partitioned into sub-systems with acyclic dependency relations among them. It is an upper bound on the diameter and a lower bound on the recurrence diameter. We prove that the compositional bounds that we provide for both properties cannot be improved.

In deriving the above results, we also provide theoretical constructs that we believe are interesting themselves. The first is a novel proof artifact, the *stitching function* (\mathfrak{k}), which can be used to merge paths from different projections of a transition system into a path in the concrete system. We also provide digraph constructions that we refer to as “flowers” and “inverted flowers”. The main feature of the flower constructions is that they are digraphs that can have arbitrarily many vertices added to them, without changing their diameters, sublist diameters, or longest paths.

We provide a new algorithm to compositionally upper-bound the diameter by using the projections' recurrence diameters. We prove that bounds computed by this algorithm cannot be improved. We also experimentally show that this algorithm significantly outperforms previous algorithms in terms of tightness of computed bounds. We combine this algorithm with a novel approach to exploiting acyclicity in the state space to form a hybrid bounding procedure. We show that this algorithm enables the decomposition of a given system into very small abstractions compared to state-of-the-art algorithms. Also experimentally, compared to existing practical approaches, ours can yield exponentially tighter bounds, which we use to significantly improve the coverage and performance of a propositional satisfiability based planner.

The other problem which we seek to solve compositionally is reachability between states. We provide a novel procedure to exploit symmetries for efficient computation of reachability between states in transition systems. This procedure differs from existing approaches in that we use symmetries to obtain an abstraction of the concrete problem description, which we call the *descriptive quotient*. The reachability computation is then done on the abstraction, and if a path is found in the descriptive quotient, it is used to synthesise a path in the concrete system. We show experimentally that this approach significantly outperforms other approaches to exploit symmetries for finding paths between states in many standard AI planning benchmarks. A drawback with this method, nonetheless, is that it is incomplete.

Our last contribution is that we provide a library of formal proofs for transition systems in the interactive theorem prover HOL4 and we use it to formally verify the algorithms that we developed. Although there are many formal proof libraries on the subject of transition systems, the library that we developed is the first we are aware of that is dedicated for systems in the factored representation.

Availability All our HOL scripts, experimental code and data are available from <https://MohammadAbdulaziz@bitbucket.org/MohammadAbdulaziz/planning.git>.

1.1.1 Publications

During the course of my PhD I wrote as a main contributor, with other contributors, the papers and manuscripts that are listed below:

Mohammad Abdulaziz and Lawrence C Paulson. An Isabelle/HOL formalisation of Green's theorem. In *International Conference on Interactive Theorem Proving*, pages 3–19. Springer, 2016.

Mohammad Abdulaziz, Charles Gretton, and Michael Norrish. Mechanising Theoretical Upper Bounds in Planning. In *Workshop on Knowledge Engineering for Planning and Scheduling*, 2014.

Mohammad Abdulaziz, Charles Gretton, and Michael Norrish. Verified Over-Approximation of the Diameter of Propositionally Factored Transition Systems. In *Interactive Theorem Proving*, pages 1–16. Springer, 2015a.

Mohammad Abdulaziz, Michael Norrish, and Charles Gretton. Exploiting symmetries by planning for a descriptive quotient. In *Proc. of the 24th International Joint Conference on Artificial Intelligence, IJCAI*, pages 25–31, 2015b.

Mohammad Abdulaziz, Charles Gretton, and Michael Norrish. A State Space Acyclicity Property for Exponentially Tighter Plan Length Bounds. In *International Conference on Automated Planning and Scheduling (ICAPS)*. In Press, 2017a.

Mohammad Abdulaziz, Michael Norrish, and Charles Gretton. Formally Verified Algorithms for Upper Bounding State Space Diameters. In *To appear in the Journal of Automated Reasoning*, 2017b.

1.2 Thesis Structure

This thesis is organised as follows. In Chapter 2 we discuss the basic concepts related to factored transition systems and the action execution semantics in our framework. In Chapter 3 we describe previous work in the literature related to computing upper bounds on topological properties of directed graphs in general as well as directed graphs that model state spaces. We then discuss more formal definitions that are more specific to compositional upper-bounding. We then describe our different compositional algorithms and results to compute upper bounds, and then we end this chapter by concluding remarks regarding our approach for compositional upper-bounding as well as some open questions that we believe are interesting. In Chapter 4 we discuss different approaches in the literature to compute reachability in general directed graphs as well as state spaces, and focus more on the approaches that try to exploit symmetry to enhance the efficiency of the reachability computation. Then we discuss our approach for compositional computation of reachability, where we formally define concepts that are relevant to our algorithm, and we focus on computing reachability within the framework of automated planning. Then we finalise that chapter with remarks summarising the merits of our compositional approach to compute reachability and possible ways of extending it. Lastly, in Chapter 5, we discuss previous work in the literature on formalisation and we focus on results related to formalising transition systems. We then discuss our formalisation of factored transition systems and our compositional algorithms in the interactive theorem prover HOL4, and then finish that chapter by some concluding remarks on our experience in formalising our algorithms.

Basic Concepts and Notations

Where we do not refer to a specific function, we shall use the symbol f . We formalise functions as sets of key-value pairs ($k \mapsto v$). We use the term *domain* mathematically. We write $\mathcal{D}(f)$ for the domain of f , i.e., $\{k \mid (k \mapsto v) \in f\}$. We write $\mathcal{R}(f)$ for the range of f , i.e., $\{v \mid (k \mapsto v) \in f\}$.

Definition 1 (Digraph). *A digraph whose vertices are labelled by values of type α (α -graph) \mathcal{G}_α is represented by the tuple $\langle V, E, f \rangle$. V is the set of vertices, which we denote as $V(\mathcal{G}_\alpha)$. E is a set of edges, that are distinct ordered pairs of vertices from $V(\mathcal{G}_\alpha)$, which we denote as $E(\mathcal{G}_\alpha)$. $f : V(\mathcal{G}_\alpha) \Rightarrow \alpha$ is a labelling function, where a vertex $u_1 \in V(\mathcal{G}_\alpha)$ has the label $f(u_1)$ of type α , which we denote as $\mathcal{G}_\alpha(u_1)$. We note that when we illustrate a digraph in a figure, a vertex u_1 is docketed with the label $f(u_1)$ instead of the vertex name u_1 . We also omit self loops. For $u_1 \in V(\mathcal{G}_\alpha)$, the set of children of u_1 is $\text{children}_{\mathcal{G}_\alpha}(u_1) = \{u_2 \mid (u_1, u_2) \in E(\mathcal{G}_\alpha)\}$. If f is injective, we use $l \in \mathcal{G}_\alpha$ to denote that l is a label of some $u_1 \in V(\mathcal{G}_\alpha)$, i.e. $l = f(u_1)$, and we use $\text{children}_{\mathcal{G}_\alpha}(l)$ to refer to the labels of the children of u_1 , i.e. $\{\mathcal{G}_\alpha(u_2) \mid u_2 \in \text{children}_{\mathcal{G}_\alpha}(u_1)\}$. For $\mathcal{G}_\alpha \equiv \langle V, E, f \rangle$, the image of a function $g : \alpha \Rightarrow \beta$ on \mathcal{G}_α is the β -graph $g(\mathcal{G}_\alpha) \equiv \langle V, E, g \circ f \rangle$, where $g \circ f$ is the composition of f and g . Effectively, the image operation on a digraph is a relabelling operation.*

We now define factored representations of digraphs, which are mainly used to compactly represent digraphs that model state spaces. In these representations a digraph is propositionally factored, where vertices correspond to “states” and sets of edges are compactly described in terms of “actions”. In this formalism edges and paths between vertices correspond to executing actions and action sequences, respectively, between states.¹

Definition 2 (States and Actions). *A state, x , is a finite map from variables—i.e., state-characterising propositions—to Booleans, i.e. a set of mappings $v \mapsto b$. We write $\mathcal{D}(x)$ to denote $\{v \mid v \mapsto b \in x\}$, the domain of x . For states x_1 and x_2 , the union, $x_1 \uplus x_2$, is defined as $\{v \mapsto b \mid \text{if } (v \in \mathcal{D}(x_1)) \text{ then } b = x_1(v) \text{ else } b = x_2(v)\}$. Note that the state x_1 takes precedence. An action is a pair of finite maps, (p, e) , where p represents the preconditions and e represents the effects. For action $\pi = (p, e)$, $\mathcal{D}(\pi) \equiv \mathcal{D}(p) \cup \mathcal{D}(e)$.*

Definition 3 (Execution). *When an action $\pi (= (p, e))$ is executed at state x , it produces a successor state $\text{ex}(x, \pi)$, formally defined as $\text{ex}(x, \pi) = \text{if } p \not\subseteq x \text{ then } x \text{ else } e \uplus x$. We lift ex to lists of actions $\vec{\pi}$ as the second argument. So $\text{ex}(x, \vec{\pi})$ denotes the state resulting from successively applying each action from $\vec{\pi}$ in turn, starting at x , which corresponds to a path in the state space.*

¹This representation is equivalent to representations commonly used in the model checking and AI planning communities (e.g. STRIPS by Fikes and Nilsson (1971) and SMV by McMillan (1993)). Whereas conventional expositions in the planning and model-checking literature would also define initial conditions and goal/safety criteria, here we omit those features from discussion. Our bounds and the different topological properties we consider are independent of those features.




Figure 2.1: The state space of δ in Example 1.

We give examples of states and actions using sets of literals. For example, $\{v_1, \overline{v_2}\}$ is a state where state variables v_1 is (maps to) true, and v_2 is false and its domain is $\{v_1, v_2\}$. $(\{v_1, \overline{v_2}\}, \{v_3\})$ is an action that if executed in a state where v_1 and $\overline{v_2}$ hold, it sets v_3 to true. $\mathcal{D}((\{v_1, \overline{v_2}\}, \{v_3\})) = \{v_1, v_2, v_3\}$.

Definition 4 (Factored System). For a set of actions δ we write $\mathcal{D}(\delta)$ for the domain of δ , which is the union of the domains of all the actions in δ . The set of valid states, written $\mathbb{U}(\delta)$, is $\{x \mid \mathcal{D}(x) = \mathcal{D}(\delta)\}$. The set of valid action sequences is the Kleene closure of δ , i.e $\delta^* = \{\vec{\pi} \mid \mathbf{set}(\vec{\pi}) \subseteq \delta\}$, where $\mathbf{set}(l)$ is the set of members in list l .

δ is the factored representation of the digraph $\mathcal{G}(\delta) \equiv \langle V, E, f \rangle$, where $V = \mathbb{U}(\delta)$, $E \equiv \{(x, \mathbf{ex}(x, \pi)) \mid x \in \mathbb{U}(\delta), \pi \in \delta\}$, and f is the identity function. We refer to $\mathcal{G}(\delta)$ as the state space of δ . For states x and x' , $x \rightsquigarrow x'$ denotes that there is a $\vec{\pi} \in \delta^*$ such that $\mathbf{ex}(x, \vec{\pi}) = x'$. Let the connected component for a state x be $\mathcal{S}(\delta, x) = \{y \mid x \rightsquigarrow y \vee y \rightsquigarrow x\}$.

Example 1. Consider the following factored representation, δ , and the digraph in Figure 2.1 that represents its state space.

$$\left\{ \begin{array}{l} p_1 = (\{\overline{v_1}, \overline{v_2}, v_3\}, \{v_1\}), p_2 = (\{v_1, \overline{v_2}, v_3\}, \{\overline{v_1}, v_2\}), p_3 = (\{\overline{v_1}, v_2, v_3\}, \{v_1\}), \\ k_1 = (\{\overline{v_3}\}, \{v_1, v_2\}), k_2 = (\{\overline{v_3}\}, \{\overline{v_1}, v_2\}), k_3 = (\{\overline{v_3}\}, \{v_1, \overline{v_2}\}), k_4 = (\{\overline{v_3}\}, \{\overline{v_1}, \overline{v_2}\}) \end{array} \right\}.$$

In the figure different states defined on the variables $\mathcal{D}(\delta) = \{v_1, v_2, v_3\}$ are shown. Interpreting δ as a transition system, it has two “modes” of operation, where the variable v_3 in the preconditions of actions in δ determines the mode of operation. Each of those modes represents one connected component of the digraph in Figure 2.1, where actions k_1, k_2, k_3, k_4 (which execute successfully if v_3 maps to false) represent edges in the clique component. For instance, action k_1 represents all incoming edges to vertex $\{v_1, v_2, \overline{v_3}\}$. On the other hand, each one of the actions p_1, p_2, p_3 (which execute successfully if v_3 maps to true) represents an edge in the simple path component.

Compositional Upper-Bounding of Topological Properties

Two fundamental properties of digraphs are the *diameter* and the length of the longest simple path (the latter also known as the *recurrence diameter* in the verification community). Knowing an upper bound on the diameter or the recurrence diameter of a state space comes up as a necessary ingredient for many practical applications and algorithms in model checking and AI planning. A prominent example is the bounded model-checking algorithm by Biere et al. (1999), whose completeness depends on having an upper bound on the diameter (for checking safety properties or plan existence) or the recurrence diameter (for checking liveness properties). Because of this practical importance a lot of effort has been spent trying to efficiently compute or approximate both properties.

However, and as we pointed out earlier to compute either property, one needs to expand the factored representation δ and construct the digraph modelling its state space, $\mathcal{G}(\delta)$, which is infeasible. This is further compounded by the fact that state-of-the-art algorithms to compute or reasonably approximate the diameter and the longest path, have worse than quadratic (e.g. Yuster (2010); Aingworth et al. (1999)) and exponential (e.g. Kroening and Strichman (2003)) run-times, respectively, in the size of the digraph. In this chapter we study applying the compositional approach to compute upper bounds on both properties for digraphs that model state spaces. In particular, we consider computing the topological property for a set of minors of the state space. Minors' properties are then composed into an upper bound on the property of the original digraph.

State Variable Dependency Informally, a state variable v_1 depends on v_2 iff the assignment of v_2 at some state can possibly affect assignment of v_1 in a current or a future state. This relation lifts to sets of state variables, where a set vs_2 depends on vs_1 iff vs_2 has a variable that depends on some variable in vs_1 . Dependency based analysis of model-checking and planning problems has been used for a long time (e.g. Rintanen and Gretton (2013); Baumgartner et al. (2002); Williams and Nayak (1997); Helmert (2006a); Kroening (2006)). A practically important class of projections are ones done on partitions of the state variables whose members are closed under mutual variable dependency, i.e. variable dependencies between different equivalence classes are acyclic. Those partitions capture the abundantly present acyclic dependencies between different modules of real-world systems, such as different circuit components.

3.1 Related Work

Computing the diameter exactly can be done via solving the All Pairs Shortest Path (APSP) problem for the (di)graph at hand. APSP cannot be solved in better than quadratic time (in the number

of vertices of the (di)graph), and existing exact solutions have a run-time close to cubic (e.g. Fredman (1976); Alon et al. (1997); Chan (2010); Yuster (2010)). Furthermore, there is strong evidence that the diameter cannot be computed in time better than quadratic (Roditty and Vassilevska Williams, 2013). This run-time can be very limiting in digraphs arising in practical applications due to their size. Accordingly, in the algorithms community a lot of work has been done on developing methods to approximate the diameter. In a seminal paper, Aingworth et al. (1999) devised an algorithm that computes a $\frac{3}{2}$ -approximation of the diameter for digraphs in $O(m\sqrt{n \log n} + n^2 \log n)$ time, where n is the number of vertices and m is the number of edges. Examples of other work studying approximation algorithms for digraph diameters include Roditty and Vassilevska Williams (2013); Chechik et al. (2014); Abboud et al. (2016).

On the other hand, computing the recurrence diameter is an NP-hard problem (Pardalos and Migdalas, 2004), with only known exact solutions in exponential time. Accordingly it is much harder to compute than the diameter, especially, for practical digraphs. Furthermore, the hardness of computing the recurrence diameter in digraphs was reaffirmed by Björklund et al. (2004), where they showed that, in the general case, it is impossible (under plausible assumptions) to get a polynomial approximation of the length of the longest path in polynomial time. Existing polynomial time approximation algorithms for the longest path, like Alon et al. (1995); Björklund and Husfeldt (2003), can only find paths of lengths logarithmic in the length of the longest path. Also, to the best of our knowledge, all approximation techniques studied in the algorithms community are lower bounding techniques.

In the graph theory and the combinatorics communities, diameters of undirected graphs have been extensively studied since 1965, where work like Moon et al. (1965); Erdős et al. (1989); Knyazev (1987) computed upper bounds for different classes of undirected graphs. However, for digraphs, work on upper-bounding diameters started more recently. In 1992 Soares (1992) provided a tight upper bound on the diameter of biregular digraphs. Then starting in 2000, Peter Dankelmann et al. treated different structures of digraphs (e.g. Dankelmann (2005); Dankelmann and Volkman (2010); Dankelmann and Dorfling (2016)).

The completeness of bounded model-checking (Biere et al., 1999, 2003) and satisfiability (SAT) based planning (Kautz and Selman, 1992) depends on having an upper bound on the diameter or the recurrence diameter (depending on the verification problem at hand), and the tighter that bound the more quickly the algorithm will likely terminate. Because of that, studying the diameter and the recurrence diameter is an active research topic in the verification community (see for example Biere et al. (1999, 2003); Kroening and Strichman (2003); Sheeran et al. (2000); Kroening et al. (2011); Bundala et al. (2012); Clarke et al. (2009)). The most dominant approach to compute the diameter or the recurrence diameter in verification applications is via encodings in SAT (see Biere et al. (1999); Kroening and Strichman (2003)). Most notably, Biere et al. (1999) conjecture that for the question of “whether for a certain digraph, a number N is its diameter”, there is not a SAT encoding of size polynomial in N . However, if the question is “whether N is the recurrence diameter”, they provide a SAT encoding of size $O(N^2)$, which was improved by Kroening and Strichman (2003) to $O(N \log N)$.

Applying the compositional approach for diameter upper-bounding was pioneered by Baumgartner et al. (2002), in the context of bounded model-checking. They showed that the diameter of the system’s state space is upper-bounded by a polynomial in the recurrence diameters of the system’s projections on sets of state variables closed under mutual dependency (i.e. a set of contractions of the state space). Also, more recently, Rintanen and Gretton applied a similar approach to upper-bound the diameters for AI planning Rintanen and Gretton (2013) problems, but using the number of the contractions’ states instead of the recurrence diameter.

3.2 Results

Our contributions are both theoretical and practical. We first discuss the theoretical contributions. First let \mathcal{G}_i denote a digraph and vs_i denote a set of state variables. Then let $\mathcal{G}_{1..n}$ denote the set of digraphs $\{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$ and $vs_{1..n}$ denote the sets of variables $\{vs_1, vs_2, \dots, vs_n\}$. For some function b , let $b_{1..n}$ denote $\{b(\mathcal{G}_1), b(\mathcal{G}_2), \dots, b(\mathcal{G}_n)\}$. For some partition $vs_{1..n}$, let \mathcal{G}_{VS} denote the dependencies between different members of $vs_{1..n}$.

Our first two results are negative. We show that recurrence diameters of projections' state spaces cannot be composed, solely based on dependencies between different sets of variables in $vs_{1..n}$, to upper-bound the recurrence diameter of the state space of the system (Theorem 2). Also, if $vs_{1..n}$ satisfies the additional requirement of having two members with a dependency between them, then the same result applies to the diameter (Theorem 1).

To circumvent these negative results we define two new topological properties that can be compositionally bounded, which are themselves upper bounds on the diameter and the recurrence diameter. Thus they can be used to compositionally upper-bound the diameter and the recurrence diameter. We consider projections induced by two types of partitions of the set of state variables of δ : arbitrary partitions and partitions whose equivalence classes are closed under mutual variable dependency.

The first new topological property, is the *traversal diameter* (td), which is one less than the largest number of vertices that can be traversed by any path in a digraph. td can be seen as a measure of branching in $\mathcal{G}(\delta)$. It is an upper bound on (and can be exponentially larger than) the recurrence diameter, and a lower bound on (and can be exponentially smaller than) the number of states in $\mathcal{G}(\delta)$ (Theorem 5). Its main advantage is that it can be compositionally bounded using projections obtained from an arbitrary partition $vs_{1..n}$. It is upper-bounded by $\prod_{1 \leq i \leq n} (td_i + 1) - 1$, if the contractions $\mathcal{G}_{1..n}$ of $\mathcal{G}(\delta)$ are state spaces of projections of δ on members of $vs_{1..n}$ (Theorem 3). Thus it provides a very flexible tool to compositionally upper-bound both the diameter and recurrence diameter. We also prove a strong statement of tightness for that bound: any function taking as input $td_{1..n}$ and \mathcal{G}_{VS} whose output can be less than $\prod_{1 \leq i \leq n} (td_i + 1) - 1$, is not a sound bound on td , for any partition. We also show that td can be computed in time that is linear in the size of the digraph (Theorem 6).

Although compositional upper-bounding using the traversal diameter has the advantage of being applicable to any partition, it can be a rather loose upper bound on the diameter. To alleviate that, we define the *sublist diameter* (ℓ), an upper bound on the diameter and a lower bound on the recurrence diameter. A main feature of it is that it can be exponentially smaller than the recurrence diameter, but it can also be exponentially larger than the diameter (Theorem 10).

We introduce a function N_{sum} that takes $\ell_{1..n}$ and \mathcal{G}_{VS} . We show that it computes an upper bound on the sublist diameter that is a polynomial in $\ell_{1..n}$, if members of $vs_{1..n}$ are closed under mutual variable dependency (Lemma 3). This is significantly better than using $td_{1..n}$ to upper-bound the diameter since ℓ can be exponentially smaller than td and because the polynomials produced by N_{sum} do not necessarily have all possible product terms, as $\prod_{1 \leq i \leq n} (td_i + 1) - 1$ does. Compared to similar earlier approaches in Baumgartner et al. (2002); Rintanen and Gretton (2013), this is a substantial improvement in two ways. Firstly, we prove that any function taking the same input as N_{sum} , and whose output can be less than computed by N_{sum} is not a sound upper bound on the diameter, i.e. N_{sum} is a lower bound on all the previous approaches (Theorem 9). Furthermore, we experimentally verify that N_{sum} computes significantly tighter bounds than previous approaches. Secondly, we use the contractions' sublist diameters which can be exponentially tighter than functions used previously (namely, the recurrence diameter and the state space size). However, we show that computing the sublist diameter is NP-hard (Theorem 12). The relation

between the different topological properties and bounding functions can be found in Figure 3.1.

Our proof of Lemma 3 is the most technical. It is constructive, and it depends on the fact that a sequence of transitions in δ (i.e. edges in $\mathcal{G}(\delta)$) is equivalent to different sequences of transitions in the different projections (i.e. edges in $\mathcal{G}_{1..n}$) that are interleaved. We show that for each transition sequence from a projection, removing redundant transitions makes it shorter than the sublist diameter of the projection. We then recombine all shortened transition sequences with a novel proof artifact, *the stitching function* (\mathbb{X}), and show that the recombined sequence both begins and ends in the same vertices as the original one, and is shorter than the bound produced by N_{sum} . Although they had different purposes from ours, in some sense this is a generalisation of the theory developed in Knoblock (1994) and in Brafman and Domshlak (2003).

Another interesting proof is that of Theorem 10. It relies on the fact that, unlike the diameter or the recurrence diameter, the sublist diameter depends on the factored representation δ rather than the underlying state space $\mathcal{G}(\delta)$. We show that it can be minimised through factoring more edges into fewer actions.

Another point to highlight is in our tightness proofs. They depend on digraph constructions that we refer to as “flowers” and “inverted flowers” (see Figure 3.6a and Figure 3.10 for an example). The main feature of those constructions is that they are digraphs that can have arbitrarily many vertices added to them, without changing their diameters, sublist diameters, or longest paths.

We combine exploitation of acyclicity in variable dependency constituted by Theorem 8, with a novel approach to exploiting acyclicity in the state space $\mathcal{G}(\delta)$, to form a hybrid bounding procedure. This algorithm enables the decomposition of a given system to abstract sub-systems that are much smaller than what is attainable using state-of-the-art algorithms, which all depend only on acyclicity in variable dependency. This provides the potential for exponentially easier computation of upper bounds. Also experimentally, compared to existing practical approaches, ours can yield exponentially tighter bounds.

This procedure is based on first finding a partition of the states in $\mathcal{G}(\delta)$ that induces an acyclic “quotient” contraction of $\mathcal{G}(\delta)$. Then, we show that the diameter of $\mathcal{G}(\delta)$ is upper-bounded by the weightiest path in the acyclic quotient of $\mathcal{G}(\delta)$. The weight of a path is the number of edges in it added to the cost of every vertex it traverses. The cost of visiting a vertex in that quotient digraph is the diameter of the subgraph of $\mathcal{G}(\delta)$ induced by the states in that quotient vertex. However, to avoid computing the acyclic quotient by constructing $\mathcal{G}(\delta)$ and the possible state space explosion, we get it through projecting δ on a subset of its state variables. That way, the acyclic quotient is the state space of the projection of δ . Also, to avoid the possible state space explosion in computing the subgraphs induced by different quotient vertices explicitly, we compute them in the factored form by snapshotting δ on different states of the projection.

We experimentally show that this algorithm significantly outperforms—both in the tightness of the bounds obtained and in the quality of decomposition—existing approaches. Using the upper bounds computed by that algorithm as horizons for a SAT-based planner, we prove the unsolvability of problems that cannot be proven using the state-of-the-art state space search planners and model-checkers. One notable example is the problem of model-checking the safety of the hotel key protocol. We also significantly improve the coverage of the SAT-based planner Madagascar MP by using upper bounds rather than ramp-up strategies Rintanen (2012).

3.3 Compositional Upper-Bounding: Negative Results

Topological properties that we consider upper-bounding are the diameter and the recurrence diameter. The diameter is the length of the longest shortest path between any two valid states. This can be formally defined as follows for a propositionally factored system.

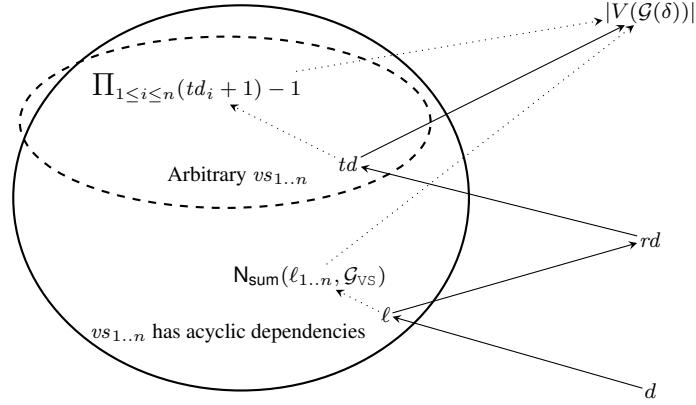


Figure 3.1: A diagram that represents the upper-bounding relations between different topological properties. An arrow from f_1 to f_2 denotes that we showed $f_1(\delta) \leq f_2(\delta)$, for any δ , and f_2 is a “tight” bound on f_1 . A solid arrow denotes that we showed an exponential separation between f_1 and f_2 . Functions in the dotted ellipse can be compositionally bounded by projections induced by arbitrary partitions of state variables, while those in the solid ellipse need the partitions’ members to be closed under mutual dependency.

Definition 5 (Diameter). Let $|l|$ be the length of a list l . The diameter is defined as follows.

$$d(\delta) = \max_{x \in \mathbb{U}(\delta)} \max_{\vec{\pi} \in \delta^*} \min\{|\vec{\pi}'| \mid \vec{\pi}' \in \delta^* \wedge \text{ex}(x, \vec{\pi}) = \text{ex}(x, \vec{\pi}')\}$$

We would like to note that instead of quantifying over pairs of states, we quantify over states and paths going outside of those states. This is to avoid the complexities arising from the situation when two states are not connected.

The recurrence diameter is the length of the longest simple path in the digraph modelling the state space. This can be formally defined as follows.

Definition 6 (Recurrence Diameter). Let $\text{distinct}(x, \vec{\pi})$ denote that all states traversed by executing $\vec{\pi}$ at x are distinct states.

$$rd(\delta) = \max_{x \in \mathbb{U}(\delta)} \max_{\vec{\pi} \in \delta^*} \max\{|\vec{\pi}| \mid \text{distinct}(x, \vec{\pi})\}$$

If there is a valid action sequence between any two states, then there is a valid action sequence between them that is no longer than the diameter. Also, it should be clear that the recurrence diameter is an upper bound on the diameter.

3.3.1 Projection and Variable Dependency

A key concept for compositional reasoning about factored representations of digraphs is *projection*. Projection of a transition system is equivalent to a sequence of vertex contractions in the digraph modelling the state space.

Definition 7 (Projection). Projecting an object (a state x , an action π , a sequence of actions $\vec{\pi}$ or a factored representation δ) on a set of variables vs restricts the domain of the object or the components of composite objects to vs . Projection is denoted as $x|_{vs}$, $\pi|_{vs}$, $\vec{\pi}|_{vs}$ and $\delta|_{vs}$ for a state, action, action sequence and factored representation, respectively. In the case of action sequences, an action with no effects after projection is dropped entirely.




Figure 3.2: The contraction of the state space equivalent to the projection of δ on $\{v_1, v_2\}$.




Figure 3.3: (a) the dependency graph of δ from Example 1, and (b) a lifted dependency graph of δ . Actions induce edges in (a): for example, v_1 and v_2 co-occur in action effects, while v_3 only happens to be in the precondition.

Example 2. Letting $vs = \{v_1, v_2\}$, below is the projection $\delta|_{vs}$, for δ from Example 1. Figure 3.2 shows $\mathcal{G}(\delta|_{vs})$.

$$\left\{ \begin{array}{l} p_1|_{vs} = (\{\overline{v_1}, \overline{v_2}\}, \{v_1\}), p_2|_{vs} = (\{v_1, \overline{v_2}\}, \{\overline{v_1}, v_2\}), p_3|_{vs} = (\{\overline{v_1}, v_2\}, \{v_1\}), \\ k_1|_{vs} = (\emptyset, \{v_1, v_2\}), k_2|_{vs} = (\emptyset, \{\overline{v_1}, v_2\}), k_3|_{vs} = (\emptyset, \{v_1, \overline{v_2}\}), k_4|_{vs} = (\emptyset, \{\overline{v_1}, \overline{v_2}\}) \end{array} \right\}$$

State variable dependency analysis captures many structures present in factored systems. It is used to obtain practically very useful projections, e.g. projections on sets of variables that are closed under mutual dependency. Also in many cases upper-bounding algorithms rely on the state variable dependencies which are described by the *dependency graph*.

Definition 8 (Dependency). A variable v_2 is dependent on v_1 in δ (written $v_1 \rightarrow v_2$) iff one of the following statements holds: ¹ (i) v_1 is the same as v_2 , (ii) there is $(p, e) \in \delta$ such that $v_1 \in \mathcal{D}(p)$ and $v_2 \in \mathcal{D}(e)$, or (iii) there is a $(p, e) \in \delta$ such that both v_1 and v_2 are in $\mathcal{D}(e)$. A set of variables vs_2 is dependent on vs_1 in δ (written $vs_1 \rightarrow vs_2$) iff: (i) vs_1 and vs_2 are disjoint, and (ii) there are $v_1 \in vs_1$ and $v_2 \in vs_2$, where $v_1 \rightarrow v_2$.

Definition 9 (Dependency Graph). This graph, sometimes called the causal graph, was described independently by Knoblock Knoblock (1994) and then Williams and Nayak Williams and Nayak (1997). $\mathcal{G}_{\mathcal{D}(\delta)}$ is a dependency graph of δ iff (i) its vertices are bijectively labelled by variables from $\mathcal{D}(\delta)$, and (ii) it has an edge from vertex u_1 to u_2 iff $v_1 \rightarrow v_2$, where v_1 and v_2 are the labels of u_1 and u_2 , respectively.

Definition 10 (Lifted Dependency Graph). \mathcal{G}_{vs} is a lifted dependency graph of δ iff (i) its vertices are bijectively labelled by members of a partition of $\mathcal{D}(\delta)$, and (ii) it has an edge from vertex u_1 to u_2 (labelled by vs_1 and vs_2 , respectively) iff $vs_1 \rightarrow vs_2$.

Example 3. Figure 3.3a and Figure 3.3b show the dependency graph and a lifted dependency graph of δ from Example 1.

As stated earlier, many compositional algorithms compute reachability in succinct digraphs successfully exploit structures in the variable dependencies of a system via computing reachability in projections. Analogously, to compositionally bound the diameter (resp. the recurrence

¹Our definition is equivalent to those in Williams and Nayak (1997); Knoblock (1994) in the context of AI planning.

diameter) one can imagine a function f that takes the projections' diameters and the dependencies between those projections and returns an upper bound on the diameter of the entire system. As arguments to f , projections' diameters and their dependencies are encoded as a \mathbb{N} -graph, $\mathcal{G}_{\mathbb{N}}$, with one vertex per projection and every edge represents a dependency between two projections. Every vertex is labelled by the diameter of the corresponding projection.

In the coming sections we show that there is no dependency structure (e.g. acyclic dependencies, symmetric dependencies, etc.) for which such a compositional upper-bounding scheme can work. This holds for both, (i) the diameter (Theorem 1), and (ii) the recurrence diameter (Theorem 2). Theorem 1 (resp. Theorem 2) shows that given projections' diameters (resp. recurrence diameters) and the dependencies between the projections, we can always construct a system with an arbitrarily large diameter, that has projections with diameters and dependencies as the given ones. Thus a function f that takes only projections' diameters and their dependencies cannot compute an upper bound on the given system's diameter, since they do not have enough information.

3.3.2 Diameter Cannot be Compositionally Upper-Bounded

Theorem 1. *For any function $f : \mathbb{N}\text{-graph} \Rightarrow \mathbb{N}$, and an \mathbb{N} -graph, $\mathcal{G}_{\mathbb{N}}$ that has at least one edge, there is a factored system δ such that: (i) $f(\mathcal{G}_{\mathbb{N}}) < d(\delta)$, and (ii) there is a lifted dependency graph \mathcal{G}_{VS} for δ , such that $\mathcal{G}_{\mathbb{N}} = \mathfrak{D}(\mathcal{G}_{VS})$, where $\mathfrak{D}(vs) = d(\delta|_{vs})$.²*

The proof is made of three main steps. Firstly, for each given projection diameter n (i.e. label in $\mathcal{G}_{\mathbb{N}}$) we construct a factored system γ_n , that is a path with diameter n . Those paths are constructed such that: i) their union is a system with a diameter that is more than $f(\mathcal{G}_{\mathbb{N}})$, and ii) they are projections of the final construction δ , so their domains are disjoint, as the projections are supposed to be performed on a partition. Secondly, for every dependency from projection γ^{u_1} to γ^{u_2} (i.e. an edge $\mathcal{G}_{\mathbb{N}}$), we construct an action that has preconditions from γ^{u_1} and effects from γ^{u_2} . Those actions are supposed to not change the state space of the final construction, they only add dependencies corresponding to the edges in $\mathcal{G}_{\mathbb{N}}$. Thirdly, we show that the union of the constructed projections and the dependency inducing actions is the required witness δ , i.e. its diameter exceeds $f(\mathcal{G}_{\mathbb{N}})$.

To facilitate our constructions, we use the following notation for states: states are annotated with a superscript, such that variables in states' domains are chosen to have different names if the superscripts of the states are different.³ Formally, for two states x_a^i and x_b^j , $\mathcal{D}(x_a^i) = \mathcal{D}(x_b^j)$ if $i = j$, and $\mathcal{D}(x_a^i) \cap \mathcal{D}(x_b^j) = \emptyset$ otherwise. A path system is defined as $\gamma_n^i = \{(x_j^i, x_{j+1}^i) \mid 1 \leq j \leq n + 1\}$, where the state space γ_n^i contains a simple path with $n + 1$ states. Note that $d(\gamma_n^i) = n$.

Proof. Let $n = f(\mathcal{G}_{\mathbb{N}}) + 1$. This is the number that the constructed system will have as its diameter. Take an arbitrary edge $(u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})$. For $u_3 \in V(\mathcal{G}_{\mathbb{N}})$, let γ^{u_3} denote the simple path $\gamma_{n+\mathcal{G}_{\mathbb{N}}(u_2)}^{u_2}$ for $u_3 = u_2$ and $\gamma_{\mathcal{G}_{\mathbb{N}}(u_3)}^{u_3}$ otherwise. Firstly, we construct a system δ_{u_2} whose state space has a simple path with $\mathcal{G}_{\mathbb{N}}(u_2)$ states and a clique with n states attached to the last state in the path. Thus, δ_{u_2} has $\mathcal{G}_{\mathbb{N}}(u_2)$ as its diameter. Formally $\delta_{u_2} = \{(x_i^{u_2}, x_{i+1}^{u_2}) \mid 1 \leq i \leq \mathcal{G}_{\mathbb{N}}(u_2)\} \cup \{(x_i^{u_2}, x_j^{u_2}) \mid \mathcal{G}_{\mathbb{N}}(u_2) \leq i, j \leq n + \mathcal{G}_{\mathbb{N}}(u_2) + 1\}$.

Consider a fresh variable z (i.e. $z \notin \mathcal{D}(\gamma^{u_3})$, for any $u_3 \in V(\mathcal{G}_{\mathbb{N}})$). Let $\delta'_{u_2} = \{(\{z\} \uplus p, e) \mid (p, e) \in \delta_{u_2}\} \cup \{(\{\bar{z}\} \uplus p, e) \mid (p, e) \in \gamma^{u_2}\}$, i.e. a transition system with two modes of operation determined by the value of z . If z is false then it operates as a path with a clique attached to its end, i.e. exactly like δ_{u_2} . Otherwise, it operates as a path of length $n + \mathcal{G}_{\mathbb{N}}(u_2) + 1$, i.e. exactly

² $\mathcal{G}_{\mathbb{N}}$ is a relabelling of \mathcal{G}_{VS} : a vertex labelled by vs in \mathcal{G}_{VS} is relabelled by the diameter of $\delta|_{vs}$

³In our proofs we use the nodes in the dependency graph as a superscript.




Figure 3.4: Referring to Example 4: (a) is the natural number labelled graph. (b), (c), (d) and (e) are the largest connected components in $\mathcal{G}(\gamma^{u_2})$, $\mathcal{G}(\gamma^{u_1})$, $\mathcal{G}(\gamma^{u_3})$, and $\mathcal{G}(\delta_{u_2})$, respectively. (f) is a lifted dependency graph of the factored system.

like γ^{u_2} . In this way the largest component in the state space of δ'_{u_2} has a short diameter if z is removed from the actions, otherwise the diameter blows up to $d(\delta'_{u_2}) = n + \mathcal{G}_{\mathbb{N}}(u_2)$.

Let $\delta_{u_1} = \{(\{\bar{z}\} \uplus p, \{\bar{z}\} \uplus e) \mid (p, e) \in \gamma^{u_1}\}$. This system operates exactly as the path γ^{u_1} when z is false. Since the value of z does not change, $d(\delta_{u_1}) = \mathcal{G}_{\mathbb{N}}(u_1)$. We now show that $\delta = \{(x_1^{u_3} \uplus x_1^{u_4}, x_2^{u_4}) \mid (u_3, u_4) \in E(\mathcal{G}_{\mathbb{N}})\} \cup \delta'_{u_2} \cup \delta_{u_1} \cup \bigcup_{u_3 \in V(\mathcal{G}_{\mathbb{N}})} \gamma^{u_3}$ is our required witness.

To show that δ satisfies the first requirement (i), consider the states $x_1 = \{\bar{z}\} \uplus \biguplus_{u_3 \in V(\mathcal{G}_{\mathbb{N}})} x_1^{u_3}$ and $x_2 = \{\bar{z}\} \uplus x_{n+\mathcal{G}_{\mathbb{N}}(u_2)+1}^{u_2} \uplus \biguplus_{u_3 \in V(\mathcal{G}_{\mathbb{N}}) \setminus u_2} x_1^{u_3}$. The action sequence $\vec{\pi} = [(\{\bar{z}\} \uplus x_1^{u_2}, x_2^{u_2}); (\{\bar{z}\} \uplus x_2^{u_2}, x_3^{u_2}) \dots; (\{\bar{z}\} \uplus x_{n+\mathcal{G}_{\mathbb{N}}(u_2)}^{u_2}, x_{n+\mathcal{G}_{\mathbb{N}}(u_2)+1}^{u_2})]$ reaches x_2 if executed at x_1 , and there is not an action sequence from δ^* that can reach x_2 from x_1 that is shorter than $\vec{\pi}$. This means that the diameter of δ is more than the length of $\vec{\pi}$, which is $n + \mathcal{G}_{\mathbb{N}}(u_2)$.

To show that δ satisfies the second requirement (ii) on the witness, consider a relabelling, \mathcal{G}_{VS} , of $\mathcal{G}_{\mathbb{N}}$, where u_1 is relabelled by $\mathcal{D}(\delta_{u_1})$ and every other every vertex u_3 is relabelled by the domain of the path γ^{u_3} . Recall that δ had the set of actions $\{(x_1^{u_3} \uplus x_1^{u_4}, x_2^{u_4}) \mid (u_3, u_4) \in E(\mathcal{G}_{\mathbb{N}})\}$ as a subset. These actions are constructed such that they add dependency from $\mathcal{D}(\gamma^{u_3})$ to $\mathcal{D}(\gamma^{u_4})$ in δ iff $(u_3, u_4) \in E(\mathcal{G}_{\mathbb{N}})$. This means that the edges of \mathcal{G}_{VS} represent the dependencies of δ and accordingly it is a lifted dependency graph of δ . Also since, $\delta|_{\mathcal{D}(\gamma^{u_1})} = \gamma^{u_1}$, for $u_1 \in V(\mathcal{G}_{\mathbb{N}})$, and since by construction $\mathcal{G}_{\mathbb{N}}$ is a relabelling of \mathcal{G}_{VS} , we have $\mathcal{G}_{\mathbb{N}} = \mathfrak{D}(\mathcal{G}_{VS})$. \square

Example 4. This example shows the previous construction for a function $f : \mathbb{N}\text{-graph} \Rightarrow \mathbb{N}$, and the graph $\mathcal{G}_{\mathbb{N}}$ shown in Figure 3.4a, where in this example $f(\mathcal{G}_{\mathbb{N}}) = 4$, and accordingly $n = 5$. In $\mathcal{G}_{\mathbb{N}}$ there are three vertices u_1 (the root), u_2 , and u_3 , labelled by the numbers 2, 3, and 2, respectively. In this case the first simple path system is $\gamma^{u_2} = \{(x_1^{u_2}, x_2^{u_2}), (x_2^{u_2}, x_3^{u_2}), (x_3^{u_2}, x_4^{u_2}), (x_4^{u_2}, x_5^{u_2}), (x_5^{u_2}, x_6^{u_2}), (x_6^{u_2}, x_7^{u_2})\}$, where $x_1^{u_2} = \{\bar{v}_3, \bar{v}_4, \bar{v}_5\}$, $x_2^{u_2} = \{\bar{v}_3, \bar{v}_4, v_5\}$, $x_3^{u_2} = \{\bar{v}_3, v_4, \bar{v}_5\}$, $x_4^{u_2} = \{\bar{v}_3, v_4, v_5\}$, $x_5^{u_2} = \{v_3, \bar{v}_4, \bar{v}_5\}$, $x_6^{u_2} = \{v_3, \bar{v}_4, v_5\}$, and $x_7^{u_2} = \{v_3, v_4, \bar{v}_5\}$. The second simple path system is $\gamma^{u_1} = \{(x_1^{u_1}, x_2^{u_1}), (x_2^{u_1}, x_3^{u_1}), (x_3^{u_1}, x_4^{u_1})\}$, where $x_1^{u_1} = \{\bar{v}_1, \bar{v}_2\}$, $x_2^{u_1} = \{\bar{v}_1, v_2\}$, $x_3^{u_1} = \{v_1, \bar{v}_2\}$, and $x_4^{u_1} = \{v_1, v_2\}$. The third simple path system is $\gamma^{u_3} = \{(x_1^{u_3}, x_2^{u_3}), (x_2^{u_3}, x_3^{u_3})\}$, where $x_1^{u_3} = \{\bar{v}_6, \bar{v}_7\}$, $x_2^{u_3} = \{\bar{v}_6, v_7\}$, and $x_3^{u_3} = \{v_6, \bar{v}_7\}$. The largest connected components of the state spaces of the three paths are shown in Figures 3.4b-3.4d, and their diameters are $d(\gamma^{u_2}) = n + \mathcal{G}_{\mathbb{N}}(u_1) = 6$, $d(\gamma^{u_1}) = \mathcal{G}_{\mathbb{N}}(u_2) = 3$, and $d(\gamma^{u_3}) = \mathcal{G}_{\mathbb{N}}(u_3) = 2$.

The system δ_{u_2} that has a simple path of length $\mathcal{G}_{\mathbb{N}}(u_2) = 3$ and a clique of size $n = 5$ attached

to its last state, $x_3^{u_2}$, is

$$\delta_{u_2} = \{(x_1^{u_2}, x_2^{u_2}), (x_2^{u_2}, x_3^{u_2})\} \cup \{(x_i^{u_2}, x_j^{u_2}) \mid 3 \leq i, j \leq 7\}$$

The largest component of its state space is shown in Figure 3.4e and $d(\delta_{u_2}) = \mathcal{G}_{\mathbb{N}}(u_2) = 3$. To construct the system that has two modes of operation, δ'_{u_2} , we use a variable z to determine the two modes of δ'_{u_2} . We add the literal z to the preconditions of every action in δ_{u_2} to ensure that they are activated when z is true, and we add the literal \bar{z} to the preconditions of every action in γ^{u_2} to ensure that they are activated when \bar{z} is false. Thus, $\delta'_{u_2} = \{(\{z\} \uplus p, e) \mid (p, e) \in \delta_{u_2}\} \cup \{(\{\bar{z}\} \uplus p, e) \mid (p, e) \in \gamma^{u_2}\}$, where for instance, the action $(x_4^{u_2}, x_5^{u_2})$ from δ_{u_2} becomes $(\{z\} \uplus x_4^{u_2}, x_5^{u_2}) = (\{z, \bar{v}_3, v_4, v_5\}, \{v_3, \bar{v}_4, \bar{v}_5\})$ and the action $(x_4^{u_2}, x_5^{u_2})$ from γ^{u_2} becomes $(\{\bar{z}\} \uplus x_4^{u_2}, x_5^{u_2}) = (\{\bar{z}\} \uplus x_4^{u_2}, x_5^{u_2}) = (\{\bar{z}, \bar{v}_3, v_4, v_5\}, \{v_3, \bar{v}_4, \bar{v}_5\})$.

To construct δ_{u_1} we add \bar{z} to the preconditions and effects of each action in γ^{u_1} . δ_{u_1} will have two modes of operation depending on the value of z . If z is true no actions at all can be executed, otherwise actions from γ^{u_1} can be executed, and accordingly $d(\delta_{u_1}) = d(\gamma^{u_1})$.

The required witness is $\delta = \{(x_1^{u_1} \uplus x_1^{u_2}, x_2^{u_2}), (x_1^{u_1} \uplus x_1^{u_3}, x_2^{u_3})\} \cup \delta'_{u_2} \cup \delta_{u_1} \cup \gamma^{u_3}$. The combined system has a diameter of 6 because it has a path isomorphic to the path in γ^{u_2} from $\{\bar{z}\} \uplus x_1^{u_2} \{\bar{z}\} \uplus x_7^{u_2}$. In particular, this path is between the state $x_1 = \{\bar{z}\} \uplus x_1^{u_2} \uplus x_1^{u_1} \uplus x_1^{u_3} = \{\bar{z}, \bar{v}_3, \bar{v}_4, \bar{v}_5, \bar{v}_1, \bar{v}_2, \bar{v}_6, \bar{v}_7\}$ and the state $x_2 = \{\bar{z}\} \uplus x_7^{u_2} \uplus x_1^{u_1} \uplus x_1^{u_3} \{\bar{z}, v_3, v_4, v_5, v_1, v_2, v_6, v_7\}$. The action sequence $\vec{\pi} = [(\bar{z} \uplus x_1^{u_2}, x_2^{u_2}); (\bar{z} \uplus x_2^{u_2}, x_3^{u_2}); (\bar{z} \uplus x_3^{u_2}, x_4^{u_2}); (\bar{z} \uplus x_4^{u_2}, x_5^{u_2}); (\bar{z} \uplus x_5^{u_2}, x_6^{u_2}); (\bar{z} \uplus x_6^{u_2}, x_7^{u_2})]$ reaches x_2 if executed at x_1 , and no action sequence shorter than $|\vec{\pi}|$ can reach x_2 from x_1 . This means that $n + \mathcal{G}_{\mathbb{N}}(u_1) = 6 \leq d(\delta)$. Also, Figure 3.4f is a lifted dependency graph of δ , as it has the actions $\{(x_1^{u_1} \uplus x_1^{u_2}, x_2^{u_2}), (x_1^{u_1} \uplus x_1^{u_3}, x_2^{u_3})\}$.

3.3.3 Recurrence Diameter Cannot be Compositionally Bounded

Theorem 2. For a function $f : \mathbb{N}\text{-graph} \Rightarrow \mathbb{N}$, and an \mathbb{N} -graph, $\mathcal{G}_{\mathbb{N}}$, there is a system δ where: (i) $f(\mathcal{G}_{\mathbb{N}}) < rd(\delta)$, and (ii) there is a lifted dependency graph \mathcal{G}_{VS} for δ , such that $\mathcal{G}_{\mathbb{N}} = \mathfrak{R}(\mathcal{G}_{VS})$, where $\mathfrak{R}(vs) = rd(\delta|_{vs})$ (i.e. $\mathcal{G}_{\mathbb{N}}$ is a relabelling of \mathcal{G}_{VS} : a vertex labelled by vs in \mathcal{G}_{VS} is relabelled by the recurrence diameter of $\delta|_{vs}$).

The proof of this theorem is similar in structure to that of Theorem 1. One difference is that instead of constructing paths for projections, we construct systems whose state spaces are “flowers” (Figure 3.5). A flower can have arbitrarily many states without changing its recurrence diameter. This is needed to construct arbitrarily long paths with distinct all states in the final construction, while keeping the recurrence diameters of projections constant. Flowers are formally described as follows. First let a petal $\eta_{l,m}^i$ be a loop of $\frac{m}{2} + 1$ states labelled with 0 and $\frac{m}{2}$ other numbers. We formally define it as $\eta_{l,m}^i = \{(x_{(l-2)c+k}^i, x_{(l-2)c+(k+1 \bmod a)}^i) \mid 0 \leq k \leq a-1\}$, where $a = \lceil \frac{m}{2} \rceil + 1$ if m is odd, and $a = \frac{m}{2} + 1$ otherwise, and $c = \lfloor \frac{m}{2} \rfloor + 1$. A flower system, $\mathfrak{I}_{l,m}^i$, denotes $\bigcup_{1 \leq j < k \leq l} \{(x_j^i, x_k^i), (x_k^i, x_j^i)\}$ (i.e. a clique) if $m = 1$, and $\eta_{l,m}^i \cup \bigcup_{1 \leq j \leq l-2} \{(x_{jc+k}^i, x_{jc+(k+1 \bmod c)}^i) \mid 0 \leq k \leq c-1\}$ otherwise. $\mathfrak{I}_{l,m}^i$ is a factored system whose state space contains l cycles (the petals). All other petals are of length $\lfloor \frac{m}{2} \rfloor + 1$, except for the last one, if m is odd its length is $\lceil \frac{m}{2} \rceil + 1$ otherwise its length is $\frac{m}{2} + 1$. All of the petals have exactly one state in common between them (the pistil), which is x_0^i . In each petal j there is a state t_j^i (the “petal tip”). $t_j^i = x_{(l-2)c+\lceil \frac{a}{2} \rceil}^i$ for the last petal (i.e. $j = l-1$), and otherwise $t_j^i = x_{jc+\lceil \frac{c}{2} \rceil}^i$. t_0^i can be reached by exactly $\lceil \frac{a}{2} \rceil$ actions from x_0^i , and x_0^i can be reached from t_0^i by exactly $\lfloor \frac{a}{2} \rfloor$ actions. Similarly, for $j \neq 1$, t_j^i can be reached by exactly $\lceil \frac{c}{2} \rceil$ actions from x_0^i , and x_0^i can be reached from t_j^i by exactly $\lfloor \frac{c}{2} \rfloor$ actions. Denoting by $\vec{\pi}_{a \rightarrow b}^i$ the shortest action sequence that joins the petal tip




Figure 3.5: The largest connected component in the state space of a flower with $2 \leq m$.

t_a^i with t_b^i , $c \leq |\vec{\pi}_{a \rightarrow b}^i|$ holds. Also $rd(\mathbb{I}_{l,m}^i) = m$ holds. Lastly we use $\mathcal{S}(\mathbb{I}^u)$ to denote the largest connected component for any state in the state space of \mathbb{I}^u , which is always unique for a flower

Proof. Let $n = f(\mathcal{G}_{\mathbb{N}})$. This is the number that the constructed system will have as its recurrence diameter. For every vertex $u_1 \in V(\mathcal{G}_{\mathbb{N}})$, let \mathbb{I}^{u_1} denote the flower $\mathbb{I}_{n+2, \mathcal{G}_{\mathbb{N}}(u_1)}^{u_1}$. Let $\delta = \{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}) \mid (u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})\} \cup \bigcup_{u_1 \in V(\mathcal{G}_{\mathbb{N}})} \mathbb{I}^{u_1}$. For each flower \mathbb{I}^{u_1} , denote its $n+2$ petal tips by $t_j^{u_1}$, for $0 \leq j \leq n+1$. Consider the two states: (i) $x_1 = \biguplus_{u_1 \in V(\mathcal{G}_{\mathbb{N}})} t_0^{u_1}$, and (ii) $x_2 = \biguplus_{u_1 \in V(\mathcal{G}_{\mathbb{N}})} t_{n+1}^{u_1}$. We show that δ satisfies requirement (i) on the witness by giving a simple path between the states x_1 and x_2 , whose length is more than n . Consider the action sequence $\vec{\pi}$ constructed by Algorithm 1.

Algorithm 1:

```

 $\vec{\pi} = []$ 
for  $0 \leq i \leq n+1$  // Looping over petals
  for  $u_1 \in V(\mathcal{G}_{\mathbb{N}})$  // Looping over flowers
     $\vec{\pi} =: \vec{\pi} \frown \vec{\pi}_{i \rightarrow i+1}^{u_1}$  4 // Go from petal tip  $t_i^{u_1}$  to  $t_{i+1}^{u_1}$ 

```

The length of $\vec{\pi}$ is at least $(n+2) \sum_{u_1 \in V(\mathcal{G}_{\mathbb{N}})} (c) - 1$, since $c \leq |\vec{\pi}_{a \rightarrow b}^{u_1}|$ for $a \neq b$. Since $c = \lfloor \frac{\mathcal{G}_{\mathbb{N}}(u_1)}{2} \rfloor + 1$ then $n \leq |\vec{\pi}|$. Also $\vec{\pi}$ is constructed such that, if $\vec{\pi}$ is executed at x_1 it only traverses distinct states, i.e. it induces a simple path.

To show that δ satisfies requirement (ii), consider a relabelling, \mathcal{G}_{VS} , of $\mathcal{G}_{\mathbb{N}}$, where every vertex u_1 is relabelled by $\mathcal{D}(\mathbb{I}^{u_1})$, the domain of the flower \mathbb{I}^{u_1} associated with the vertex. Recall that

⁴For two lists l_1 and l_2 , $l_1 \frown l_2$ denotes their concatenation.

δ had the set of actions $\{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}) \mid (u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})\}$ as a subset. These actions are constructed such that they add dependency from $\mathcal{D}(\mathbb{I}^{u_1})$ to $\mathcal{D}(\mathbb{I}^{u_2})$ in δ iff $(u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})$. This means that the edges of \mathcal{G}_{VS} represent the dependencies of δ and accordingly it is a lifted dependency DAG of δ . Also since $\delta|_{\mathcal{D}(\mathbb{I}^{u_1})} = \mathbb{I}^{u_1}$, and by construction $\mathcal{G}_{\mathbb{N}}$ is a relabelling of \mathcal{G}_{VS} , we have $\mathcal{G}_{\mathbb{N}} = \mathfrak{R}(\mathcal{G}_{\text{VS}})$. \square

Example 5. *This example shows the previous construction for a function $f : \mathbb{N}\text{-graph} \Rightarrow \mathbb{N}$, and the graph $\mathcal{G}_{\mathbb{N}}$ shown in Figure 3.4a, where $f(\mathcal{G}_{\mathbb{N}}) = 2$. In this graph there are three vertices u_1 (the root), u_2 , and u_3 , labelled by the numbers 2, 3, and 2, respectively. We construct three flowers, one per vertex. For u_2 the constructed flower is $\mathbb{I}_{4,3}^{u_2} = \{(x_0^{u_2}, x_4^{u_2}), (x_4^{u_2}, x_5^{u_2}), (x_5^{u_2}, x_0^{u_2})\} \cup \bigcup \{(x_0^{u_2}, x_l^{u_2}), (x_l^{u_2}, x_0^{u_2})\} \mid 1 \leq l \leq 3\}$ where the states are defined as follows $x_0^{u_2} = \{\bar{v}_4, \bar{v}_5, \bar{v}_6\}$ (the pistil), $x_1^{u_2} = \{\bar{v}_4, \bar{v}_5, v_6\}$ (the first petal tip), $x_2^{u_2} = \{\bar{v}_4, v_5, \bar{v}_6\}$ (the second petal tip), $x_3^{u_2} = \{\bar{v}_4, v_5, v_6\}$ (the third petal tip), $x_4^{u_2} = \{v_4, \bar{v}_5, \bar{v}_6\}$, $x_5^{u_2} = \{v_4, \bar{v}_5, v_6\}$ (the fourth petal tip). Note that since for the flower $\mathbb{I}_{4,3}^{u_2}$, $m = 3$ (i.e. odd), the last petal has one more state in it (the petal at the bottom of Figure 3.6a). For u_3 the constructed flower is $\mathbb{I}_{4,2}^{u_3} = \bigcup \{(x_0^{u_3}, x_l^{u_3}), (x_l^{u_3}, x_0^{u_3})\} \mid 1 \leq l \leq 4\}$, where the states are defined as follows $x_0^{u_3} = \{\bar{v}_7, \bar{v}_8, \bar{v}_9\}$ (the pistil), $x_1^{u_3} = \{\bar{v}_7, \bar{v}_8, v_9\}$ (the first petal tip), $x_2^{u_3} = \{\bar{v}_7, v_8, \bar{v}_9\}$ (the second petal tip), $x_3^{u_3} = \{\bar{v}_7, v_8, v_9\}$ (the third petal tip), and $x_4^{u_3} = \{v_7, \bar{v}_8, \bar{v}_9\}$ (the fourth petal tip). For u_1 the constructed flower is $\mathbb{I}_{4,2}^{u_1} = \bigcup \{(x_0^{u_1}, x_l^{u_1}), (x_l^{u_1}, x_0^{u_1})\} \mid 1 \leq l \leq 4\}$, where the states are defined as follows $x_0^{u_1} = \{\bar{v}_1, \bar{v}_2, \bar{v}_3\}$ (the pistil), $x_1^{u_1} = \{\bar{v}_1, \bar{v}_2, v_3\}$ (the first petal tip), $x_2^{u_1} = \{\bar{v}_1, v_2, \bar{v}_3\}$ (the second petal tip), $x_3^{u_1} = \{\bar{v}_1, v_2, v_3\}$ (the third petal tip), and $x_4^{u_1} = \{v_1, \bar{v}_2, \bar{v}_3\}$ (the fourth petal tip).*

The required witness $\delta = \mathbb{I}_{4,2}^{u_1} \cup \mathbb{I}_{4,3}^{u_2} \cup \mathbb{I}_{4,2}^{u_3} \cup \{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}), (x_0^{u_1} \uplus x_0^{u_3}, x_1^{u_3})\}$ where the actions $\{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}), (x_0^{u_1} \uplus x_0^{u_3}, x_1^{u_3})\}$ add to δ dependencies equivalent to the edges of $\mathcal{G}_{\mathbb{N}}$, i.e. the dependencies shown in Figure 3.6d.

Consider the states $x_1 = x_1^{u_1} \uplus x_1^{u_2} \uplus x_1^{u_3} = \{\bar{v}_1, \bar{v}_2, v_3, \bar{v}_4, \bar{v}_5, v_6, \bar{v}_7, \bar{v}_8, v_9\}$ and $x_2 = x_4^{u_1} \uplus x_4^{u_2} \uplus x_4^{u_3} = \{v_1, \bar{v}_2, \bar{v}_3, \bar{v}_4, \bar{v}_5, v_6, v_7, \bar{v}_8, \bar{v}_9\}$. Following Algorithm 1, the resulting action sequence is $\vec{\pi} = \vec{\pi}_{1 \rightarrow 2} \frown \vec{\pi}_{1 \rightarrow 2} \frown \vec{\pi}_{1 \rightarrow 2} \frown \vec{\pi}_{2 \rightarrow 3} \frown \vec{\pi}_{2 \rightarrow 3} \frown \vec{\pi}_{2 \rightarrow 3} \frown \vec{\pi}_{2 \rightarrow 3} \frown \vec{\pi}_{3 \rightarrow 4} \frown \vec{\pi}_{3 \rightarrow 4} \frown \vec{\pi}_{3 \rightarrow 4}$ and its length is 18. $\vec{\pi}$ will reach x_2 if executed at x_1 , while traversing all distinct states. The largest connected component of $\mathcal{G}(\delta)$ and the path traversed by executing $\vec{\pi}$ from x_1 are shown in Figure 3.6e.

3.3.4 Discussion

In both proofs above, it might seem that the constructions and accordingly the theorem statements are based on a bad choice of projections. However, the point that these results make is subtle but more profound: any compositional algorithm cannot know whether the given projections are “sound” just by “looking” in the dependency graph. In other words, if a compositional algorithm only uses the dependency graph, then there will always be a rouge projection. Thus to guarantee that the given projections can be soundly used for upper bounding, an algorithm needs more information than just what is available in the dependency graph. This accordingly shows that all existing compositional bounding algorithms that use only dependency structures (which are the majority) need some non-trivial extensions that look into more than just dependencies (e.g. they need to look into state space structures) in order for them to bound the system’s (recurrence) diameter using projections’ (recurrence) diameters.

3.4 The Traversal Diameter

The traversal diameter is one less than the largest number of states that can be traversed by any path. It is defined as follows.

Definition 11 (Traversal Diameter). *Let $\mathbf{ss}(x, \vec{\pi})$ be the set of states traversed by executing $\vec{\pi}$ from x .*

$$td(\delta) = \max_{x \in \mathbb{U}(\delta)} \max_{\vec{\pi} \in \delta^*} |\mathbf{ss}(x, \vec{\pi})| - 1.$$

Example 6. *For $\delta|_{vs}$ from Example 2 the traversal diameter is 3, since the longest path in its state space has 4 states.*

The most appealing feature of the traversal diameter is that it is an upper bound on the recurrence diameter (and accordingly the diameter) that can be compositionally bounded with projections from arbitrary partitions of the state variables, as shown in the following theorem.

Theorem 3. *For a factored representation δ and a partition $vs_{1..n}$ of $\mathcal{D}(\delta)$, $td(\delta) \leq \prod_{vs \in vs_{1..n}} (td(\delta|_{vs}) + 1) - 1$.*

Proposition 1. *For some k , if for every $x \in \mathbb{U}(\delta)$ and $\vec{\pi} \in \delta^*$, $|\mathbf{ss}(x, \vec{\pi})| \leq k+1$, then $td(\delta) \leq k$.*

Proposition 2. *For a set of states S , let $S|_{vs}$ denote $\{x|_{vs} \mid x \in S\}$. Let $\mathbf{sat-pre}(x, \vec{\pi})$ denote that preconditions of every action in $\vec{\pi}$ are satisfied, if $\vec{\pi}$ is executed from x . If $\mathbf{sat-pre}(x, \vec{\pi})$, then $\mathbf{ss}(x, \vec{\pi})|_{vs} = \mathbf{ss}(S|_{vs}, \vec{\pi}|_{vs})$.*

Proposition 3. *For any partition $vs_{1..n}$ of $\mathcal{D}(\delta)$, $|\mathbf{ss}(x, \vec{\pi})| \leq \prod_{vs \in vs_{1..n}} |\mathbf{ss}(x, \vec{\pi})|_{vs}$.*

The proof of the above proposition comes from the fact that a set of states is a subset of the cartesian product of its own projections, given that the projection is on a partition of the state variables.

Proof of Theorem 3. Consider $x \in \mathbb{U}(\delta)$ and without loss of generality, an action sequence $\vec{\pi} \in \delta^*$ such that $\mathbf{sat-pre}(x, \vec{\pi})$. From Definition 11, for any vs , $x|_{vs} \in \mathbb{U}(\delta|_{vs})$ and $\vec{\pi}|_{vs} \in \delta|_{vs}^*$, we have $|\mathbf{ss}(x|_{vs}, \vec{\pi}|_{vs})| - 1 \leq td(\delta|_{vs})$. Theorem 3 then follows from Proposition 2, Proposition 3 and Proposition 1. \square

The bound in the previous theorem is a polynomial that has all possible terms from the traversal diameters of the projections. This bound cannot be improved regardless of the dependency structure inducing the projections, as follows.

Theorem 4. *For any \mathbb{N} -graph, $\mathcal{G}_{\mathbb{N}}$, there is a factored system δ such that: (i) $\prod_{u_1 \in V(\mathcal{G}_{\mathbb{N}})} (\mathcal{G}_{\mathbb{N}}(u_1) + 1) - 1 \leq td(\delta)$, and (ii) there is a lifted dependency graph \mathcal{G}_{VS} for δ , such that $\mathcal{G}_{\mathbb{N}} = \mathfrak{T}(\mathcal{G}_{VS})$, where $\mathfrak{T}(vs) = td(\delta|_{vs})$.*

Proof. For $u \in V(\mathcal{G}_{\mathbb{N}})$, let \mathbb{I}^u denote the flower $\mathbb{I}_{\mathcal{G}_{\mathbb{N}}(u), 2}^u$. For all $u \in V(\mathcal{G}_{\mathbb{N}})$, $x \rightsquigarrow y$ holds for any $x, y \in \mathcal{S}(\mathbb{I}^u)$, thus $td(\mathbb{I}^u) = |\mathcal{S}(\mathbb{I}^u)| - 1 = \mathcal{G}_{\mathbb{N}}(u)$. Let $\delta = \{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}) \mid (u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})\} \cup \bigcup_{u \in V(\mathcal{G}_{\mathbb{N}})} \mathbb{I}^u$. We now show that δ satisfies requirement (i). First let $\mathcal{S}(\delta)$ denote the largest connected component in the state space of δ , which is unique. Since $x \rightsquigarrow y$ holds for any $x, y \in \mathcal{S}(\mathbb{I}^u)$, then $x \rightsquigarrow y$ holds for any $x, y \in \mathcal{S}(\delta)$, and therefore there is a path that traverses every member of $\mathcal{S}(\delta)$. Since for $u_1 \neq u_2$ we have $\mathbb{I}^{u_1} \cap \mathbb{I}^{u_2} = \emptyset$, we have $|\mathcal{S}(\delta)| = \prod_{u \in V(\mathcal{G}_{\mathbb{N}})} |\mathcal{S}(\mathbb{I}^u)|$. Since $\prod_{u \in V(\mathcal{G}_{\mathbb{N}})} |\mathcal{S}(\mathbb{I}^u)| = \prod_{u \in V(\mathcal{G}_{\mathbb{N}})} (\mathcal{G}_{\mathbb{N}}(u) + 1)$, we have that $\prod_{u \in V(\mathcal{G}_{\mathbb{N}})} (\mathcal{G}_{\mathbb{N}}(u) + 1) - 1 \leq td(\delta)$.

To show that δ satisfies requirement (ii), consider a relabelling, \mathcal{G}_{VS} , of $\mathcal{G}_{\mathbb{N}}$, where every vertex u is relabelled by the domain of the flower \mathbb{I}^u . Recall that δ had the set of actions $\{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}) \mid (u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})\}$ as a subset. These actions are constructed such that they add dependency from $\mathcal{D}(\mathbb{I}^{u_1})$ to $\mathcal{D}(\mathbb{I}^{u_2})$ in δ iff $(u_1, u_2) \in E(\mathcal{G}_{\mathbb{N}})$. Accordingly edges of \mathcal{G}_{VS} represent the dependencies of δ and accordingly it is a lifted dependency graph of δ . Also since $\delta|_{\mathcal{D}(\mathbb{I}^u)} = \mathbb{I}^u$, for $u \in V(\mathcal{G}_{\mathbb{N}})$, and since by construction $\mathcal{G}_{\mathbb{N}}$ is a relabelling of \mathcal{G}_{VS} , we have $\mathcal{G}_{\mathbb{N}} = \mathfrak{T}(\mathcal{G}_{\text{VS}})$. \square

Example 7. This is an example of the construction from Theorem 4, for the natural number labelled DAG in Figure 3.4a. We construct three flowers, one per vertex, shown in Figures 3.7a-3.7c. For u_2 the constructed flower is $\mathbb{I}_{3,2}^{u_2} = \{(x_0^{u_2}, x_1^{u_2}), (x_0^{u_2}, x_2^{u_2}), (x_0^{u_2}, x_3^{u_2}), (x_1^{u_2}, x_0^{u_2}), (x_2^{u_2}, x_0^{u_2}), (x_3^{u_2}, x_0^{u_2})\}$. The states are defined as follows $x_0^{u_2} = \{\overline{v_3}, \overline{v_4}\}$ (the pistil), $x_1^{u_2} = \{\overline{v_3}, v_4\}$ (the first petal tip), $x_2^{u_2} = \{v_3, \overline{v_4}\}$ (the second petal tip), $x_3^{u_2} = \{v_3, v_4\}$ (the third petal tip). For u_3 the constructed flower is $\mathbb{I}_{2,2}^{u_3} = \{(x_0^{u_3}, x_1^{u_3}), (x_0^{u_3}, x_2^{u_3}), (x_1^{u_3}, x_0^{u_3}), (x_2^{u_3}, x_0^{u_3})\}$. The states are defined as follows $x_0^{u_3} = \{\overline{v_5}, \overline{v_4}\}$ (the pistil), $x_1^{u_3} = \{\overline{v_5}, v_4\}$ (the first petal tip) and $x_2^{u_3} = \{v_5, \overline{v_4}\}$ (the second petal tip). For u_1 the constructed flower is $\mathbb{I}_{2,2}^{u_1} = \{(x_0^{u_1}, x_1^{u_1}), (x_0^{u_1}, x_2^{u_1}), (x_1^{u_1}, x_0^{u_1}), (x_2^{u_1}, x_0^{u_1})\}$. The states are defined as follows $x_0^{u_1} = \{\overline{v_1}, \overline{v_2}\}$ (the pistil), $x_1^{u_1} = \{\overline{v_1}, v_2\}$ (the first petal tip), and $x_2^{u_1} = \{v_1, \overline{v_2}\}$ (the second petal tip). For any $x_0, x_1 \in \mathcal{S}(\mathbb{I}^{u_1})$, $x_0 \rightsquigarrow x_1$ holds and accordingly $td(\mathbb{I}^{u_1}) = 2$. Similarly, $td(\mathbb{I}^{u_2}) = 3$ and $td(\mathbb{I}^{u_3}) = 2$.

The required witness is $\delta = \{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}), (x_0^{u_1} \uplus x_0^{u_3}, x_1^{u_3})\} \cup \mathbb{I}^{u_1} \cup \mathbb{I}^{u_2} \cup \mathbb{I}^{u_3}$, where the actions $\{(x_0^{u_1} \uplus x_0^{u_2}, x_1^{u_2}), (x_0^{u_1} \uplus x_0^{u_3}, x_1^{u_3})\}$ add dependencies equivalent to the edges of $A_{\mathbb{N}}$, i.e. the dependencies shown in Figure 3.7d. Also, in the constructed witness, for all states $x_0, x_1 \in \mathcal{S}(\delta)$ (shown in Figure 3.7) $x_0 \rightsquigarrow x_1$ holds, and accordingly $td(\delta) = 35$.

There are two reasons the traversal diameter and the above theorems are interesting. Firstly, the traversal diameter is a non-trivial upper bound on the recurrence diameter, in the sense that it can be exponentially smaller than the number of states, which is the best known upper bound on the recurrence diameter. Secondly, the traversal diameter can be computed in linear time, unlike the NP-hard recurrence diameter.

3.4.1 Tightness of the Traversal Diameter

Theorem 5. There are infinitely many factored systems whose traversal diameter is (i) equal to their recurrence diameter, and (ii) exponentially smaller (in the number of state variables) than their state space. There are also infinitely many factored systems whose traversal diameter is (i) exponentially larger (in the number of state variables) than their recurrence diameter, and (ii) equal to the size of the largest connected component on their state space.

Theorem 5 follows from Lemma 2 and Lemma 1:

Lemma 1. There are infinitely many factored systems whose traversal diameters are exponentially smaller (in the number of state variables) than the size of their state spaces.

Proof. For an arbitrary number $n \in \mathbb{N}$, we construct a system whose state space size is a factor of n more than its traversal diameter. Consider the system $\{(x_0^i, x_i^1) \mid 1 \leq i \leq n\}$. The traversal diameter of this system is 1 since, the only possible transitions are from state x_0^1 to a state x_i^1 , for $1 \leq i \leq n$. However the system's state space has at least $n + 1$ states. \square

Lemma 2. There are infinitely many factored systems whose recurrence diameters are exponentially smaller (in the number of state variables) than their traversal diameters.

Proof. We show that for an arbitrary number $n \in \mathbb{N}$, there is a system whose traversal diameter is a factor of n more than its recurrence diameter. Consider the flower system $\mathbb{I}_{2n,2}^1$. The recurrence diameter of this flower is 2 since the length of each of its petals is 1, and any action sequence that traverses more than 3 states will traverse x_0^1 (the pistil) more than once. Its traversal diameter on the other hand is $2n$, since for any $x, y \in \mathcal{S}(\mathbb{I}_{2n,2}^1)$, $x \rightsquigarrow y$ holds and $|\mathcal{S}(\mathbb{I}_{2n,2}^1)| = 2n + 1$. \square

3.4.2 Computing the Traversal Diameter

An important aspect of td is that it can be computed in linear time using the following algorithm. However, we first formally define the *quotient* of a labelled digraph that is induced by repetitive vertex contraction operations, and define a function that computes the longest path in a given digraph weighted by a given function on the vertices.

Definition 12 (Digraph Quotient). *For a digraph \mathcal{G}_α , and a partition P of $V(\mathcal{G}_\alpha)$, the quotient \mathcal{G}_α/P of \mathcal{G}_α is the digraph:*

- (i) $V(\mathcal{G}_\alpha/P) = P$,
- (ii) $E(\mathcal{G}_\alpha/P) = \{(U, W) \mid U, W \in P \wedge \exists u_2 \in U, u_3 \in W. (u_2, u_3) \in E(\mathcal{G}_\alpha)\}$, and
- (iii) the label $\mathcal{G}_\alpha/P(U)$ is $\{\mathcal{G}_\alpha(u_1) \mid u_1 \in U\}$, for $U \in V(\mathcal{G}_\alpha/P)$.

Definition 13 (Weighted Longest Path). *For a digraph \mathcal{G}_α , let the function $b : \alpha \Rightarrow \mathbb{N}$ be a function that assigns a natural number for the label of every vertex in $V(\mathcal{G}_\alpha)$. \mathbf{S} is*

$$\mathbf{S}\langle b \rangle(u_1, \mathcal{G}_\alpha) = b(\mathcal{G}_\alpha(u_1)) + \max_{u_1' \in \text{children}_{\mathcal{G}_\alpha}(u_1)} (\mathbf{S}\langle b \rangle(u_1', \mathcal{G}_\alpha) + 1)$$

Then, let $\mathbf{S}_{\max}\langle b \rangle(\mathcal{G}_\alpha) = \max_{u_1 \in V(\mathcal{G}_\alpha)} \mathbf{S}\langle b \rangle(u_1, \mathcal{G}_\alpha)$.

\mathbf{S}_{\max} is only well-defined if \mathcal{G}_α is acyclic. The definition of \mathbf{S}_{\max} follows the scheme of an algorithm that finds the length of the longest path in a DAG, so its runtime is linear in the size of $V(\mathcal{G}_\alpha)$, if the values of \mathbf{S} for different vertices are memoised.⁵ However, \mathbf{S}_{\max} computes the weightiest path, where an edge weighs one, and a vertex u_1 weighs $b(\mathcal{G}_\alpha(u_1))$.

Algorithm 2: TRAVDIAM(δ)

$SCC :=$ set of strongly connected components of $\mathcal{G}(\delta)$
 $\mathcal{G} := \mathcal{G}(\delta)/SCC$
 Return $\mathbf{S}_{\max}\langle \mathbb{C} \rangle(\mathcal{G}) - 1$, where $\mathbb{C}(scc) = |scc|$ for $scc \in SCC$

Theorem 6. $\text{TRAVDIAM}(\delta) = td(\delta)$.

Proof. For notational brevity, let $\mathcal{G} = \mathcal{G}(\delta)/SCC$, and for a strongly connected component scc , $\mathbf{S}(scc) = \mathbf{S}\langle \mathbb{C} \rangle(scc, \mathcal{G})$ and $\text{children}(scc) = \text{children}_{\mathcal{G}}(scc)$. Since \mathcal{G} is a DAG, its vertices can be topologically ordered in a list l_{SCC} .

Firstly, we prove $\text{TRAVDIAM}(\delta) \leq td(\delta)$. We show that for any strongly connected component $scc \in \mathcal{G}$ there is an action sequence $\vec{\pi}_{scc} \in \delta^*$ and a state $x_{scc} \in scc$, such that $\mathbf{S}(scc) \leq |\mathbf{SS}(x_{scc}, \vec{\pi}_{scc})|$, which from Definitions 11 and 13, implies $\text{TRAVDIAM}(\delta) \leq td(\delta)$. We prove this by induction on l_{SCC} . The base case, $l_{SCC} = []$, is straightforward. For the step

⁵This assumes that b is at most of linear complexity.

case $l_{SCC} = scc :: l'_{SCC}$,⁶ and for any $scc' \in l'_{SCC}$, there is a state $x' \in scc'$ and an action sequence $\vec{\pi}' \in \delta^*$ where $\mathbf{S}(scc') \leq |\mathbf{SS}(x', \vec{\pi}')|$. Since scc is a strongly connected component of states in $\mathcal{G}(\delta)$, then there is $\vec{\pi}'_{scc} \in \delta^*$ and a state $x_{scc} \in scc$, where $\vec{\pi}'_{scc}$ traverses exactly all the states in scc , if executed at $x_{scc} \in scc$, i.e. $\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) = scc$. We have two cases:

Case 1 ($\text{children}(scc) = \emptyset$). From Definition 13, $\mathbf{S}(scc) = |scc| = |\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})|$ holds for this case. Accordingly the required witness $\vec{\pi}'_{scc}$ is the same as $\vec{\pi}'_{scc}$.

Case 2 ($\text{children}(scc) \neq \emptyset$). Let scc_{\max} be the strongly connected component such that $\forall scc' \in \text{children}(scc)$. $\mathbf{S}(scc) \leq \mathbf{S}(scc_{\max})$. Because $scc_{\max} \in \text{children}(scc)$ we have $scc_{\max} \in l'_{SCC}$, and accordingly from the inductive hypothesis there are $x_{\max} \in scc_{\max}$ and $\vec{\pi}'_{\max} \in \delta^*$ such that $\mathbf{S}(scc_{\max}) \leq |\mathbf{SS}(x_{\max}, \vec{\pi}'_{\max})|$.^(†) Also, from $scc_{\max} \in \text{children}(scc)$ and the fact that both scc and scc_{\max} are strongly connected components, there must be $\vec{\pi}' \in \delta^*$, such that $\text{ex}(x_{scc}, \vec{\pi}'_{scc} \frown \vec{\pi}') = x_{\max}$. We now show that $\vec{\pi}'_{scc} = \vec{\pi}'_{scc} \frown \vec{\pi}' \frown \vec{\pi}'_{\max}$ is the required witness. First it is easy to see that $\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \cup \mathbf{SS}(\text{ex}(x_{scc}, \vec{\pi}'_{scc}), \vec{\pi}') \cup \mathbf{SS}(x_{\max}, \vec{\pi}'_{\max}) = \mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})$. Since $scc_{\max} \in \text{children}(scc)$, we have that $\mathbf{SS}(x_{\max}, \vec{\pi}'_{\max})$ is disjoint with scc and accordingly $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| + |\mathbf{SS}(x_{\max}, \vec{\pi}'_{\max})| \leq |\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})|$. From this, (†), and Definition 13 we have $\mathbf{S}(scc) \leq |\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})|$.

Secondly, we prove $td(\delta) \leq \text{TRAVDIAM}(\delta)$ by showing that for any $scc \in \mathcal{G}$, $x_{scc} \in scc$, and $\vec{\pi}'_{scc} \in \delta$, we have $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq \mathbf{S}(scc)$. Our proof is again by induction on the list l_{SCC} . The base case, $l_{SCC} = []$, is straightforward. The step case $l_{SCC} = scc :: l'_{SCC}$, and we have that for any strongly connected component $scc' \in l'_{SCC}$, $x' \in scc'$, and $\vec{\pi}' \in \delta^*$, $|\mathbf{SS}(x', \vec{\pi}')| \leq \mathbf{S}(scc')$ holds. We have two cases:

Case 1 ($\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \subseteq scc$). Since $\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \subseteq scc$ then $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq |scc|$. From Definition 13, we know that $|scc| \leq \mathbf{S}(scc)$, and accordingly $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq \mathbf{S}(scc)$.

Case 2 ($\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \not\subseteq scc$). Since $\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \not\subseteq scc$, then there are $\vec{\pi}'_{scc}$, π , and $\vec{\pi}'_{\text{children}}$ such that: (i) $\vec{\pi}'_{scc} = \vec{\pi}'_{scc} \frown \pi :: \vec{\pi}'_{\text{children}}$, (ii) $\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc}) \subseteq scc$, and (iii) letting $x_{\text{children}} = \text{ex}(\text{ex}(x_{scc}, \vec{\pi}'_{scc}), \pi)$, $x_{\text{children}} \in scc_{\text{children}}$ holds, for some $scc_{\text{children}} \in \text{children}(scc)$. Using the same argument as the last case, we have $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq |scc|$.^(*) Since $x_{\text{children}} \in scc_{\text{children}}$, and from the inductive hypothesis, we have that $|\mathbf{SS}(x_{\text{children}}, \vec{\pi}'_{\text{children}})| \leq \mathbf{S}(scc_{\text{children}})$. Then using (*) and since $\mathbf{SS}(x_{\text{children}}, \vec{\pi}'_{\text{children}})$ and scc are disjoint, we have $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq |scc| + \mathbf{S}(scc_{\text{children}})$. From Definition 13, we have $|scc| + \mathbf{S}(scc_{\text{children}}) \leq \mathbf{S}(scc)$ and accordingly $|\mathbf{SS}(x_{scc}, \vec{\pi}'_{scc})| \leq \mathbf{S}(scc)$. \square

3.5 Using Structural Knowledge for Better Bounds: Acyclic Dependency

In this section we investigate restricting the type of abstractions to get better compositional bounding properties of different diameter functions. We consider abstractions induced by acyclicity in the variable dependency graph.

Although both, the diameter and the recurrence diameter are not susceptible to compositional bounding as shown in Theorem 2 and Theorem 1, Theorem 3 suggests they can be bounded via

⁶For a list l , $h :: l$ is l with the element h appended to its front.

compositional bounding of the traversal diameter. This guarantees the main utility of compositional bounding: restricting computations to be done using small abstractions, under very general conditions. However, bounds computed using this technique can be very loose, especially for the diameter, and Theorem 4 shows that it is impossible to improve them.

Restricting the projections to be induced by acyclic dependency graphs gives rise to a new possibility: upper-bounding the diameter through computing the recurrence diameter of projections. We propose a new way of upper-bounding the diameter of a factored system through computing the recurrence diameters of projections. This has two advantages over just using the aforementioned technique of compositional bounding using the traversal diameter.

- (i) We produce polynomials that are substantially tighter than the product expression from Theorem 3.
- (ii) The recurrence diameter can be exponentially smaller than the traversal diameter.

On the other hand, computing the recurrence diameter is an NP-hard problem, while computing the traversal diameter can be done in linear time.

3.5.1 Upper-Bounding the Diameter using Abstractions' Recurrence Diameters

Previous attempts to employ the compositional approach to upper-bound the diameter computed an upper bound on the diameter of the system's abstractions and then combined them into an upper bound on the diameter of the concrete system. For instance, Baumgartner et al. (2002) used the recurrence diameters of abstractions of the system at hand. A similar approach was tried by Rintanen and Gretton (2013), where they computed upper-bounds on the size of the state spaces of the abstractions. Two common features of these approaches are:

- (i) both approaches used algorithms defined recursively “bottom-up” on acyclic dependency graphs (or equivalently, design netlists)
- (ii) the algorithms map every dependency graph to a polynomial, where the terms of the polynomial depend only on the structure of the dependency graph.

In this section we formulate the underlying algorithm of the two previous approaches as a polynomial generator recursively defined on acyclic dependency graphs, which we call M_{sum} . We then introduce a new polynomial generator, N_{sum} , that does the recursion in the opposite direction (i.e. “top-down”), and experimentally show that it significantly dominates the bottom-up approach in terms of tightness. We then show that bounds computed by N_{sum} are the best any function that is a sound bound on the diameter can compute using abstractions' recurrence diameters and dependencies, given that the dependencies are acyclic.

The Bottom-Up Approach

Before we model the bottom-up calculation, we need to define the concepts of *ancestors* and *leaves*.

Definition 14 (Leaves). *We define the set of leaves $\text{leaves}(\mathcal{G}_{VS})$ to contain those vertices of \mathcal{G}_{VS} from which there are no outgoing edges.*

Definition 15 (Ancestors). *We write $\text{ancestors}(vs)$ to denote the set of ancestor vertices of vs in \mathcal{G}_{VS} . It is the set $\{vs_0 \mid vs_0 \in \mathcal{G}_{VS} \wedge vs_0 \rightarrow^+ vs\}$, where \rightarrow^+ is the transitive closure of \rightarrow .*

Definition 16 (Bottom-Up Acyclic Dependency Compositional Bound).

$$\mathbf{M}\langle b \rangle(vs, \delta, \mathcal{G}_{VS}) = b(\delta|_{vs}) + (1 + b(\delta|_{vs})) \sum_{a \in \text{ancestors}(vs)} \mathbf{M}\langle b \rangle(\delta|_a)$$

Then, let $\mathbf{M}_{\text{sum}}\langle b \rangle(\delta, \mathcal{G}_{VS}) = \sum_{vs \in \text{leaves}(\mathcal{G}_{VS})} \mathbf{M}\langle b \rangle(vs, \delta, \mathcal{G}_{VS})$.

The functional parameter b is used to bound abstract subproblems, \mathcal{G}_{VS} is a lifted dependency graph of δ used to identify abstract subproblems, δ is the system of interest. Valid instantiations of b are the recurrence diameter (used in Baumgartner et al. (2002)) and the size of state space $2^{|\mathcal{D}(\delta)|}$ (used in Rintanen and Gretton (2013)). Also, \mathbf{M}_{sum} is recursively defined on the structure of \mathcal{G}_{VS} , and accordingly it is only well-defined if \mathcal{G}_{VS} is a DAG. In Baumgartner et al. (2002), it was shown that $\mathbf{M}_{\text{sum}}\langle rd \rangle$ can be used to upper-bound the diameter, in case δ has an acyclic lifted dependency graph. A restatement of Theorem 1 from Baumgartner et al. (2002) is as follows.

Theorem 7. For any factored representation δ with acyclic lifted dependency graph A_{VS} , $d(\delta) \leq \mathbf{M}_{\text{sum}}\langle rd \rangle(\delta, A_{VS})$.

The following example illustrates the operation of \mathbf{M}_{sum} .

Example 8. Figure 3.8 shows a lifted dependency DAG, A_{VS} , of some factored system δ . Since A_{VS} is a DAG, this implies that for $1 \leq i \leq 4$, the set of variables vs_i is closed under mutual dependency, and $\mathcal{D}(\delta) = \bigcup vs_i$. Given b , and letting b_i be $b(\delta|_{vs_i})$ and $\mathbf{M}\langle b \rangle(vs_i, \delta, A_{VS}) = M_i$, we have

$$(i) \quad M_1 = b_1,$$

$$(ii) \quad M_2 = b_2,$$

$$(iii) \quad M_3 = b_3 + M_1 + b_3 M_1 = b_3 + b_1 + b_1 b_3,$$

$$(iv) \quad M_4 = b_4 + (1 + b_4)(M_1 + M_2 + M_3) \\ = 2b_1 + b_2 + b_3 + b_4 + b_1 b_3 + 2b_1 b_4 + b_2 b_4 + b_3 b_4 + b_1 b_3 b_4.$$

Since vs_4 is the only leaf in the dependency graph, the polynomial evaluated by \mathbf{M}_{sum} will be M_4 .

The previous example should make it clear that \mathbf{M}_{sum} can be viewed as a polynomial generating function *recursively* defined on a DAG. The terms of the polynomial \mathbf{M}_{sum} returns depends *only* on the structure of A_{VS} (i.e. the number of vertices and their connectivity), regardless of δ or the values of b on different projections. However, \mathbf{M}_{sum} has a problem: it repeatedly adds the terms $\mathbf{M}\langle b \rangle(vs_i)$ as many times as there are children for vs_i . Except for the first $\mathbf{M}\langle b \rangle(vs_i)$ term it adds, all those terms are redundant as we show in the next section. We also note that the function \mathbf{M}_{sum} is monotonic: for a bounding function b_1 that is an upper bound on another bounding function b_2 , $\mathbf{M}_{\text{sum}}\langle b_2 \rangle(\delta, A_{VS}) \leq \mathbf{M}_{\text{sum}}\langle b_1 \rangle(\delta, A_{VS})$ holds.

The Top-Down Approach

In this section we define a new polynomial generator. The motivation in defining this polynomial generator is to avoid redundant terms as highlighted in the previous section. We shall also see that it is easy to formally verify it as an upper bound on the diameter.

Definition 17 (Acyclic Dependency Compositional Bound).

$$\mathbf{N}\langle b \rangle(vs, \delta, \mathcal{G}_{VS}) = b(\delta|_{vs})(1 + \sum_{c \in \text{children}_{\mathcal{G}_{VS}}(vs)} \mathbf{N}\langle b \rangle(c, \delta, \mathcal{G}_{VS}))$$

Then, let $\mathbf{N}_{\text{sum}}\langle b \rangle(\delta, \mathcal{G}_{VS}) = \sum_{vs \in \mathcal{G}_{VS}} \mathbf{N}\langle b \rangle(vs, \delta, \mathcal{G}_{VS})$.

\mathbf{N}_{sum} is recursively defined on the structure of \mathcal{G}_{VS} , and accordingly it is only well-defined if \mathcal{G}_{VS} is a DAG. We note that \mathbf{N}_{sum} is monotonic: for a bounding function b_1 that is an upper bound on another bounding function b_2 , $\mathbf{N}_{\text{sum}}\langle b_2 \rangle(\delta, A_{VS}) \leq \mathbf{N}_{\text{sum}}\langle b_1 \rangle(\delta, A_{VS})$ holds, for an acyclic dependency graph A_{VS} .

Theorem 8. For any factored representation δ with an acyclic lifted dependency graph A_{VS} , $d(\delta) \leq \mathbf{N}_{\text{sum}}\langle rd \rangle(\delta, A_{VS})$.

We postpone the proof of this theorem to the next section. However, we now compare the bounds computed by $\mathbf{N}_{\text{sum}}\langle b \rangle$ with the ones computed using $\mathbf{M}_{\text{sum}}\langle b \rangle$.

Example 9. Again we refer to A_{VS} from Figure 3.8. Given a topological property b , and letting $b(\delta|_{vs_i})$ be b_i and $\mathbf{N}\langle b \rangle(vs_i, \delta, A_{VS}) = \mathbf{N}_i$, we have

(i) $\mathbf{N}_4 = b_4$,

(ii) $\mathbf{N}_3 = b_3 + b_3b_4$,

(iii) $\mathbf{N}_2 = b_2 + b_2b_4$,

(iv) $\mathbf{N}_1 = b_1 + b_1\mathbf{N}_3 + b_1\mathbf{N}_4 = b_1 + b_1b_3 + b_1b_3b_4 + b_1b_4$, and the polynomial returned by \mathbf{N}_{sum} is

(v) $\mathbf{N}_{\text{sum}}\langle b \rangle(\delta, A_{VS}) = b_1 + b_2 + b_3 + b_4 + b_1b_3 + b_1b_4 + b_2b_4 + b_3b_4 + b_1b_3b_4$.

The value of $\mathbf{M}_{\text{sum}}\langle b \rangle$ has an extra b_1 term and an extra b_1b_4 term, over that of $\mathbf{N}_{\text{sum}}\langle b \rangle$. This is because $\mathbf{M}_{\text{sum}}\langle b \rangle$ counts every ancestor vertex in the lifted dependency graph as many times as the size of its posterity. We found this phenomenon to be highly prevalent in standard benchmarks, where \mathbf{N}_{sum} substantially outperforms \mathbf{M}_{sum} in terms of the tightness of the computed bounds. Figure 3.9 shows the computed bounds of $\mathbf{N}_{\text{sum}}\langle b \rangle$ versus $\mathbf{M}_{\text{sum}}\langle b \rangle$ with the function $2^{|\mathcal{D}(\delta)|} - 1$ as the base function for a 1030 different International Planning Competition benchmarks. It shows that $\mathbf{M}_{\text{sum}}\langle b \rangle$ computes looser bounds as it repeats counting the ancestor vertices unnecessarily.

Theoretical Guarantees on $\mathbf{N}_{\text{sum}}\langle rd \rangle$ as a Bound on the Diameter

The next theorem shows that for every dependency graph, there is a system where the diameter lower bounds $\mathbf{N}_{\text{sum}}\langle rd \rangle$. This, together with Theorem 8 that states that $\mathbf{N}_{\text{sum}}\langle rd \rangle$ upper-bounds the diameter, show that the top-down algorithm is tight and cannot be improved.

Theorem 9. For any \mathbb{N} -DAG $A_{\mathbb{N}}$, there is a factored system δ with a lifted dependency DAG, A_{VS} , such that (i) $\mathbf{N}_{\text{sum}}\langle rd \rangle(\delta, A_{VS}) \leq d(\delta)$, and (ii) $A_{\mathbb{N}} = \mathfrak{R}(A_{VS})$, where $\mathfrak{R}(vs) = rd(\delta|_{vs})$.

We use an ‘‘inverted flower’’ system, defined as follows.

$$\bar{v}_{l,m}^i = \begin{cases} \bigcup_{1 \leq j < k \leq l} \{(x_j^i, x_k^i), (x_k^i, x_j^i)\} & \text{if } m = 1 \\ \bigcup_{2 \leq j \leq l+1} \{(x_1^i, x_j^i), (x_j^i, x_1^i)\} & \text{if } m = 2 \\ \gamma_{m-2}^i \cup \bigcup_{0 \leq j \leq l-1} \{(x_{m-1}^i, x_{m+j}^i), (x_{m+j}^i, x_1^i)\} & \text{o/w} \end{cases}$$

$\vec{l}_{l,m}^i$ is a system whose state space contains a simple path with $m - 1$ states and l additional states (petals). Each petal has an outgoing edge to the first state in the path and an incoming edge from the last state in the path. For $\vec{l}_{l,m}^i$ we denote the shortest action sequence that joins x_a^i with x_b^i , with $\vec{\pi}_{a \rightarrow b}^i$, for $m \leq a, b \leq m + l - 1$. Since for $m \leq a, b \leq m + l - 1$ the length of $\vec{\pi}_{a \rightarrow b}^i$ will always be m , we have $d(\vec{l}_{l,m}^i) = rd(\vec{l}_{l,m}^i) = m$.

Proof. Our proof is constructive. The first step of is to construct an inverted flower relabel for every $u_1 \in V(A_{\mathbb{N}})$, so we denote by \vec{l}^{u_1} the inverted flower $\vec{l}_{p_{u_1}, A_{\mathbb{N}}(u_1)}^{u_1}$. The relabelling inverted flowers are constructed with two properties in mind. Firstly, the recurrence diameter of an inverted flower label of u_1 matches $A_{\mathbb{N}}(u_1)$, the original label of u_1 . Secondly, the number of petals, p_{u_1} , in the inverted flower labelling a vertex u_1 must be two more than the actions in all of inverted flowers labelling the children of u_1 , i.e. $p_{u_1} = 2 + \sum_{u_2 \in \text{children}_{A_{\mathbb{N}}}(u_1)} |\vec{l}^{u_2}|$. Having that many petals is crucial to the proof: we add a different petal as a precondition in every action in the inverted flowers of the children of a vertex u_1 . This is to necessitate the execution of $A_{\mathbb{N}}(u_1)$ actions from \vec{l}^{u_1} between every action from the inverted lotus associated with every child of u_1 , so that the constructed system has the required diameter.

The next step in this construction is to build the required witness system δ . If we follow the same strategy as the ones in the constructions of Theorem 1 or Theorem 2, δ can be formed by the union of all the inverted flowers and then redundant actions with superfluous preconditions can be added to δ , to give the required dependencies. This would guarantee that δ satisfies condition ii. However, δ has to be constructed such that there is an action sequence $\vec{\pi} \in \delta^*$ that is $N_{\text{sum}}(rd)(\delta, A_{\mathbb{V}\mathbb{S}})$ actions long, and such that there are two states $x_1 \in \mathbb{U}(\delta)$ and $x_2 = \text{ex}(\vec{\pi}, x_1)$, with no action sequence shorter than $\vec{\pi}$ being able to reach x_2 from x_1 . To guarantee this, actions with preconditions added (for the purpose of adding dependencies) must replace the original actions, they cannot be redundant. Additionally, those preconditions have to guarantee that in $\vec{\pi}$, between every pair of actions from an inverted flower associated with a vertex u_1 , there needs to be as many action sequences as there are parents to u_1 , each of which is as long as the recurrence diameter of the corresponding parent's inverted flower. To achieve that δ and $\vec{\pi}$ are constructed by Algorithm 3.

Algorithm 3:

```

 $\vec{\pi} := []$ ;  $\delta := \bigcup_{u_1 \in V(A_{\mathbb{N}})} (\vec{l}^{u_1})$ 
// Loop over  $V(A_{\mathbb{N}})$  in reverse topological order
for  $u_1 \in l_{\mathbb{N}}$ 
   $i := A_{\mathbb{N}}(u_1)$ ; children :=  $\bigcup_{u_2 \in \text{children}_{A_{\mathbb{N}}}(u_1)} \vec{l}^{u_2}$ 
  // Loop over every action in  $\vec{\pi}$ , in execution order
  for  $(p, e) \in \vec{\pi}$ 
    // check if the current action belongs to a child of  $u_1$ 
    if  $\mathcal{D}(e) \cap \mathcal{D}(\text{children}) \neq \emptyset$ 
      // Add preconditions to ensure it needs  $\vec{\pi}_{i \rightarrow i+1}$ 
       $\delta := \delta \setminus \{(p, e)\}$ ;  $\delta := \delta \cup \{(x_{i+1}^{u_1} \uplus p, e)\}$ 
      // Add  $\vec{\pi}_{i \rightarrow i+1}$  before7( $p, e$ );
       $\vec{\pi} := \text{pre}(\vec{\pi}, (p, e)) \frown \vec{\pi}_{i \rightarrow i+1}^v \frown (p, e) :: \text{suf}(\vec{\pi}, (p, e))$ 
      // Increment the petal counter
       $i := i + 1$ 
  // Add an action sequence to visit one more petal
   $\vec{\pi} := \vec{\pi} \frown \vec{\pi}_{i \rightarrow i+1}^{u_1}$ 

```

We now show that δ satisfies requirement i. Consider two states, a state at which every inverted flower is at the first petal $x_1 = \bigsqcup_{u_1 \in V(A_N)} x_{A_N(u_1)}^{u_1}$ and a state at which every inverted flower is at the last petal $x_2 = \bigsqcup_{v \in V(A_{VS})} x_{A_N(u_1)+p_{u_1}+1}^{u_1}$. By construction, $\text{ex}(x_1, \vec{\pi}) = x_2$. Since every action in $\vec{\pi}$ had a different petal added as a precondition from every parent, for any $\vec{\pi}' \in \delta^*$ that is shorter than $\vec{\pi}$, $\text{ex}(x_1, \vec{\pi}') \neq x_2$. Accordingly the diameter of δ is at least $|\vec{\pi}|$. Since for an inverted flower \mathbb{I}^{u_1} , the length of $\vec{\pi}_{a \rightarrow b}^{u_1}$ is $\text{rd}(\mathbb{I}^{u_1}) = A_N(u_1)$, then the length of $\vec{\pi}$ is at least $N_{\text{sum}}(\text{rd})(\delta, A_{VS})$, and δ satisfies conclusion i.

We now show that δ satisfies requirement ii. Consider a relabelling, A_{VS} , of A_N , where every vertex u_1 is relabelled by $\mathcal{D}(\mathbb{I}^{u_1})$, the domain of the inverted flower \mathbb{I}^{u_1} associated with the vertex. Since for a vertex u_1 , actions added from \mathbb{I}^{u_1} to δ had preconditions added to them from the petals of inverted flowers labelling all the parents of u_1 and only the parents, edges of A_{VS} represent dependencies of δ . Accordingly A_{VS} is a lifted dependency DAG of δ . Also since $\delta|_{\mathcal{D}(\mathbb{I}^{u_1})} = \mathbb{I}^{u_1}$, and \mathcal{G}_N is a relabelling of \mathcal{G}_{VS} , then $\mathcal{G}_N = \mathfrak{R}(\mathcal{G}_{VS})$. \square

Example 10. For the \mathbb{N} -DAG shown in Figure 3.4a, a reverse topological ordering of the vertices in this graph is the list $[u_2; u_3; u_1]$, where leaves u_1 and u_3 come first. We first build an inverted flower system to label each of the vertices (Figures 3.11a-3.11c). As u_2 is a leaf, $p_{u_2} = 0$, so we construct an inverted flower $\mathbb{I}_{2,3}^{u_2}$ with two petals and a path with two states. Concretely, $\mathbb{I}_{2,3}^{u_2} = \{(x_1^{u_2}, x_2^{u_2}), (x_2^{u_2}, x_3^{u_2}), (x_2^{u_2}, x_4^{u_2}), (x_3^{u_2}, x_1^{u_2}), (x_4^{u_2}, x_1^{u_2})\}$. The states are $x_1^{u_2} = \{\bar{v}_4, \bar{v}_5\}$, $x_2^{u_2} = \{\bar{v}_4, v_5\}$, $x_3^{u_2} = \{v_4, \bar{v}_5\}$ (the first petal), $x_4^{u_2} = \{v_4, v_5\}$ (the second petal). For u_3 we construct $\mathbb{I}_{2,2}^{u_3} = \{(x_1^{u_3}, x_3^{u_3}), (x_1^{u_3}, x_2^{u_3}), (x_3^{u_3}, x_1^{u_3}), (x_2^{u_3}, x_1^{u_3})\}$. The states are $x_1^{u_3} = \{\bar{v}_6, \bar{v}_7\}$, $x_2^{u_3} = \{\bar{v}_6, v_7\}$ (the first petal), and $x_3^{u_3} = \{v_6, \bar{v}_7\}$ (the second petal). For the root vertex u_1 , $p_{u_1} = 7$, as inverted flowers labelling its children have 5 actions, so $\mathbb{I}_{7,2}^{u_1}$ will have 7 petals. Concretely, $\mathbb{I}_{7,2}^{u_1} = \bigcup \{(x_1^{u_1}, x_j^{u_1}), (x_j^{u_1}, x_1^{u_1}) \mid 2 \leq j \leq 8\}$. The states are $x_1^{u_1} = \{\bar{v}_1, \bar{v}_2, \bar{v}_3\}$, $x_2^{u_1} = \{\bar{v}_1, \bar{v}_2, v_3\}$, $x_3^{u_1} = \{\bar{v}_1, v_2, \bar{v}_3\}$, $x_4^{u_1} = \{\bar{v}_1, v_2, v_3\}$, $x_5^{u_1} = \{v_1, \bar{v}_2, \bar{v}_3\}$, $x_6^{u_1} = \{v_1, \bar{v}_2, v_3\}$, $x_7^{u_1} = \{v_1, v_2, \bar{v}_3\}$, and $x_8^{u_1} = \{v_1, v_2, v_3\}$.

After constructing the inverted flowers we construct the system δ and the action sequence $\vec{\pi}$. Based on Algorithm 3, we start from the leaves, and initially $\vec{\pi} = []$ and $\delta = \mathbb{I}_{2,3}^{u_2} \cup \mathbb{I}_{2,2}^{u_3} \cup \mathbb{I}_{7,2}^{u_1}$. u_2 has no children, i.e. $\text{children} = \emptyset$, so we skip the inner for-loop and concatenate the shortest action sequence that reaches the second petal from the first one in the inverted flower associated with u_2 , so $\vec{\pi} = \vec{\pi}_{3 \rightarrow 4}^{u_2} = [(x_3^{u_2}, x_1^{u_2}); (x_1^{u_2}, x_2^{u_2}); (x_2^{u_2}, x_4^{u_2})]$. Similarly, since u_3 has no children, we concatenate $\vec{\pi}_{1 \rightarrow 2}^{u_3}$ to $\vec{\pi}$, so $\vec{\pi} = \vec{\pi}_{3 \rightarrow 4}^{u_2} \frown \vec{\pi}_{2 \rightarrow 3}^{u_3} = [(x_3^{u_2}, x_1^{u_2}); (x_1^{u_2}, x_2^{u_2}); (x_2^{u_2}, x_4^{u_2}); (x_2^{u_3}, x_1^{u_3}); (x_1^{u_3}, x_3^{u_3})]$.

For u_1 , $\text{children} = \mathbb{I}_{2,3}^{u_2} \cup \mathbb{I}_{2,2}^{u_3}$, which means that the inner for-loop executes, iterating over each action in $\vec{\pi}$. The first action in $\vec{\pi}$, $(x_3^{u_2}, x_1^{u_2})$ comes from the inverted flower labelling u_2 , which is a child of u_1 . Accordingly, we remove $(x_3^{u_2}, x_1^{u_2})$ from δ , add to its precondition the second petal $(x_3^{u_1})$ in the inverted flower labelling u_1 , which makes the action $(x_3^{u_2} \sqcup x_3^{u_1}, x_1^{u_2})$ and add the augmented action to δ . Then we add before that action, the shortest action sequence joining the first and the second petals in the inverted flower of u_1 , so $\vec{\pi} = \vec{\pi}_{2 \rightarrow 3}^{u_1} \frown [(x_3^{u_1} \sqcup x_3^{u_2}, x_1^{u_2}); (x_1^{u_2}, x_2^{u_2}); (x_2^{u_2}, x_4^{u_2}); (x_2^{u_3}, x_1^{u_3}); (x_1^{u_3}, x_3^{u_3})]$. This is repeated for the remaining actions, so $\vec{\pi} = \vec{\pi}_{2 \rightarrow 3}^{u_1} \frown [(x_3^{u_1} \sqcup x_3^{u_2}, x_1^{u_2})] \frown \vec{\pi}_{3 \rightarrow 4}^{u_2} \frown [(x_4^{u_1} \sqcup x_1^{u_2}, x_2^{u_2})] \frown \vec{\pi}_{4 \rightarrow 5}^{u_1} \frown [(x_5^{u_1} \sqcup x_2^{u_2}, x_4^{u_2})] \frown \vec{\pi}_{5 \rightarrow 6}^{u_1} \frown [(x_6^{u_1} \sqcup x_2^{u_3}, x_1^{u_3})] \frown \vec{\pi}_{6 \rightarrow 7}^{u_1} \frown [(x_7^{u_1} \sqcup x_1^{u_3}, x_3^{u_3})]$. All the modified actions were replaced in δ with themselves, but with the added preconditions. After the inner for loop ends, $\vec{\pi}_{7 \rightarrow 8}^{u_1}$ is concatenated, so $\vec{\pi} = \vec{\pi}_{2 \rightarrow 3}^{u_1} \frown [(x_3^{u_1} \sqcup$

⁷For lists l , l_1 , and l_3 , if $l = l_1 \frown a \frown l_2$, then $\text{pre}(l, a)$ denotes l_1 and $\text{suf}(l, a)$ denotes l_2 .

$x_3^{u_2}, x_1^{u_2}) \curvearrowright \vec{\pi}_{3 \rightarrow 4}^{u_1} \curvearrowright [(x_4^{u_1} \uplus x_1^{u_2}, x_2^{u_2})] \curvearrowright \vec{\pi}_{4 \rightarrow 5}^{u_1} \curvearrowright [(x_5^{u_1} \uplus x_2^{u_2}, x_4^{u_2})] \curvearrowright \vec{\pi}_{5 \rightarrow 6}^{u_1} \curvearrowright [(x_6^{u_1} \uplus x_2^{u_3}, x_1^{u_3})] \curvearrowright \vec{\pi}_{6 \rightarrow 7}^{u_1} \curvearrowright [(x_7^{u_1} \uplus x_1^{u_3}, x_3^{u_3})] \curvearrowright \vec{\pi}_{7 \rightarrow 8}^{u_1}$.

Consider the states $x_1 = x_2^{u_1} \uplus x_3^{u_2} \uplus x_2^{u_3} = \{\bar{v}_1, \bar{v}_2, v_3, \bar{v}_4, v_5, \bar{v}_6, v_7\}$ and $x_2 = x_8^{u_1} \uplus x_4^{u_2} \uplus x_3^{u_3} = \{v_1, v_2, \bar{v}_3, \bar{v}_4, v_5, \bar{v}_6, v_7\}$. $\text{ex}(x_1, \vec{\pi}) = x_2$, and the path traversed by $\vec{\pi}$ in $\mathcal{G}(\delta)$ is shown in Figure 3.11e. There is not an action sequence in δ^* shorter than $\vec{\pi}$ (17) that can reach x_2 from x_1 . Accordingly the diameter of δ is at least 17. Letting $\mathbf{N}(vs) = \mathbf{N}\langle rd \rangle(vs, \delta, A_{VS})$, we have $\mathbf{N}(\{v_4, v_5\}) = rd(\delta|_{\{v_4, v_5\}}) = rd(\mathbb{I}_{2,3}^{u_2}) = 3$, $\mathbf{N}(\{v_4, v_5\}) = rd(\delta|_{\{v_6, v_7\}}) = rd(\mathbb{I}_{2,2}^{u_3}) = 2$, $\mathbf{N}(\{v_1, v_2, v_3\}) = rd(\delta|_{\{v_1, v_2, v_3\}})(1 + \mathbf{N}(\{v_4, v_5\}) + \mathbf{N}(\{v_6, v_7\})) = 2(1 + 3 + 2) = 12$, and $\mathbf{N}_{\text{sum}}\langle rd \rangle(\delta, A_{VS}) = \mathbf{N}(\{v_1, v_2, v_3\}) + \mathbf{N}(\{v_4, v_5\}) + \mathbf{N}(\{v_6, v_7\}) = 12 + 3 + 2 = 17$, so $\mathbf{N}_{\text{sum}}\langle rd \rangle(\delta, A_{VS}) \leq d(\delta)$.

The domain of δ is $\{v_1, v_2, v_3, v_4, v_5, v_6, v_7\}$ and a partition of it is $\{\mathcal{D}(\mathbb{I}_{7,2}^{u_1}), \mathcal{D}(\mathbb{I}_{2,3}^{u_2}), \mathcal{D}(\mathbb{I}_{2,2}^{u_3})\} = \{\{v_1, v_2, v_3\}, \{v_4, v_5\}, \{v_6, v_7\}\}$. A lifted dependency graph of δ , A_{VS} , has the following labels: $A_{VS}(u_1) = \{v_1, v_2, v_3\}$, $A_{VS}(u_2) = \{v_4, v_5\}$, and $A_{VS}(u_3) = \{v_6, v_7\}$ (shown in Figure 3.11d). It should be clear that $A_{VS} = \mathcal{D}(A_\Delta)$.

3.6 The Sublist Diameter

In this section we prove the validity of $\mathbf{N}_{\text{sum}}\langle rd \rangle$ as an upper bound on the diameter in the case of an acyclic dependency graph. To do so we define the *sublist diameter*.

Definition 18 (Sublist Diameter). Recall that a list l' is a sublist of l , written $l' \preceq l$, iff all the members of l' occur in the same order in l . The sublist diameter, $\ell(\delta)$, is the length of the longest shortest equivalent sublist to any execution $\vec{\pi} \in \delta^*$ starting at any state $x \in \mathbb{U}(\delta)$. Formally,

$$\ell(\delta) = \max_{x \in \mathbb{U}(\delta)} \max_{\vec{\pi} \in \delta^*} \min\{|\vec{\pi}'| \mid \vec{\pi}' \preceq \vec{\pi} \wedge \text{ex}(x, \vec{\pi}') = \text{ex}(x, \vec{\pi})\}$$

It should be clear that the sublist diameter is an upper bound on the diameter and a lower bound on the recurrence diameter. We prove that it can be compositionally bounded by $\mathbf{N}_{\text{sum}}\langle \ell \rangle$, Theorem 8 then follows. We also investigate the usage of the sublist diameter to compositionally upper-bound the diameter instead of using the recurrence diameter (i.e. using $\mathbf{N}_{\text{sum}}\langle \ell \rangle$ instead of $\mathbf{N}_{\text{sum}}\langle rd \rangle$). This can be very advantageous, since we show that the sublist diameter is equal to the diameter in infinitely many systems, where the recurrence diameter is exponentially looser than the diameter (Theorem 10). However, computing the sublist the diameter is at least as hard as computing the recurrence diameter, which is NP-hard (Theorem 12).

3.6.1 Compositionally Bounding the Sublist Diameter Under Acyclic Dependency

Lemma 3. For any factored system δ with an acyclic lifted dependency DAG A_{VS} , $\ell(\delta) \leq \mathbf{N}_{\text{sum}}\langle \ell \rangle(\delta, A_{VS})$.

Proof Summary of Lemma 3. To prove Lemma 3 we use a construction which, given any action sequence $\vec{\pi} \in \delta^*$ violating the stated bound and a state $s \in \mathbb{U}(\delta)$, produces a sublist, $\vec{\pi}'$, of $\vec{\pi}$ satisfying that bound and $\text{ex}(s, \vec{\pi}) = \text{ex}(s, \vec{\pi}')$. The proof is by induction on the structure of A_{VS} , where we repeatedly consider every set of variables $P \in A_{VS}$ and each of its children $C \in \text{children}_{A_{VS}}(P)$, and for each parent child pair we do the following. We first consider the projected action sequence $\vec{\pi}|_C$. From the definition of ℓ , there is a sublist $\vec{\pi}'|_C \preceq \vec{\pi}|_C$ satisfying $|\vec{\pi}'|_C| \leq \ell(\delta|_C)$. Moreover, this guarantees that $\vec{\pi}'|_C$ is equivalent, in terms of the execution outcome, to

$\vec{\pi} \downarrow_C$. Hereupon, if for an action $\pi = (p, e)$, $e \subseteq vs$ is true, we call π a vs -action. The stitching function described in Figure 3.12a is then used to remove the C -actions in $\vec{\pi}$ whose projections on C are not in $\vec{\pi}'_C$. Thus our construction arrives at the action sequence $\vec{\pi}'' = \vec{\pi}'_C \uparrow \vec{\pi}$ with at most $\ell(\delta \downarrow_C)$ C -actions. Since $C \in \text{children}_{A_{VS}}(P)$, then $C \not\rightarrow P$ and accordingly, actions with variables from C in their effects never include P variables in their effects. Accordingly, $\text{ex}(s, \vec{\pi}) = \text{ex}(s, \vec{\pi}'')$, and in $\vec{\pi}''$, there will be a list, $\vec{\pi}_P$, of only P -actions between every consecutive C -actions. However, again, because of how ℓ is defined, there is a $\vec{\pi}'_P \preceq \vec{\pi}_P$, whose execution outcome is the same as $\vec{\pi}_P$ and whose length is no more than $\ell(\delta \downarrow_P)$. Repeating this for all sequences of only P -actions completes our construction. \square

The above construction and the usage of the stitching function is illustrated in the following example.

Example 11. Consider the following factored system, whose dependency graph is shown in Figure 3.12b.

$$\left\{ \begin{array}{l} \pi_1 = (\emptyset, \{v_3\}), \pi_2 = (\{v_3\}, \{v_4\}), \pi_3 = (\{v_3\}, \{\bar{v}_1\}), \pi_4 = (\{v_3\}, \{v_2\}), \\ \pi_5 = (\{v_4\}, \{v_1\}), \pi_6 = (\{v_2, v_4\}, \{v_5\}), \pi_7 = (\{\bar{v}_3\}, \{v_4, v_5\}) \end{array} \right\}$$

$\mathcal{D}(\delta) = c \cup p$, where $c = \{v_1, v_4, v_5\}$ are called the “child” variables, and $p = \{v_2, v_3\}$, and $c \not\rightarrow p$ are called the “parent” variables. In δ , the actions $\pi_2, \pi_3, \pi_5, \pi_6, \pi_7$ are c -actions, and π_1, π_4 are p -actions. An action sequence $\vec{\pi} \in \delta^*$ is $[\pi_1; \pi_1; \pi_2; \pi_3; \pi_4; \pi_4; \pi_5; \pi_6]$ that reaches the state $\{v_1, v_2, v_3, v_4, v_5\}$ from $\{v_1, \bar{v}_2, \bar{v}_3, \bar{v}_4, \bar{v}_5\}$. When $\vec{\pi}$ is projected on c it becomes $[\pi_2 \downarrow_c; \pi_3 \downarrow_c; \pi_5 \downarrow_c; \pi_6 \downarrow_c]$, which is in $(\delta \downarrow_c)$. A shorter action sequence, $\vec{\pi}_c$, achieving the same result as $\vec{\pi} \downarrow_c$ is $[\pi_2 \downarrow_c; \pi_6 \downarrow_c]$. Since $\vec{\pi}_c$ is a sublist of $\vec{\pi} \downarrow_c$, we can use the stitching function to obtain a shorter action sequence in δ^* that reaches the same state as $\vec{\pi}$. In this case, $\vec{\pi}_c \uparrow \vec{\pi}$ is $[\pi_1; \pi_1; \pi_2; \pi_4; \pi_4; \pi_6]$. The second step is to contract the pure p segments which are $[\pi_1; \pi_1]$ and $[\pi_4; \pi_4]$, which are contracted to $[\pi_1]$ and $[\pi_4]$, respectively. The final constructed action sequence is $[\pi_1; \pi_2; \pi_4; \pi_6]$, which achieves the same state as $\vec{\pi}$.

Theorem 8 follows directly from Lemma 3 and the fact that the sublist diameter is both: an upper bound on the diameter and a lower bound on the recurrence diameter.

Also because the recurrence diameter upper bounds the sublist diameter, and the sublist diameter upper-bounds the diameter, and since N_{sum} is monotonic, the following lemma can be derived from Theorem 9 and Lemma 3.

Lemma 4. For any \mathbb{N} -DAG (a DAG whose vertices are labelled with numbers), $A_{\mathbb{N}}$, there is a δ that has a lifted dependency DAG, A_{VS} , such that (i) $D_{\mathbb{N}} = \mathfrak{L}(A_{VS})$, where $\mathfrak{L}(vs) = \ell(\delta \downarrow_{vs})$ (i.e. $D_{\mathbb{N}}$ is a relabelling of A_{VS} , where every vertex labelled by vs in A_{VS} , is relabelled by the sublist diameter of the projection $\delta \downarrow_{vs}$), and (ii) $\ell(\delta) = N_{\text{sum}}(\ell)(\delta, A_{VS})$.

3.6.2 The Sublist Diameter as a Compositional Upper Bound on the Diameter

Initially we defined the sublist diameter as a tool for proving Theorem 8. However, Lemma 3 shows that we can use the sublist diameter (instead of the recurrence diameter) as a base case function with the top-down approach to upper-bound the diameter. In this section we compare the sublist diameter to the recurrence diameter as a base case function for upper-bounding the diameter. We compare two aspects: their tightness as bounds, and the hardness of computing each of them.

Tightness of the Sublist Diameter

In this section we show that in addition to being a lower bound on the recurrence diameter, the sublist diameter can be exponentially smaller than the recurrence diameter. To do that, we exploit a rather interesting insight: the value of sublist diameter depends on the factored representation: two factored representations of the same state space can have different sublist diameters. Indeed minimising (or *factoring*) the size of a factored representation (i.e. the number of actions) used to represent the same system can minimise the sublist diameter. This is in contrast to the diameter, the recurrence diameter and the traversal diameter, which are agnostic to the factored representation: they only depend on the state space of the system. The following example demonstrates using factoring to reduce the sublist diameter.

Example 12. Consider $\delta_{\{v_1, v_2\}}$ from Example 2 and $\delta_1 = \{k_1 \downarrow_{\{v_1, v_2\}}, k_2 \downarrow_{\{v_1, v_2\}}, k_3 \downarrow_{\{v_1, v_2\}}, k_4 \downarrow_{\{v_1, v_2\}}\}$. δ_1 is a factoring of $\delta_{\{v_1, v_2\}}$ and accordingly, $\mathcal{G}(\delta_1)$ and $\mathcal{G}(\delta_{\{v_1, v_2\}})$ are identical. Accordingly $d(\delta_1) = d(\delta_{\{v_1, v_2\}}) = 1$, and $rd(\delta_1) = rd(\delta_{\{v_1, v_2\}}) = 3$. $\ell(\delta_1) = 1$, because for any non empty action sequence $\vec{\pi} \in \delta_1^*$, the last action π in $\vec{\pi}$ reaches the same state as $\vec{\pi}$, and $[\pi] \preceq \vec{\pi}$. In contrast, $\ell(\delta_{\{v_1, v_2\}}) = 3$, since no shorter sublist of the action sequence $[p_1 \downarrow_{\{v_1, v_2\}}; p_2 \downarrow_{\{v_1, v_2\}}; p_3 \downarrow_{\{v_1, v_2\}}]$ starts at $\{\bar{v}_1, \bar{v}_2\}$ and results in $\{v_1, v_2\}$.

We now formally demonstrate the utility of factoring to reduce the sublist diameter.

Definition 19 (Factoring). Factored representation δ_1 is a factoring of representation δ_2 iff all the following is true:

- (i) $\mathcal{D}(\delta_1) \subseteq \mathcal{D}(\delta_2)$
- (ii) $|\delta_1| \leq |\delta_2|$
- (iii) $\forall \pi_2 \in \delta_2, x \in \mathbb{U}(\delta_2). \exists \pi_1 \in \delta_1. \text{ex}(x, \pi_1) = \text{ex}(x, \pi_2)$
- (iv) there is a function $f : \delta_1 \Rightarrow 2^{\delta_2}$, where
 - (a) $\forall (p_1, e_1) \in \delta_1, (p_2, e_2) \in f(p_1, e_1). p_1 \subseteq p_2$,
 - (b) $\forall \pi_1 \in \delta_1, x \in \mathbb{U}(\delta_2). \exists \pi_2 \in f(\pi_1). \text{ex}(x, \pi_1) = \text{ex}(x, \pi_2)$, and
 - (c) $\forall \pi_1 \in \delta_1. f(\pi_1) \neq \emptyset$

Proposition 4. If δ_1 is a factoring of δ_2 then the largest connected components of $\mathcal{G}(\delta_1)$ and $\mathcal{G}(\delta_2)$ are isomorphic.

Corollary 1. If δ_1 is a factoring of δ_2 then $d(\delta_1) = d(\delta_2)$ and $rd(\delta_1) = rd(\delta_2)$.

Lemma 5. There are infinitely many factored systems that have factorings with exponentially smaller sublist diameters, that are equal to the diameter.

Before we begin our proof, let $K_n^i = \{(\emptyset, x_j^i) \mid 1 \leq j \leq n + 1\}$, i.e. a factored system with the only non-singleton component of its state space being a clique with $n + 1$ states.

Proof. We show that for every $n \in \mathbb{N}$ there is a system whose sublist diameter can be reduced from n to 1 through factoring. Consider the factored representations $K = K_n^1$, and $\gamma = \gamma_n^1$. $\ell(K) = 1$ and $\ell(\gamma) = n$ hold. Note that since in the definition of K and γ we used the same superscript, $\mathcal{D}(K) = \mathcal{D}(\gamma) = \text{vs}(\cdot)$. Accordingly the state space of $K \cup \gamma$ will have only one non-singleton component that is a clique of order $n + 1$. However, although $K \cup \gamma$ is a factored system whose

state space is a clique, its sublist diameter will be n , because there is not a shorter sublist of the action sequence $[(x_1^1, x_2^1); (x_2^1, x_3^1); \dots; (x_n^1, x_{n+1}^1)]$ that can reach the state x_{n+1}^1 , starting from the state x_1^1 . Finally, from (*) and since γ does not add any more edges in the state space of K , then K is a factoring of $K \cup \gamma$, which finishes our proof. \square

Theorem 10. *There are infinitely many factored systems whose sublist diameter is (i) equal to their diameter, and (ii) exponentially smaller (in the number of state variables) than their recurrence diameter. There are also infinitely many factored systems whose sublist diameter is (i) exponentially larger (in the number of state variables) than their diameter, and (ii) equal to their recurrence diameter.*

This theorem follows immediately from Corollary 1 and Lemma 5. Furthermore we show that factoring a representation can only reduce the value of the sublist diameter.

Theorem 11. *If δ_1 is a factoring of δ_2 then $\ell(\delta_1) \leq \ell(\delta_2)$.*

Proof. We prove the theorem by showing that for any $x \in \mathbb{U}(\delta_1)$ and $\vec{\pi} \in \delta_1^*$ there is a sublist of $\vec{\pi}$, $\vec{\pi}'$, such that $\text{ex}(s, \vec{\pi}') = \text{ex}(s, \vec{\pi})$ and $|\vec{\pi}'| \leq \ell(\delta_2)$. Using the mapping f from Definition 19, there must be an action sequence $\vec{\pi}_2 \in \delta_2^*$, such that $\text{ex}(x, \vec{\pi}_2) = \text{ex}(x, \vec{\pi})$. From Definition 18 we then obtain a sublist, $\vec{\pi}'_2$, of $\vec{\pi}_2$ whose length is bounded by $\ell(\delta_2)$ and, such that $\text{ex}(x, \vec{\pi}'_2) = \text{ex}(x, \vec{\pi}_2)$. We reverse the mapping of actions to transform $\vec{\pi}'_2$ into a sublist $\vec{\pi}'$ of $\vec{\pi}$. From Definition 19, $\text{ex}(x, \vec{\pi}) = \text{ex}(x, \vec{\pi}')$, but $|\vec{\pi}'| = |\vec{\pi}'_2| \leq \ell(\delta_2)$, which finishes our proof. \square

The Complexity of Computing the Sublist Diameter

Theorem 12. *Computing the Sublist Diameter is NP-hard.*

Proof. We prove this theorem by showing that for any digraph \mathcal{G} , one can compute the length of the longest path in \mathcal{G} by computing the sublist diameter of a system, δ , that can be constructed in polynomial time from \mathcal{G} . Let $\mathcal{D} = \{v_1, v_2, \dots, v_{\lceil \log_2 |V(\mathcal{G})| \rceil}\}$ and let $2^{\mathcal{D}}$ denote the set of all states representable by \mathcal{D} . Since $|V(\mathcal{G})| \leq 2^{|\mathcal{D}|}$, then there is an injective function $x : V(\mathcal{G}) \Rightarrow 2^{\mathcal{D}}$. Consider the system $\delta = \{(x(u_1), x(u_2)) \mid (u_1, u_2) \in E(\mathcal{G})\}$. It should be clear that the largest connected component in $\mathcal{G}(\delta)$ is isomorphic to \mathcal{G} and accordingly the length of the longest path in \mathcal{G} is the same as the length of the longest path in $\mathcal{G}(\delta)$ which is equal to $rd(\delta)$. Now observe that for every $(p, e) \in \delta$, the precondition and the effect are both complete assignments of \mathcal{D} , i.e. $p \in \mathbb{U}(\delta)$ and $e \in \mathbb{U}(\delta)$. Accordingly every action in δ represents exactly one edge from \mathcal{G} . From that we have $rd(\delta) = \ell(\delta)$, which finishes our proof. \square

3.7 Exploiting State Space Acyclicity

The practical utility of dependency graph based decompositions (like N_{sum}) provides a good motivation to pursue other structures, like state space acyclicity. State space acyclicity is independent of acyclicity in variable dependency. Thus, methods previously developed cannot be used to exploit the former in compositional upper-bounding. We demonstrate this using the well-studied hotel key protocol as a case-study.

3.7.1 Hotel Key Protocol

We now consider the hotel key protocol from Jackson (2006). Reasoning about safe and unsafe versions of this protocol is challenging for state-of-the-art AI planners and model-checkers. For example, a version of the protocol was shown unsafe for an instance with 1 room, 2 guests and 4 keys using a counterexample generator in Blanchette and Nipkow (2010). The problem becomes more challenging for the safe version of the protocol, where the only feasible approach is using interactive theorem provers, as in Nipkow (2006).

We describe the factored transition system corresponding to that protocol. The system models a hotel with R rooms, G guests, and K keys per room, which guests can use to enter rooms (Figure 3.13 shows an example with $R = 2$, $G = 2$ and $K = 3$). The state characterising propositions are: (i) $\mathbb{1}k_{r,k}$, reception last issued key k for room r , for $0 < r \leq R$ and $(r - 1)K < k \leq rK$; (ii) $\mathbb{c}k_{r,k}$, room r can be accessed using key k , for $0 < r \leq R$ and $(r - 1)k < k \leq rK$; (iii) $\mathbb{g}k_{g,k}$, guest g has key k , for $0 < g \leq G$, $0 < k \leq RK$; and (iv) \mathbb{s}_r , is an auxiliary variable that means that room r is “safely” delivered to some guest. The protocol actions are as follows: (i) guest g can check-in to room r , receiving key k —($\{\mathbb{1}k_{r,k_1}\}, \{\mathbb{g}k_{g,k_2}, \mathbb{1}k_{r,k_2}, \overline{\mathbb{1}k_{r,k_1}}, \overline{\mathbb{s}_r}\}$); and (ii) where room r was previously entered using key k , guest g can enter room r using key k' —($\{\mathbb{g}k_{g,k'}, \mathbb{1}k_{r,k}\}, \{\mathbb{c}k_{r,k'}, \overline{\mathbb{c}k_{r,k}}, \mathbb{s}_r\}$). Thus, guests can retain keys indefinitely, and there is no direct communication between rooms and reception.

For completeness, we note that this protocol was formulated in the context of checking safety properties. Safety is violated only if a guest enters a room occupied by another guest. Formally, the safety of this protocol is checked by querying if there exists a room r , guest g and keys $k \neq k'$, so that $\mathbb{1}k_{r,k'} \wedge \mathbb{c}k_{r,k} \wedge \mathbb{g}k_{g,k'} \wedge \mathbb{s}_r$. The initial state asserts that guests possess no keys, and the reception issued the first key for each room, and each room opens with its first key. Formally, this is represented by asserting $\mathbb{1}k_{r,(r-1)K} \wedge \mathbb{c}k_{r,(r-1)K}$ is true for $1 \leq r \leq R$, $(r - 1)K < k \leq rK$, and that all other state variables are false.

We adopt some shorthand notations in order to provide examples of concepts in terms of the hotel key protocol. A variable name is written in upper case to refer to a particular assignment, where the only variable that is true is given by the indices. For example, the assignment $\{\overline{\mathbb{c}k_{1,1}}, \mathbb{c}k_{1,2}, \overline{\mathbb{c}k_{1,3}}\}$ —indicating room 1 can be accessed using key 2—is indicated by writing $\mathbb{C}K_{1,2}$. We refer to sets of variables by omitting an index term. For example, $\mathbb{1}k_1$ indicates the variables $\{\mathbb{1}k_{1,i} \mid 1 \leq i \leq 3\}$. The following examples illustrate the concepts of projection and state variable dependency in the context of the hotel key protocol.

Example 13. Consider the set of variables $\text{ROOM1} \equiv \mathbb{1}k_1 \cup \mathbb{c}k_1 \cup \{\mathbb{g}k_{1,2}, \mathbb{g}k_{1,3}, \mathbb{g}k_{2,2}, \mathbb{g}k_{2,3}\}$. The variables ROOM1 model system state relevant to the 1st hotel room. Figure 3.13c shows the projected system $\delta|_{\text{ROOM1}}$.

Example 14. Figure 3.13b shows a dependency graph associated with the system from Figure 3.13a. Let $\text{ROOM2} \equiv \mathbb{1}k_2 \cup \mathbb{c}k_2 \cup \{\mathbb{g}k_{1,5}, \mathbb{g}k_{1,6}, \mathbb{g}k_{2,5}, \mathbb{g}k_{2,6}\}$. Figure 3.13b depicts two connected components induced by the sets ROOM1 and ROOM2 , respectively. One lifted dependency graph would have exactly two unconnected vertices, one being a contraction of the vertices from ROOM1 , and the other a contraction of those from ROOM2 . Due to the disconnected structure of the dependency graph, intuitively the sum of bounds for $\delta|_{\text{ROOM1}}$ and $\delta|_{\text{ROOM2}}$ can be used to upper-bound the diameter of the concrete system.

3.7.2 State Space Acyclicity Compositional Bounding Constructs

To exploit state space acyclicity we formalise it as follows.

Definition 20 (Acyclic Transition System). δ is acyclic iff $\forall x, x' \in \mathbb{U}(\delta)$. $x \neq x'$ then $x \not\rightarrow x'$ or $x' \not\rightarrow x$.

In the next example we show that state space acyclicity is independent of acyclicity in variable dependency, and thus N_{sum} or other methods cannot be used to exploit state space acyclicity for compositional upper-bounding.

Example 15. $\delta \downarrow_{cK_1}$ is acyclic. For example, no state satisfying $CK_{1,2}$ can be reached from a state satisfying $CK_{1,3}$. Now consider $\delta \downarrow_{\text{ROOM1}}$ from Example 13. The dependency graph of $\delta \downarrow_{\text{ROOM1}}$ is comprised of one strongly connected component (SCC). Thus, acyclicity in the assignments of cK_1 cannot be exploited in $\delta \downarrow_{\text{ROOM1}}$ by analysing its dependency graph.

To be able to exploit state space acyclicity, we now introduce a new abstraction concept: *snapshot*. A snapshot models the system when we fix the assignment to a subset of the state variables, removing actions whose preconditions or effects contradict that assignment.

Definition 21 (Snapshot). We write $|X|$ to denote the cardinality of the set X . For states x and x' , let $\text{agree}(x, x')$ denote $|\mathcal{D}(x) \cap \mathcal{D}(x')| = |x \cap x'|$, i.e. a variable that is in the domains of both x and x' has the same assignment in x and x' . For δ and a state x , the snapshot of δ at x is

$$\delta \downarrow_x \equiv \{(p, e) \mid (p, e) \in \delta \wedge \text{agree}(p, x) \wedge \text{agree}(e, x)\} \downarrow_{\mathcal{D}(\delta) \setminus \mathcal{D}(x)}$$

Example 16. $\delta \downarrow_{\text{ROOM1}} \downarrow_{cK_{1,2}}$ is shown in Figure 3.13d.

We now investigate how such acyclicity can be used for bounding. First, let $\Sigma(x) = \delta \downarrow_x$, i.e. it is a function that maps a state x to the snapshot of δ on that state. Consider a system δ where for some variables vs we have that $\delta \downarrow_{vs}$ is acyclic – i.e. the state space of $\delta \downarrow_{vs}$ forms a directed acyclic graph. In that case, we define the following version of S_{max} that operates on acyclic abstraction state spaces.

Definition 22 (Acyclic System Compositional Bound). For an arbitrary bounding function b and for a state $x \in \mathbb{U}(\delta \downarrow_{vs})$,

$$\mathbf{S}\langle b \rangle(x, vs, \delta) = \mathbf{S}\langle b \rangle(x, \Sigma(|\mathcal{G}(\delta \downarrow_{vs})|))$$

Then, let $\mathbf{S}_{\text{max}}\langle b \rangle(vs, \delta) = \max_{x \in \mathbb{U}(\delta \downarrow_{vs})} \mathbf{S}\langle b \rangle(x, vs, \delta)$.

This function returns the longest path in the state space of the projection $\delta \downarrow_{vs}$ weighted by the diameters of the snapshots of δ on different states in $\mathbb{U}(\delta)$. To make this concrete, consider the following example.

Example 17. Since $\delta \downarrow_{cK_1}$ is acyclic, and $CK_{1,i} \in \mathbb{U}(\delta \downarrow_{cK_1})$, then $\mathbf{S}\langle d \rangle(CK_{1,i}, cK_1, \delta)$ is well-defined, for $i \in \{1, 2, 3\}$. Denoting $d(\delta \downarrow_{cK_{1,i}})$ with $d_{1,i}$ and $\mathbf{S}\langle d \rangle(CK_{1,i}, cK_1, \delta)$ with $\mathbf{S}_{1,i}$, we have $\mathbf{S}_{1,3} = d_{1,3}$ because $\text{succ}(CK_{1,3}, \delta \downarrow_{cK_1}) = \emptyset$. We also have $\mathbf{S}_{1,2} = d_{1,2} + 1 + \mathbf{S}_{1,3} = d_{1,2} + 1 + d_{1,3}$ and $\mathbf{S}_{1,1} = d_{1,1} + 1 + \mathbf{S}_{1,2} = d_{1,1} + 1 + d_{1,2} + 1 + d_{1,3} = d_{1,1} + d_{1,2} + d_{1,3} + 2$.

Theorem 13. If $\delta \downarrow_{vs}$ is acyclic and b bounds d , then $d(\delta) \leq \mathbf{S}_{\text{max}}\langle b \rangle(vs, \delta)$.

Proposition 5. If $\delta \downarrow_{vs}$ is acyclic and $x \in \mathbb{U}(\delta \downarrow_{vs})$, if $x' \in \text{succ}(x, \delta \downarrow_{vs})$, then $b(\delta \downarrow_x) + 1 + \mathbf{S}\langle b \rangle(x', vs, \delta) \leq \mathbf{S}\langle b \rangle(x, vs, \delta)$, for a base case function b .

Definition 23 (Subsystem Trace). For a state x , action sequence $\vec{\pi}$, and set of variables vs , let $\partial(x, \vec{\pi}, vs)$ be:

$$\begin{aligned} \partial(x, [], vs) &= [] \\ \partial(x, \pi :: \vec{\pi}, vs) &= \begin{cases} x' :: \partial(x', \vec{\pi}, vs) & \text{if } x \downarrow_{vs} \neq x' \downarrow_{vs} \\ \partial(x', \vec{\pi}, vs) & \text{otherwise} \end{cases} \end{aligned}$$

where $x' = \text{ex}(x, \pi)$.

Proposition 6. For any $x, \vec{\pi}$, and vs , if $\partial(x, \vec{\pi}, vs) = []$ then: (i) $x \downarrow_{vs} = \text{ex}(x, \vec{\pi}) \downarrow_{vs}$ and (ii) there is $\vec{\pi}'$ where $\text{ex}(x, \vec{\pi}) = \text{ex}(x, \vec{\pi}')$ and $|\vec{\pi}'| \leq d(\delta \downarrow_{x \downarrow_{vs}})$.

Proposition 7. For two states x and x' , a sequence of states \vec{x} , a set of variables vs , and an action sequence $\vec{\pi}$, if $\partial(x, \vec{\pi}, vs) = x' :: \vec{x}$, then there are $\vec{\pi}_1, \pi$ and $\vec{\pi}_2$ such that (i) $\vec{\pi} = \vec{\pi}_1 \frown \pi :: \vec{\pi}_2$, (ii) $\partial(x, \vec{\pi}_1, vs) = []$, (iii) $\text{ex}(\text{ex}(x, \vec{\pi}_1), \pi) = x'$, and (iv) $\text{ex}(x', \vec{\pi}_2) = \text{ex}(x, \vec{\pi})$.

Proposition 8. For any $x, \vec{\pi}_1, \vec{\pi}_2$, and vs , we have that $\partial(x, \vec{\pi}_1 \frown \vec{\pi}_2, vs) = \partial(x, \vec{\pi}_1, vs) \frown \partial(\text{ex}(x, \vec{\pi}_1), \vec{\pi}_2, vs)$.

Lemma 6. For any δ and vs where $\delta \downarrow_{vs}$ is acyclic, $s \in \mathbb{U}(\delta)$, and $\vec{\pi} \in \delta^*$, there is $\vec{\pi}'$ such that $\text{ex}(s, \vec{\pi}) = \text{ex}(s, \vec{\pi}')$ and $|\vec{\pi}'| \leq \mathbf{S}\langle d \rangle(s \downarrow_{vs}, vs, \delta)$.⁸

Proof. The proof is by induction on $\partial(s, \vec{\pi})$. The base case, where, $\partial(s, \vec{\pi}) = []$, is trivial. In the step case we have that $\partial(s, \vec{\pi}) = s' :: \vec{x}$ and the induction hypothesis: for any $s^* \in \mathbb{U}(\delta)$, and $\vec{\pi}^* \in \delta^*$ if $\partial(s^*, \vec{\pi}^*) = \vec{x}$ then there is $\vec{\pi}^{*'}$ where $\text{ex}(s^*, \vec{\pi}^*) = \text{ex}(s^*, \vec{\pi}^{*'})$ and $|\vec{\pi}^{*'}| \leq \mathbf{S}\langle d \rangle(s^* \downarrow_{vs})$.

Since $\partial(s, \vec{\pi}) = s' :: \vec{x}$, we have $\vec{\pi}_1, \pi$ and $\vec{\pi}_2$ satisfying the conclusions of Proposition 7. Based on conclusion i, ii, and iii of Proposition 7 and Proposition 8 we have $\partial(s', \vec{\pi}_2) = \vec{x}$. Accordingly, letting s^* , and $\vec{\pi}^*$ from the inductive hypothesis be s' , and $\vec{\pi}_2$, respectively, there is $\vec{\pi}_2'$ such that $\text{ex}(s', \vec{\pi}_2) = \text{ex}(s, \vec{\pi}_2)$ and $|\vec{\pi}_2'| \leq \mathbf{S}\langle d \rangle(s' \downarrow_{vs})$.[†]

From conclusion ii of Proposition 7 and conclusion ii of Proposition 6 there is $\vec{\pi}_1'$ where $\text{ex}(s, \vec{\pi}_1) = \text{ex}(s, \vec{\pi}_1')$ and $|\vec{\pi}_1'| \leq d(\delta \downarrow_{s \downarrow_{vs}})$. Letting $\vec{\pi}' = \vec{\pi}_1' \# \pi :: \vec{\pi}_2'$, from conclusions iii and iv of Proposition 7 and [†] we have $\text{ex}(s, \vec{\pi}) = \text{ex}(s, \vec{\pi}')$ and $|\vec{\pi}'| \leq d(\delta \downarrow_{s \downarrow_{vs}}) + 1 + \mathbf{S}\langle d \rangle(s' \downarrow_{vs})$.[‡]

Lastly, from conclusion i of Proposition 6 and conclusion ii of Proposition 7 we have $s \downarrow_{vs} = \text{ex}(s, \vec{\pi}_1) \downarrow_{vs} = \text{ex}(s, \vec{\pi}_1') \downarrow_{vs}$ and accordingly $\text{ex}(s \downarrow_{vs}, \pi \downarrow_{vs}) = s' \downarrow_{vs}$. Based on that we have $s' \downarrow_{vs} \in \text{children}_{\delta \downarrow_{vs}}(s \downarrow_{vs})$. Then from Proposition 5 and [‡] we have $|\vec{\pi}'| \leq \mathbf{S}\langle d \rangle(s \downarrow_{vs})$. \square

Theorem 13 follows from Lemma 6 and Definitions 6 and 13.

In closing, it is worth noting that for the above example, the dependency graph of $\delta \downarrow_{\text{ROOM1}}$ is comprised of one SCC. Therefore there is no lifted dependency graph providing further decomposition. Indeed, exploiting acyclicity in dependency between variables alone, one cannot further decompose the subproblem $\delta \downarrow_{\text{ROOM1}}$. We were able to achieve a more fine grained decomposition of that component above, by exploiting state space acyclicity.

3.8 A Practical Algorithm for Upper-Bounding

Theorem 13 suggests the possibility of compositional upper-bounding of the diameter given the presence of acyclicity in a transition system's state space. In this section we investigate building a practical compositional algorithm to upper-bound the diameter based on Theorem 13. One straightforward approach is given by Algorithm 4.

In PUR, Ω is an oracle that returns a set of strict subsets of $\mathcal{D}(\delta)$, where $\forall vs \in \Omega(\delta). \delta \downarrow_{vs}$ is acyclic. PUR terminates because a snapshot has fewer variables than the concrete system. In PUR the function UPBND provides upper bounds for the diameters of “base-case” problems – i.e.

⁸In the rest of this proof we omit the vs and/or δ arguments from $\partial(, ,)$ and \mathbf{S} as they do not change.

Algorithm 4: PUR(δ)

$$S = \min(\{\mathbf{S}_{\max}(\text{PUR})(vs, \delta) \mid vs \in \Omega(\delta)\} \cup \infty)$$
if $S = \infty$ **return** UPBND(δ) **else return** S

problems that are not further decomposed. Given this assumption and Theorem 13, PUR itself computes valid upper bounds for the diameter of the whole problem.

A main question for a practical implementation of PUR is the choice of Ω . The trivial choice of all strict subsets of $\mathcal{D}(\delta)$ is impractical. A pragmatic solution which we have adopted, is to take the situation that elements in $\mathcal{D}(\delta)$ model individual assignments in the SAS+ model generated using Fast-Downward's preprocessing step Helmert (2006a). Each element in $\Omega(\delta)$ then corresponds to a set of elements from $\mathcal{D}(\delta)$ that model one multi-valued state variable whose domain transition graph is acyclic.

A source of intractability in PUR comes from the \min operator. For a full evaluation, \mathbf{S}_{\max} is recursively called as many as $|\Omega(\delta)|!$ times. In practice we only evaluate \mathbf{S}_{\max} on one *arbitrarily* chosen element from $\Omega(\delta)$. Our experimentation never uncovered a problem where a full evaluation of the \min , where computationally feasible, produced a better bound. A second source of computational expense comes from the definition of \mathbf{S}_{\max} : PUR can be recursively called a number of times that is linear in the size of the state space of δ . This happens if $\Omega(\delta)$ is a partition of $\mathcal{D}(\delta)$. Although this worst case scenario is contrived, in practice $\Omega(\delta)$ can cover sufficient elements from $\mathcal{D}(\delta)$ to render PUR impractical. This is demonstrated in the following example.

Example 18. Taking $\mathcal{D}(\delta)$ associated with the hotel key protocol example, the Fast-Downward preprocessor Helmert (2006a) identifies partition:

$$\left\{ \begin{array}{l} ck_1, ck_2, lk_1, lk_2, \\ \{gk_{1,2}\}, \{gk_{1,3}\}, \{gk_{1,5}\}, \{gk_{1,6}\}, \\ \{gk_{2,2}\}, \{gk_{2,3}\}, \{gk_{2,5}\}, \{gk_{2,6}\}, \\ \{s_1\}, \{s_2\} \end{array} \right\}$$

as SAS+ variable assignments. Let $\Omega(\delta)$ denote that set, excluding $\{s_1\}$ and $\{s_2\}$. Note, $\forall vs \in \Omega(\delta)$ we have that $\delta|_{vs}$ is acyclic. Consequently, we have that PUR(δ) evaluates after $\prod_{vs \in \Omega(\delta)} |vs|$, i.e. $3^4 = 81$, calls to \mathbf{S}_{\max} .

3.8.1 Hybrid Algorithm

We have just observed a situation where PUR can exhibit a runtime that is linear in the size of the state space. That is favourable compared to exact calculations of diameter, which in our opening remarks we noted to have worse-than-quadratic runtime. Nevertheless this is unacceptable in our factored setting, and we now seek to alleviate this computational burden by applying \mathbf{S}_{\max} to abstract sub-systems obtained using projections that motivated Definition 17. Such abstractions can be significantly smaller than the concrete systems, thus motivating a hybrid approach that can exponentially reduce bound computation times.

Example 19. Consider applying the approach outlined in Example 14 to compute PUR only on the abstractions $\delta|_{\text{ROOM1}}$ and $\delta|_{\text{ROOM2}}$. PUR($\delta|_{\text{ROOM1}}$) can be evaluated in $\prod_{vs \in \Omega(\delta|_{\text{ROOM1}})} |vs|$ calls to \mathbf{S}_{\max} , where $\Omega(\delta|_{\text{ROOM1}}) = \{ck_1, lk_1, \{gk_{1,2}\}, \{gk_{1,3}\}, \{gk_{2,2}\}, \{gk_{2,3}\}\}$. The same observation can be made for the evaluation time of PUR($\delta|_{\text{ROOM2}}$). Thus the product expression in Example 18 is split into a sum if PUR is called on projections.

We now give an upper-bounding algorithm, HYB, that combines exploitation of acyclic variable dependency with exploitation of acyclicity in state spaces.

Algorithm 5: HYB(δ)

Compute the dependency graph $\mathcal{G}_{\mathcal{D}(\delta)}$ of δ and its SCCs
 Compute the lifted dependency graph \mathcal{G}_{VS}
if $2 \leq |\mathcal{G}_{VS}.V|$ **return** $N_{\text{sum}}\langle \text{HYB} \rangle(\delta, \mathcal{G}_{VS})$
else if $\Omega(\delta) \neq \emptyset$ **return** $S_{\text{max}}\langle \text{HYB} \rangle(ch(\Omega(\delta)), \delta)$
else return UPBND(δ)

In HYB, ch is an arbitrary choice function. Note that in HYB, S_{max} is only applied to the given transition system δ if there is no non-trivial projection (i.e. if $\mathcal{G}_{\mathcal{D}(\delta)}$ has one SCC), and UPBND is applied only to base-cases. Also note that $\mathcal{G}_{\mathcal{D}(\delta)}$ is constructed and analysed with every recursive call to HYB, as snapshotting in earlier calls can remove variable dependencies as a result of removing actions, leading to the breaking of the SCCs in $\mathcal{G}_{\mathcal{D}(\delta)}$, as shown in Example 20.

Example 20. As shown in Figure 3.13b, the dependency graph of $\delta|_{\text{ROOM1}}$ has a single SCC, and thus not susceptible to dependency analysis. Taking a snapshot of $\delta|_{\text{ROOM1}}$ at the assignment $CK_{1,2}$ yields a system with one SCC in its dependency graph as well, as shown in Figure 3.13e. However, taking the snapshot of $\delta|_{\text{ROOM1}}|_{CK_{1,2}}$ at the assignment $\{\overline{1k_{1,1}}, 1k_{1,2}, \overline{1k_{1,3}}\}$, denoted by $LK_{1,2}$, yields a system with an acyclic dependency graph as shown in Figure 3.13f.

We prove HYB is sound by proving it is sound for as tight a base function as possible. Then soundness for using UPBND as a base function follows. As discussed above, d cannot be used, because $N_{\text{sum}}\langle d \rangle$ is not a valid upper bound on d . However, using ℓ as a base-case function is sound. To prove that, we derive an analogue of Theorem 13 for ℓ , where $S_{\text{max}}\langle \ell \rangle(vs, \delta) = S_{\text{max}}\langle \ell \rangle(\Sigma(\mathcal{G}(\delta|_{vs})))$.

Theorem 14. If $\delta|_{vs}$ is acyclic and b bounds ℓ , then $\ell(\delta) \leq S_{\text{max}}\langle b \rangle(vs, \delta)$.

This theorem follows from an argument analogous to that provided for Theorem 13, taking ℓ to be d . Using this theorem, and Lemma 3, the validity of HYB as an upper bound on ℓ (and accordingly, the diameter) follows.

Proposition 9. If UPBND bounds ℓ , then $\ell(\delta) \leq \text{HYB}(\delta)$.

3.9 Empirical Evaluations

We first discuss the practicalities of implementing HYB. Following Rintanen and Getton (2013), we take a base-case function, UPBND, which gives the cardinality of the state space. This choice is pragmatic, taken in light of the fact that computation of alternatives, such as recurrence and sublist diameters, is NP-hard. To optimise computing N_{sum} and S_{max} , we use memoisation, where we compute N or S once for every projection or snapshot, respectively, and store it in a look-up table. This reduced the bound computation time by 70% on average. Our evaluation considers problems from previous International Planning Competitions (IPC), and the unsolvability IPC, and open Qualitative Preference Rovers benchmarks from IPC2006. Below, the latter are referred to as NEWOPEN.

3.9.1 Quality of HYB Bounds

Two measurements related to a compositional upper-bounding algorithm are indicative of its quality. First, we seek an indication of the degree of decompositionality provided by the algorithm. An indication is provided by comparing the size of the domain of the concrete problem—i.e. $|\mathcal{D}(\delta)|$ —with that of the largest base-case. A strong decomposition is indicated when the domain of the base-case is small relative to the concrete problem. Second, we seek an approach that is able to produce bounds that grow sub-exponentially with the size of the problem, when they exist. Thus, we measure how the upper bounds scale in domains as the size of the problem instances grow. If the bounds scale gracefully, this indicates an effective compositional approach.

We report our measurements of the performance of HYB in these terms. Our experiments were conducted on a uniform cluster with time and memory limits of 30minutes and 4GB, respectively. Figure 3.14b shows the domain size of the largest base-case compared to the size of the concrete problem. IPC domains with instances remarkably susceptible to decomposition by HYB are: ROVERS (both solvable and unsolvable), STORAGE, TPP (both solvable and unsolvable), LOGISTICS, NEWOPEN, NOMYSTERY (both solvable and over-subscribed), UNSOLVABLE MYSTERY, VISITALL, SATELLITES, ZENO TRAVEL, and ELEVATORS. For those problems, the size of the largest base-case is significantly smaller than the size of the concrete problem, as shown in Figure 3.14b. One IPC domain that is particularly amenable to decomposition is the ROVERS domain, where many of its instances are decomposed to have largest base-cases modelling a single Boolean state-variable. Also, for domains susceptible to decomposition, the bounds computed by HYB grow sub-exponentially with the number of state variables, as shown in Figure 3.14a. We also note that out of those domains, LOGISTICS, NOMYSTERY, SATELLITES, ZENO TRAVEL and ELEVATORS, have linear (or almost linear) growth of the bounds with the size of the problem.

We also ran HYB on a PDDL McDermott et al. (1998) encoding of the hotel key protocol, with the parameters G , k , and R ranging between 1 and 10 (i.e. 1000 instances of the protocol). As shown in Figure 3.14b (and in the examples earlier), this protocol is particularly amenable to decomposition by HYB. All instances had a largest base-case modelling a single Boolean state-variable. Additionally, the bounds computed by HYB for this set of benchmarks are constant in the number of guests G , grow linearly in the number of rooms R , and quadratically in the number of keys per room K .

3.9.2 Comparison of HYB and N_{sum}

We compared the performance of the hybrid compositional bounding algorithm HYB with N_{sum} , which we showed earlier to dominate other state-of-the-art compositional bounding algorithms. Our experimental cluster and settings are as above. Our analysis and experimentation shows that HYB significantly outperforms N_{sum} , both in terms of decomposition quality and the tightness of computed bounds. This is particularly the case for the domains: NEWOPEN, NOMYSTERY, ROVERS, HYP, TPP, VISITALL, and BOTTLENECK. The success of HYB in our experimentation reveals something of an abundance of problems with acyclicity in their state space. Figure 3.15b indicates that HYB is more successful in decomposing problems compared to N_{sum} , where the largest base-cases for HYB are smaller than those for N_{sum} in 71% of the IPC problems. This observation is reinforced, considering that the 1000th largest bound computed by HYB is 50,534, while the 1000th largest bound computed by N_{sum} is more than 10^6 . In the HOTELKEY domain, the difference is even more pronounced. The bound computed by HYB is at most 990 for all the 1000 instances, while for N_{sum} only 285 instances have bounds less than 10^6 .

Figure 3.15a shows the computational cost of this improved bounding performance. HYB typically required more computation time than N_{sum} . However, HYB terminated in 60 seconds,

or less, for 93% of the benchmarks. Thus, we have not observed a significant time penalty. We note that the improved compositionality over N_{sum} exhibited here has further application yet to be explored. Should we take UPBND to be a more expensive operator, such as the NP-hard recurrence or sublist diameters, the stronger decomposition indicates that UPBND is invoked for relatively small instances when using HYB compared to N_{sum} . Thus, computing UPBND can be exponentially easier for decompositions computed by HYB compared to decompositions from N_{sum} .

3.9.3 Planning with HYB

To evaluate the practical utility of the bounds calculated using HYB, we take them as the queried horizon using the MP version of the SAT-based planner Madagascar Rintanen (2012). In our experiments we limited the time and memory for planners to 1 hour (inclusive of bound computation) and 4GB. The resulting planner proves the safety of 635 instances of the hotel key protocol, where the instance with 9 rooms, 7 guests, and 45 keys, takes the longest to prove safe – it took just under 30 minutes. This is a substantial improvement over the size of instances automatically proven safe in earlier work. We also ran AIDOS 1 Seipp et al. (2014) (unsolvability IPC winner) on the hotel key instances and it proved the safety of only 285 of them, where the instance with 2 rooms, 5 guests, and 10 keys, took the longest to prove safe – in 17 minutes. For the IPC benchmarks, our planner proved that 53 instances are unsolvable, 27 of which could not be proven unsolvable by AIDOS 1. The 27 instances are from BOTTLENECK (7 problems), 3UNSAT (4 problems), ELEVATORS (5 problems), and NEWOPEN (11 problems). We also note that compared to the system from Rintanen and Gretton (2013), we are additionally able to close the heretofore open 7th and 8th Qualitative Preference problem from IPC2006. We also found our bounds useful in solving satisfiable benchmarks. It allowed MP to solve 162 instances that it could not with its default query strategy. Those instances are from ELEVATORS (150 problems), DIAGNOSIS (8 problems), ROVERS (1 problem) and SLIDING-TILES (3 problems).

3.10 Conclusion and Open Questions

The practical incompleteness of SAT-based planning and bounded model checking has for some years been noted as a significant problem Clarke et al. (2004). It is perceived as a deficiency of SAT methods in making comparisons with state based methods. This is due to the absence of effective upper-bounding methods for topological properties of digraphs that model state spaces. We have addressed that deficiency by advancing the compositional approach to upper-bounding, virtually, the only practical approach. Although the compositional approach was intensely investigated to solve other graph theoretic questions concerning state spaces, most notably reachability, there is not as much work on compositionally upper-bounding topological properties of state spaces, except for Baumgartner et al. (2002); Rintanen and Gretton (2013). We considered a compositional approach based on a well studied abstraction, which is *projection*, and a well studied system structure, which is *state variable dependency*. We showed that projection and variable dependencies cannot be used to compositionally upper-bound two interesting topological properties: the diameter and the recurrence diameter.

To mitigate this, we defined new topological properties, the *sublist diameter* and the *traversal diameter*, which can be upper-bounded compositionally, and accordingly can be used to upper-bound the diameter and/or the recurrence diameter. In addition to being compositionally bounded, those new properties can be exponentially smaller than existing properties (like the size of the state space or the recurrence diameter) that were used in previous works to compositionally upper-

bound the diameter and the recurrence diameter. We also showed that those new parameters have interesting properties. For example, the sublist diameter exploits the factored representation: its value depends on the factored representation rather than the digraph modelling the state space. The traversal diameter on the other hand can be computed in linear time, and thus can be generally used as a completeness threshold instead of the NP-hard recurrence diameter.

To be able to deal with realistic problems we introduced a compositional procedure that exploits *acyclicity in the state space* and that is based on a new abstraction that we defined, *snapshooting*. An advantage of this new approach is that both, the diameter and the recurrence diameter can be compositionally bounded using it. Combined with projection based compositional bounding, the hybrid procedure produces very fine grained abstractions and very tight bounds compared to existing compositional upper-bounding procedures. Those benefits comes with the risk of an exponential explosion in the number of abstractions considered by the algorithm. The runtime measurements we made experimentally suggest that this theoretical risk is not realised in practice.

An important contribution was our methodology of defining new topological properties and new abstractions to indirectly compositionally upper-bound interesting topological properties, that themselves cannot be upper-bounded compositionally. We believe that the following questions are interesting for future research.

- Are there certain classes of temporal logic formulae, for which the sublist diameter or the traversal diameter is a tight completeness threshold?
- Does every digraph have a factored representation with a sublist diameter equal to the diameter? If this is possible, obtaining this factoring must be at least an NP-hard problem, since computing the sublist diameter is NP-hard, while computing the diameter is in P. If this is not the case, how close can the sublist diameter be, to the diameter, after factoring?
- N_{sum} computes bounds on the diameter that *cannot* be improved, if the sublist diameters or the recurrence diameters of projections, and the dependencies are given. Would N_{sum} still be tight given more knowledge, like the projections themselves instead of their sublist diameters or recurrence diameters.
- As we stated earlier, the polynomial returned by N_{sum} depends only on the structure of the given DAG. We conjecture that for *any* polynomial, ρ , there is a DAG for which N_{sum} returns a polynomial isomorphic with ρ .




Figure 3.6: Referring to Example 5. (a), (b), and (c) are the largest connected components in the state spaces of the flowers $\mathbb{I}_{4,2}^{u_1}$, $\mathbb{I}_{4,2}^{u_2}$ and $\mathbb{I}_{4,2}^{u_3}$, respectively, (d) a lifted dependency graph of the factored system δ , and (e) is the largest connected component in the state space of δ . In this graph our notation is flattened, where for instance the state $x_i^{u_1} \uplus x_j^{u_2} \uplus x_k^{u_3}$ is denoted by ijk . States x_1 and x_2 are indicated by the blue and green vertices, respectively. The red vertices indicate the states traversed by $\vec{\pi}$ if executed on x_1 . Note that (a), (b), and (c) are contractions of (e).




Figure 3.7: Referring to Example 7, (a), (b), and (c) are the largest connected components in the state spaces of the flowers $\hat{I}_{3,2}^{u_2}$, $\hat{I}_{2,2}^{u_3}$ and $\hat{I}_{2,2}^{u_1}$, respectively, (d) a lifted dependency graph of the factored system δ from Example 7, and (e) is the largest connected component in the state space of δ .




Figure 3.8: A lifted dependency graph for a factored digraph that has four sets of variables closed under mutual dependency.




Figure 3.9: The bounds computed by $N_{\text{sum}}(b)$ versus $M_{\text{sum}}(b)$ with $2^{|\mathcal{D}(\delta)|} - 1$ as a base function.




Figure 3.10: An inverted flower with $3 \leq m$.

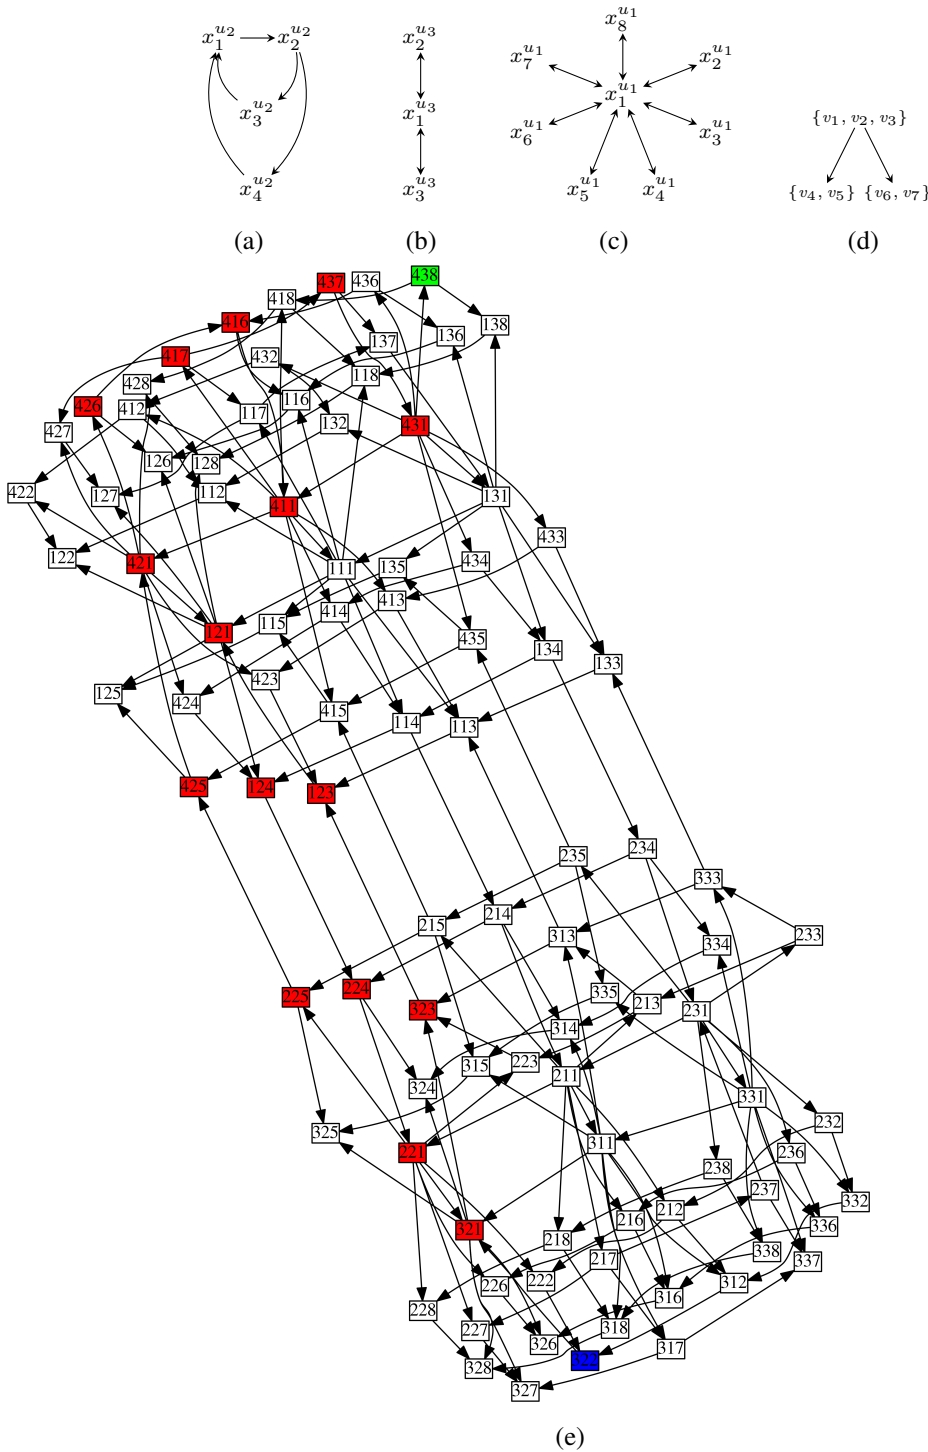


Figure 3.11: Referring to Example 10, (a), (b), and (c) are the largest connected components in the state spaces of the inverted flowers $\overleftarrow{f}_{2,3}^{u_2}$, $\overleftarrow{f}_{2,2}^{u_3}$ and $\overleftarrow{f}_{7,2}^{u_1}$, respectively (i.e. $\mathcal{G}(\overleftarrow{f}_{2,3}^{u_2})$, $\mathcal{G}(\overleftarrow{f}_{2,2}^{u_3})$, and $\mathcal{G}(\overleftarrow{f}_{7,2}^{u_1})$), (d) a lifted dependency graph of the factored system δ , and (e) is the largest connected component in the state space of δ . States x_1 and x_2 are indicated by the blue and green vertices, respectively. The red vertices indicate the states traversed by $\overrightarrow{\pi}$ if executed on x_1 . Note that (a), (b), and (c) are contractions of (e).

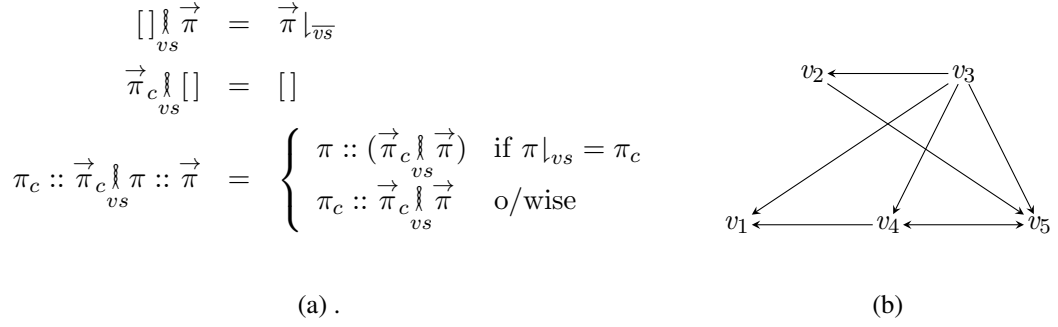



Figure 3.12: (a)The definition of the stitching function ($\Big|$), and (b) is the dependency graph of the system in Example 11.

; check in to a room (at reception), receiving a new key
 $(\{lk_{1,1}\}, \{gk_{1,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{1,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{lk_{1,1}\}, \{gk_{2,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{2,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{lk_{2,4}\}, \{gk_{1,5}, lk_{1,5}, \overline{lk_{1,4}}, \overline{s_2}\}), (\{lk_{2,5}\}, \{gk_{1,6}, lk_{1,6}, \overline{lk_{1,5}}, \overline{s_2}\}),$
 $(\{lk_{2,4}\}, \{gk_{2,5}, lk_{1,5}, \overline{lk_{1,4}}, \overline{s_2}\}), (\{lk_{2,5}\}, \{gk_{2,6}, lk_{1,6}, \overline{lk_{1,5}}, \overline{s_2}\})$

; enter a room with new key
 $(\{gk_{1,2}\}, \{ck_{1,2}, \overline{ck_{1,1}}, s_1\}), (\{gk_{2,2}\}, \{ck_{1,2}, \overline{ck_{1,1}}, s_1\}),$
 $(\{gk_{1,3}\}, \{ck_{1,3}, \overline{ck_{1,2}}, s_1\}), (\{gk_{2,3}\}, \{ck_{1,3}, \overline{ck_{1,2}}, s_1\}),$
 $(\{gk_{1,5}\}, \{ck_{2,5}, \overline{ck_{2,4}}, s_2\}), (\{gk_{2,5}\}, \{ck_{2,5}, \overline{ck_{2,4}}, s_2\}),$
 $(\{gk_{1,6}\}, \{ck_{2,6}, \overline{ck_{2,5}}, s_2\}), (\{gk_{2,6}\}, \{ck_{2,6}, \overline{ck_{2,5}}, s_2\})$

(a)




(b)

$(\{lk_{1,1}\}, \{gk_{1,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{1,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{lk_{1,1}\}, \{gk_{2,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{2,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{gk_{1,2}\}, \{ck_{1,2}, \overline{ck_{1,1}}, s_1\}), (\{gk_{2,2}\}, \{ck_{1,2}, \overline{ck_{1,1}}, s_1\}),$
 $(\{gk_{1,3}\}, \{ck_{1,3}, \overline{ck_{1,2}}, s_1\}), (\{gk_{2,3}\}, \{ck_{1,3}, \overline{ck_{1,2}}, s_1\})$


(c)

$(\{lk_{1,1}\}, \{gk_{1,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{1,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{lk_{1,1}\}, \{gk_{2,2}, lk_{1,2}, \overline{lk_{1,1}}, \overline{s_1}\}), (\{lk_{1,2}\}, \{gk_{2,3}, lk_{1,3}, \overline{lk_{1,2}}, \overline{s_1}\}),$
 $(\{gk_{1,2}\}, \{s_1\}), (\{gk_{2,2}\}, \{s_1\})$

(d)




(e)




(f)

Figure 3.13: (a) shows the actions of a *transition system* δ representing the hotel key protocol with 2 rooms, 2 guests and 3 keys per room; room 1 is associated with keys 1–3; room 2 with keys 4–6. (b) is the *dependency graph* for that system. (c) is the *projection* of the system on an abstraction that models only the changes related to room 1. (d) is the *snapshot* of $\delta \Big|_{\text{ROOM1}}$ on $CK_{1,2}$, an abstraction that only analyses the changes related to room 1 when its door recognises key 2 as the current key. (e) and (f) are the dependency graphs of snapshots that we use for illustrative purposes in the examples.




(a)




(b)

Figure 3.14: Measurements related to HYB on benchmarks. (a) Scatter plot of the bound (horizontal axis) computed by HYB, and the size (i.e. $|\mathcal{D}(\delta)|$) of the concrete problem (vertical). (b) Scatter plot of the size of the largest base-case (horizontal), and the size of the concrete problem (vertical).



(a)



(b)

Figure 3.15: Comparison of the bounding performance of N_{sum} and HYB on benchmarks. Legend is provided in Figure 3.14a. (a) Scatter plot of the bounds computed by HYB (horizontal axis) and the state-of-the-art bounding algorithm N_{sum} (vertical). (b) Scatter plot of the size (i.e. $|\mathcal{D}(\delta)|$) of the largest base-case using HYB (horizontal axis) and N_{sum} (vertical).




Figure 3.16: Scatter plot of computation time (in seconds) of HYB (horizontal axis) and N_{sum} (vertical) for benchmarks. Legend is provided in Figure 3.14a.

Compositional Computation of Reachability in the Presence of Repetitive Symmetry

Testing whether one vertex (or a set of vertices) in a digraph can be reached from another vertex is one of the oldest and most well studied problems in computer science and graph theory. In AI planning and model checking, this problem is of immense importance, where the digraph is taken to model the state space of the problem, and the question is whether a goal state can be reached from a given initial state. In AI planning, the problem of reachability between states is the main problem whose solution is sought. An initial state and a set of goal states are given, and the quest is to find a valid action sequence (a *plan*) that if executed in the initial state will result in one of the goal states. Similarly, in model checking, the problem of reachability between states corresponds to the problem of checking safety properties of systems. An initial state of the system is given, along with a formula characterising desired safety properties, and a sequence of transitions (a bug trace) from the initial state to a state violating the given formula is searched for.

Although many polynomial time algorithms (e.g. Bellman-Ford algorithm) were introduced to solve the reachability problem for digraphs, they all assume that the digraph is explicit. Accordingly, for propositionally factored transition systems, if one uses any of those algorithms naively, he needs to construct the digraph modelling the state space to solve the reachability problem. Thus, the complexity of reachability in factored transition systems is PSPACE-complete Sistla and Clarke (1985); Bylander (1994), since the digraph modelling the state space can be exponentially larger than the factored representation. The compositional approach alleviates can alleviate the practical complexity of computing reachability, making the problem feasible to solve in many cases, where it otherwise may not be. In this approach, digraphs modelling state spaces of abstractions (which are minors of the original state space) are searched for paths, which are then used to synthesise a path in the digraph modelling the state space of the concrete system, a path that connects the required states. In this chapter we consider applying that approach to compute reachability, where we consider minors of the state space corresponding to the state space of an abstraction of the factored system that we refer to as a “descriptive quotient”. In particular, we exploit a certain type of symmetry, which we call *repetitive symmetries*, in the factored system to obtain those abstractions.

State Variable Symmetry Informally, in a factored system δ , two state variables v_1 and v_2 are symmetric if they can be swapped by a permutation of the state variables of δ , and the resulting system is the same (up to isomorphism) as δ . This state variable symmetry relation induces a partition of the set of state variables and a partition of the actions. We find those partitions most

useful for use with quotienting within our work.

4.1 Related Work

Algorithms for finding a path or the shortest path between two vertices in a digraph were devised since the middle of last century. The earliest algorithm that we are aware of, the Bellman-Ford algorithm Bellman (1958), finds the shortest path between two vertices in a time linear in the product of the edges and vertices of a digraph. Many algorithms were introduced after that like the ones in Dijkstra (1959); Fredman and Tarjan (1987); Thorup (2003). All of those algorithms assume that the digraph is represented explicitly, and also they are all at least linear in both, the number of vertices and the number of edges. Accordingly a naive application of these algorithms to finding paths in state spaces of factored systems would be infeasible.

Compositional Planning and Model Checking

The compositional approach was employed to find solutions to AI planning and model checking problems in work like Knoblock (1994); Williams and Nayak (1997); Berezin et al. (1998); Case et al. (2009); Kroening (2006); Amir and Engelhardt (2003); Brafman and Domshlak (2006); Kelareva et al. (2007); Guere and Alami (2001); Helmert et al. (2014); Sievers et al. (2015). In those works, different types of abstractions and structures were used, and only the state space of a projected system is searched for a path. For example, abstractions based on projection were used in Knoblock (1994); Williams and Nayak (1997), which exploited acyclicity in variable dependencies. Another approach for compositional path search, known as *factored planning*, is used by Amir and Engelhardt (2003); Brafman and Domshlak (2006); Kelareva et al. (2007), in which the factored system δ is abstracted into multiple abstractions, referred to as “factors”. These factors are obtained by an abstraction that, given a partition of the set of state variables $vs_{1..n}$, it produces a set of sets of actions, one per $vs_i \in vs_{1..n}$, such that the set of actions refers only to variables in vs_i . This partition of state variables is based on a tree decomposition of an undirected version of the variable dependency graph.

Exploiting Symmetries

Symmetries that occur in factored transition systems have for some time been exploited for efficient search for paths between states, both by AI planning and model checking communities. The quintessential planning scenario which exhibits symmetries is GRIPPER. This comprises a robot whose left and right grippers can be used interchangeably in the task of moving a set of N indistinguishable packages from a source location to a goal location. Intuitively the left and right grippers are symmetric because if we changed their names, by interchanging the terms left and right in the problem description, we are left with an identical problem. Packages are also interchangeable and symmetric.

One method to exploit symmetry is to perform search in a minor of the digraph modelling the state space, referred to as a *quotient system*, which corresponds to a (sometimes exponentially) smaller bisimulation of the system at hand. Planning in a quotient system, a state x is represented by a *canonical* element from its *orbit*, the set of states which are symmetric to x . Giving integer labels to packages in GRIPPER, when the search encounters a state where the robot is holding 1 package using 1 gripper, this is represented using the canonical state where, for example, the left gripper is holding package with identity “1”. Orbit search explores the quotient system by

simulating actions from known canonical states, and then computing the canonical representations of resultant states. That canonicalisation step requires the solution to the *constructive orbit problem* Clarke et al. (1998) which is *NP-hard* Eugene (1993). A key weakness, is that for each state encountered by the search an intractable canonicalisation operation is performed. This is mitigated in practice by using approximate canonicalisation. By forgoing exact canonicalisation, one encounters a much larger search problem than necessary. For a GRIPPER instance with 42 packages, the breadth-first orbit search with approximate canonicalisation by Pochter et al. (2011) reportedly performs $60.5K$ state expansion operations, far more than necessary. This method of exploiting symmetry was firstly used in solving model checking problems, which is reviewed by Wahl and Donaldson (2010). Also, this method was recently adapted for planning and studied by Pochter et al. (2011), Domshlak et al. (2012, 2013) and Shleyfman et al. (2015). Related work about state-based planning in equivalence classes of symmetric states includes Guere and Alami (2001); Fox and Long (1999, 2002).

Following the seminal work by Crawford et al. (1996) and Brown et al. (1996), when searching for a path (in the context of AI planning) *via* constraint satisfaction, known symmetries are exploited by:

- (i) including symmetry breaking constraints, either directly as part of the problem expression Joslin and Roy (1997); Rintanen (2003), or
- (ii) otherwise dynamically as suggested by Miguel (2001) as part of a *nogood* mechanism.

In GRIPPER, we can statically require that if no package is yet retrieved, then any retrieval targets package 1 with the left gripper. Dynamically, having proved that no 3-step plan exists retrieving package 1 with the left gripper, then no plan of that length exists retrieving package $i \neq 1$ using either gripper. Searching using the proposed dynamic approach is quite weak, as symmetries are only broken as nogood information becomes available. A weakness of both approaches is that problems expressed as CSPs include variables describing the state of the transition system at different plan steps. Existing approaches do not break symmetries across steps, and can therefore waste effort exploring partial assignments that express executions which visit symmetric states at different steps.

Another method to exploit symmetries is via abstracting the factored system at hand. Some authors perform search in the state space of a *counter abstraction*, as surveyed by Wahl and Donaldson (2010). That approach treats a *transition system isomorphic* to the quotient system, but it avoids solving the NP-hard constructive orbit problem. Symmetries were also explored in factored planning in the context of *merge-and-shrink* heuristics Helmert et al. (2014). Sievers et al. (2015) developed a symmetry guided *merging* operation which yields relatively compact heuristic models, improving the scalability of that approach. Guere and Alami (2001) propose to plan via a *shape-graph*, a compact description of the problem state space in which states are represented by equivalence classes of symmetric states. As well as planning by searching in that graph, using the diameter of the *shape-graph* Guere and Alami are able to calculate tight upper bounds for the highly symmetric GRIPPER and BLOCKS-WORLD domains.

4.2 Results

We consider the AI planning problem defined on factored transition systems, i.e. given an initial state $I \in \mathbb{U}(\delta)$ and a set of goal states G , is there an action sequence that reaches one of the goal states if it is executed at the initial state.

Our contribution is a new abstraction of factored transition systems, the result of which we call a *descriptive quotient*. We provide conditions under which searching for a path between the initial state and a goal state (i.e. a plan) in the digraph modelling the state space of the descriptive quotient can be used to synthesise a path between the initial state and some goal state in the original system. The first condition is that the partition used to obtain the descriptive quotient is induced by the symmetry relation between state variables (a.k.a. orbits of a subgroup of the automorphism group). The second condition is that the descriptive quotient is isomorphic to a “sub-system” of the original factored system. Informally, those two conditions mean that planning via descriptive quotients is a way to exploit “repetitive symmetries” in a factored system, where by repetitive symmetries we mean that the factored system is constituted by a union of isomorphic sub-systems.

Based on that we provide a novel procedure for domain-independent planning with symmetries. Following, e.g., Pochter et al., in a first step we infer knowledge about the symmetries between state variable. Then departing from existing approaches, our second step uses that knowledge to obtain a quotient of the concrete factored system. Called the descriptive quotient, this describes any element in the set of isomorphic subsystems which abstractly model the concrete system. Third, we invoke a planner once to solve the small reachability problem posed by that descriptive quotient. In the fourth and final step, a concrete plan is synthesized by concatenating instantiations of that plan for each isomorphic subproblem.

The non-existence of a plan for the descriptive quotient does not exclude the possibility of a plan in the concrete system. Although sound, in that respect our approach is incomplete. Having an optimal plan for the quotient does not guarantee optimality in the concatenated plan. Aside from computing a plan for the descriptive quotient, the computationally challenging steps occur in preprocessing:

- (i) Identification of state variable symmetries from the original description, a problem as hard as graph isomorphism, which is not known to be tractable, and
- (ii) Computing an appropriate set of subsystems isomorphic to the quotient.

We introduce the general version of the latter problem: for an undirected graph \mathcal{G} and a partition of its vertices P , is the quotient \mathcal{G}/P isomorphic to a subgraph of \mathcal{G} , such that the morphism maps every set of vertices from P to one of its members? We show that this problem is NP-complete. We also show that if P is a set of *orbits* of a symmetry group for \mathcal{G} , then there is set of morphisms from \mathcal{G}/P to \mathcal{G} that covers all the vertices of \mathcal{G} , and whose size is logarithmic in the vertices of \mathcal{G} .

Unlike existing approaches, our search for a plan does not need to reason about symmetries between concrete states and the effects of actions on those. Plan search can be performed by an off-the-shelf SAT/CSP system, in which case *symmetry breaking* constraints are not required. Alternatively, using a state-based planner we avoid repeated (approximate) solution to the intractable canonicalisation problem, a clear bottleneck of recent planning algorithms. In this respect, our approach is similar to searching in a *counter abstraction*, as surveyed by Wahl and Donaldson (2010). Also, viewing our approach as one that decomposes a problem into subproblems, it is related to factored planning Amir and Engelhardt (2003); Brafman and Domshlak (2006); Kelareva et al. (2007).

The most important advantage of searching in the digraph modelling the state space of the descriptive quotient is that it is a minor that can be very compact relative to the digraph modelling the state space of the original factored system. For instance, a counter abstraction has a state space which can be exponentially larger than that of a descriptive quotient—i.e., the descriptive quotient will model 1 object, whereas the quotient transition system models N symmetric objects. Also,




Figure 4.1: The largest connected component of the state space of the transition system underlying the problem from Example 21. It shows the presence of symmetries between different states.

existing state-based methods plan in that relatively large *quotient system*, and as just mentioned, face an additional intractable problem for every encountered state. By employing approximate canonicalisation, such methods face a state space much larger than that posed by the quotient system, which can be exponentially larger than that posed by a descriptive quotient. To give an idea of how compact can the size of the state space of the descriptive quotient, the descriptive quotient of the aforementioned 42-package GRIPPER instance is solved by breadth-first search expanding 6 states, and a concrete plan is obtained in under a second. On the other hand, the approach by Pochter et al. (2011) expands $60.5K$ states and takes about 28 seconds.

4.3 Planning Problems and Additional Notation

We formally define *planning problems*.

Definition 24 (Planning Problem). A *planning problem* Π is a 3-tuple $\langle I, \delta, G \rangle$, with I the initial state of the problem, G a description of goal states (another finite map from variables to Booleans), and δ a set of actions. We define the domain of the problem ($\mathcal{D}(\Pi)$) to be domain of the set of actions ($\mathcal{D}(\delta)$). The set of valid states with respect to a planning problem, which we write as $\mathbb{U}(\Pi)$ is $\mathbb{U}(\delta)$. Problem Π is valid if $\mathcal{D}(G) \subseteq \mathcal{D}(I)$, and $I \in \mathbb{U}(\Pi)$. We only consider valid problems. Hereafter we refer to the initial state, actions or goal of problem Π as $\Pi.I$, $\Pi.\delta$ or $\Pi.G$ respectively. We may also omit the Π if it is clear from the context, e.g. I for $\Pi.I$ and δ_i for $\Pi_i.\delta$. An action sequence $\vec{\pi}$ is a plan/solution for a planning problem Π iff $\vec{\pi} \in \delta^*$ and $G \subseteq \text{ex}(I, \vec{\pi})$.

Definition 25 (Subproblem). Problem Π_1 is a subproblem of Π_2 , written $\Pi_1 \subseteq \Pi_2$, if $I_1 \subseteq I_2$, and if $\delta_1 \subseteq \delta_2$.

Definition 25 purposefully only treats the description of system dynamics, and does not consider problem goals. We later consider subproblems with a variety of goals which do not all correspond to concrete problem goals. We form a concrete plan by concatenating plans for subproblems with such varieties of goals. For example, we will consider subproblems from GRIPPER with only one package and one gripper, where the goal is to: (i) relocate that package according to the concrete goal, (iii) additionally free the gripper, and (ii) additionally have the robot relocate to its starting position. In this case a concrete plan is formed by concatenating plans for subproblems given for each distinct package.

Lastly, we require a few notations. For a set of actions δ , we define the set of preconditions $\text{pre}(\delta)$ as $\bigcup_{(p,e) \in \delta} p$. Let m be a finite map, e.g., a state x , etc.. Then, let $f(m)$ be the image of m under function f : the map $\bigcup_{(k \mapsto v) \in m} \{f(k) \mapsto v\}$. This is well-defined if $f(k_1) = f(k_2)$ implies that $m(k_1) = m(k_2)$. We lift this notion of image to other composite types. For example, we write $f(\Pi)$ for the image of Π under f , where all finite maps in Π are transformed by f .



Figure 4.2: The quotient state space of the system from Example 21 in which the traditional orbit search algorithm, ideally, would search for a path.



Figure 4.3: The largest connected component in the state space of the transition system underlying the descriptive quotient of the problem in Example 21.

4.4 Computing Problem Symmetries

To exploit problem symmetries we must first discover them. We follow the discovery approach from Pochter et al. (2011), restricting ourselves to Boolean-valued variables in $\mathcal{D}(\Pi)$. We assume familiarity with *groups*, *subgroups*, *group actions*, and *graph automorphism groups*. Symmetries in a problem description are defined by an *automorphism group*.

Definition 26 (Problem Automorphism Group). *The automorphism group of Π is: $Aut(\Pi) = \{\sigma \mid \sigma(\Pi) = \Pi\}$. Members of $Aut(\Pi)$ are permutations on $\mathcal{D}(\Pi)$ and they induces a partition of $\mathcal{D}(\Pi)$ to which we refer as orbits.*

Example 21. *Consider the planning problem Π_1 with*

$$\begin{aligned} \Pi_1.I &= \{v_3, v_1, v_2, v_4, v_5\} \\ \Pi_1.\delta &= \{(\{v_3, v_1\}, \{\overline{v_4}, \overline{v_3}\}), (\{v_3, v_2\}, \{\overline{v_5}, \overline{v_3}\}), (\emptyset, \{v_3\})\} \\ \Pi_1.G &= \{\overline{v_4}, \overline{v_5}\} \end{aligned}$$

The largest connected component of $\mathcal{G}(\delta_1)$ is shown in Figure 4.1. $Aut(\Pi_1)$ is the closure of this set of permutations under composition: $\{\{v_1 \mapsto v_2, v_2 \mapsto v_1\}, \{v_4 \mapsto v_5, v_5 \mapsto v_4\}\}$. Let $v_{s_{1..n}}$ be $\{p_1 = \{v_1, v_2\}, p_2 = \{v_3\}, p_3 = \{v_4, v_5\}\}$. This partition of $\mathcal{D}(\Pi_1)$ is the set of orbits induced by $Aut(\Pi_1)$. This variable symmetry induces symmetries between the states and accordingly vertices of $\mathcal{G}(\delta_1)$ as shown in Figure 4.1, where state 1 is symmetric with 2, and 5 is symmetric with 6. The quotient system in which, ideally, traditional approaches would search for a path to exploit symmetry is shown in Figure 4.2, and it is clearly smaller in size than $\mathcal{G}(\delta_1)$.

A graphical representation of Π is constructed so vertex symmetries in it correspond to variable symmetries in Π . We follow the graphical representation introduced in Pochter et al. (2011).

Definition 27 (Undirected Graph). *An undirected graph \mathcal{G} is a tuple $\langle V, E \rangle$, where V is the set of vertices of \mathcal{G} and E is the set edges of \mathcal{G} which is a set of unordered pairs from V . We write $V(\mathcal{G})$ for the set of vertices, and $E(\mathcal{G})$ for edges of a graph \mathcal{G} .*

Definition 28 (Problem Description Graph (PDG)). *The (undirected) graph $\mathcal{G}(\Pi)$ for a planning problem Π , is defined as follows:*

- (i) $\mathcal{G}(\Pi)$ has two vertices, u_1^\top and u_1^\perp , for every variable $v \in \mathcal{D}(\Pi)$; two vertices, a_p and a_e , for every action $(p, e) \in \delta$; and vertex u_{1I} for I and u_{1G} for G .
- (ii) $\mathcal{G}(\Pi)$ has an edge between u_1^\top and u_1^\perp for all $v \in \mathcal{D}(\Pi)$; between a_p and a_e for all $(a, e) \in \delta$; between a_p and v^* if $(v \mapsto *) \in p$, and between a_e and v^* if $(v \mapsto *) \in e$; between v^* and v_I if $(v \mapsto *)$ occurs in I ; and between v^* and v_G if $(v \mapsto *)$ occurs in G .

The automorphism group of the PDG, $Aut(\mathcal{G}(\Pi))$, is identified by solving an undirected graph isomorphism problem. The action of a subgroup of $Aut(\mathcal{G}(\Pi))$ on $V(\mathcal{G}(\Pi))$ induces a *partition*, called the *orbits*, of $V(\mathcal{G}(\Pi))$. We can now define our *quotient* structures based on partitions $vs_{1..n}$ of $\mathcal{D}(\Pi)$.

Definition 29 (Quotient Undirected Graph). *For graph \mathcal{G} and a partition P of its vertices, the quotient \mathcal{G}/P is the digraph:*

- (i) $V(\mathcal{G}/P) = P$, and
- (ii) $E(\mathcal{G}/P) = \{\{U, W\} \mid U, W \in P \wedge \exists u_2 \in U, u_3 \in W. \{u_2, u_3\} \in E(\mathcal{G})\}$.

Definition 30 (Quotient Problem). *Given partition $vs_{1..n}$ of $\mathcal{D}(\Pi)$, let \mathcal{Q} map members of $\mathcal{D}(\Pi)$ to their equivalence class in $vs_{1..n}$. The descriptive quotient is $\Pi/vs_{1..n} = \mathcal{Q}(\Pi)$, the image of Π under \mathcal{Q} . This is well-defined if $vs_{1..n}$ is a set of orbits. We assume that quotient problems are well-defined.*

Example 22. *Consider the planning problem Π_1 from Example 21, and the partition $vs_{1..n}$. The descriptive quotient associated with $vs_{1..n}$ is as follows.*

$$\begin{aligned} (\Pi_1/vs_{1..n}).I &= \{p_1, p_2, p_3\} \\ (\Pi_1/vs_{1..n}).\delta &= \{(\{p_1, p_2\}, \{\overline{p_2}, \overline{p_3}\}), (\emptyset, \{p_2\})\} \\ (\Pi_1/vs_{1..n}).G &= \{\overline{p_3}\} \end{aligned}$$

The largest component of the minor $\mathcal{G}((\Pi_1/vs_{1..n}).\delta)$ is shown in Figure 4.3. It is clearly smaller than the original state space shown in Figure 4.1 as well as the quotient system shown in Figure 4.2.

To ensure correspondence between PDG symmetries and problem symmetries, we must ensure incompatible descriptive elements do not coalesce in the same orbit. For example, we cannot have action precondition vertices and state-variables in the same orbit.

Definition 31 (Well-formed Partitions). *A partition of $V(\mathcal{G}(\Pi))$ is well-formed iff:*

- (i) *Positive (u_1^\top) and negative (u_1^\perp) variable assignment vertices only coalesce with ones of the same parity;*
- (ii) *Precondition (a_p) and effect (a_e) vertices only coalesce with preconditions and effects respectively, and*
- (iii) *Both u_{1I} and u_{1G} are always in a singleton.*

A well-formed partition \hat{P} defines a corresponding partition $vs_{1..n}$ of $\mathcal{D}(\Pi)$, so that $\mathcal{G}(\Pi)/\hat{P} = \mathcal{G}(\Pi/vs_{1..n})$

To ensure well-formedness, vertex symmetry is calculated using the *coloured* graph-isomorphism procedure (CGIP). Vertices of distinct colour cannot coalesce in the same orbit. Vertices of $\mathcal{G}(\Pi)$ are coloured to ensure the orbits correspond to a well-formed partition.

4.5 Computing the Set of Instantiations

Recapping, symmetries in Π are the basis of a partition $vs_{1..n}$ of its domain $\mathcal{D}(\Pi)$ into orbits. That exposes the *descriptive quotient*, $\Pi/vs_{1..n}$, an abstract problem whose variables correspond

to orbits of concrete symbols. Our task now is to compute a set of *instantiations* of the quotient which *cover* all the goal variables $\mathcal{D}(G)$. Called a covering set of isomorphic subproblems, by instantiating a quotient plan for each subproblem and concatenating the results we intend to arrive at a concrete plan. We describe a pragmatic approach to obtaining that covering set. We establish that a covering set is not guaranteed to exist, and give an approach to *refining* partitions to mitigate that fact. We derive a theoretical bound on the necessary size of a covering set, and prove that a general graph formulation of the problem of computing a covering instantiation is NP-complete.

Definition 32 (Transversals, Instantiations). *A transversal \mathfrak{h} is an injective choice function over a set of sets S such that $\mathfrak{h}(c) \in c$ for every equivalence class $c \in S$. A transversal covers v if $v \in \mathcal{R}(\mathfrak{h})$. If $vs_{1..n}$ is a partition of $\mathcal{D}(\Pi)$, we refer to a transversal \mathfrak{h} of $vs_{1..n}$ as an instantiation of $\Pi/vs_{1..n}$. An instantiation \mathfrak{h} is consistent with a concrete problem Π if $\mathfrak{h}(\Pi/vs_{1..n}) \subseteq \Pi$. When we use the term “problem” discussing an instantiation \mathfrak{h} , we intend $\mathfrak{h}(\Pi/vs_{1..n})$. Note that $\mathcal{D}(\mathfrak{h}(\Pi/vs_{1..n})) = \mathcal{R}(\mathfrak{h})$.*

Example 23. *Consider the set $s = \{a, b, c, d, e, f\}$, and the equivalence classes $c_1 = \{a, b\}$, $c_2 = \{c, d\}$ and $c_3 = \{e, f\}$ of its members. For the partition $P = \{c_1, c_2, c_3\}$ of s , $\mathfrak{h}_1 = \{a, c, e\}$ and $\mathfrak{h}_2 = \{b, c, f\}$ are two transversals.*

For each goal variable in the problem, our approach shall need to find a consistent instantiation of the quotient which covers that. We thus face the following problem.

Problem 1. *Given Π and a partition $vs_{1..n}$ of $\mathcal{D}(\Pi)$, is there an instantiation of $\Pi/vs_{1..n}$ consistent with Π that covers a variable $v \in \mathcal{D}(\Pi)$?*

Example 24. *Consider the planning problem Π_1 from Example 22. Let instantiation \mathfrak{h} be $\{p_1 \mapsto v_3, p_2 \mapsto v_1, p_3 \mapsto v_4\}$. Then, \mathfrak{h} covers v_4 , and $\mathfrak{h}(\Pi_1/vs_{1..n})$ has $I = \{v_3, v_1, v_4\}$, $A = \{(\{v_3, v_1\}, \{\overline{v_4}, \overline{v_3}\}), (\emptyset, \{v_3\})\}$ and $G = \{\overline{v_4}\}$. Thus, $\mathfrak{h}(\Pi_1/vs_{1..n}) \subseteq \Pi_1$, making \mathfrak{h} a consistent instantiation and so a solution to Problem 1 with inputs Π_1 , \mathcal{P} and v_4 .*

4.5.1 Finding Instantiations: Practice

We now describe how we obtain a set of isomorphic subproblems of Π that covers the goal variables. We compute a set of instantiations, Δ , of the quotient $\Pi/vs_{1..n}$. Our algorithm first initialises that set of instantiations, $\Delta := \emptyset$. Every iteration of the *main loop* computes a consistent instantiation that covers at least one variable $v \in \mathcal{D}(\Pi.G)$. This is done as follows: we create a new (partial) instantiation $\mathfrak{h} = p_v \mapsto v$, where p_v is the set in $vs_{1..n}$ containing v . Then we determine whether \mathfrak{h} can be completed, to instantiate every set in $vs_{1..n}$, while being consistent at the same time. This determination is achieved by posing the problem in the theory of uninterpreted functions, and using a satisfiability modulo theory (SMT) solver. Our encoding in SMT is constructive, and if a completion of \mathfrak{h} is possible the solver provides it. If successful, we set $\Delta := \Delta \cup \{\mathfrak{h}\}$ and $G := G \setminus \mathcal{R}(\mathfrak{h})$, and loop. In the case the SMT solver reports failure, the concrete problem cannot be covered by instantiations of the descriptive quotient, and we report failure. In the worst case of $\mathcal{D}(\Pi.G) = \mathcal{D}(\Pi)$, our *main loop* executes $\sum_{p \in vs_{1..n}} |p| - |vs_{1..n}| + 1$ times, and hence we have that many instantiations.

Scenarios need not admit a consistent instantiation.

Example 25. *Take Π_2 with*

$$\begin{aligned} \Pi_2.I &= \{\overline{v_1}, \overline{v_2}, v_3\} \\ \Pi_2.\delta &= \{(\{\overline{v_1}\}, \{v_2, \overline{v_3}\}), (\{\overline{v_2}\}, \{v_1, \overline{v_3}\})\} \\ \Pi_2.G &= \{\overline{v_3}, \overline{v_1}, \overline{v_2}\} \end{aligned}$$

Let $vs_{1..n} = \{p_1 = \{v_1, v_2\}, p_2 = \{v_3\}\}$. The quotient $\Pi_2/vs_{1..n}$ has

$$\begin{aligned} (\Pi_2/vs_{1..n}).I &= \{\overline{p_1}, p_2\} \\ (\Pi_2/vs_{1..n}).\delta &= \{(\{\overline{p_1}\}, \{p_1, \overline{p_2}\})\} \\ (\Pi_2/vs_{1..n}).G &= \{\overline{p_1}, \overline{p_2}\} \end{aligned}$$

There are no consistent instantiations of $\Pi_2/vs_{1..n}$ because v_1 and v_2 are in the same equivalence class and also occur together in the same action.

Our example demonstrates a common scenario in the IPC benchmarks, where a partition $vs_{1..n}$ of $\mathcal{D}(\Pi)$ does not admit a consistent instantiation because variables that occur in the same action coalesce in the same member of $vs_{1..n}$. We resolve this situation by refining the partition produced using CGIP to avoid having such variables coalesce in the same orbit. For a $vs \in vs_{1..n}$, consider the graph $\mathcal{G}(vs)$ with vertices vs and edges $\{\{v_1, v_2\} \mid \exists \pi \in \Pi.\delta \wedge \{v_1, v_2\} \subseteq \mathcal{D}(\pi) \wedge v_1 \neq v_2\}$. The *chromatic number* N of $\mathcal{G}(vs)$ gives us the number of colours needed to colour the corresponding vertices $\hat{v}s$ in the PDG. Where two variables occur in the same action, their vertices in $\hat{v}s$ are coloured differently. For every $vs \in vs_{1..n}$, we use an SMT solver to calculate the required chromatic numbers N and graph colourings for $\mathcal{G}(vs)$, then we colour the corresponding vertices in the PDG according to the computed N -colouring of $\mathcal{G}(vs)$. Lastly, we again pass the PDG to a CGIP after it is recoloured. In 4 benchmark sets the thus-revised partition admits a consistent instantiation where the initial partition does not. Although the *chromatic number* problem is NP-complete, this step is not a bottleneck in practice because the size of the instances that need to be solved is bounded above by the largest number of literals stated in any action.

4.5.2 Finding Instantiations: Theory

We first develop a theoretical bound on the number of required instantiations—by treating abstract covering transversals—that is tight compared to our pragmatic solution above. To characterise the complexity of our instantiation problem, we then study the general problem of computing covering transversals.

Theorem 15 (B. McKay, 2014). *Let $-$ be a group acting on a set V (e.g., $\mathcal{D}(\Pi)$). Suppose \mathfrak{h} is a transversal of O , a set of orbits induced by the action of $-$ on V . Take $S = \sum_{o \in O} |o|$ and $M = \max_{o \in O} |o|$. Where \bullet is composition, there will always be a set of transversals \mathcal{T} with size $\leq M \ln(S)$ such that*

- (i) *Each element $\mathfrak{h}' \in \mathcal{T}$ satisfies $\mathfrak{h}' = \sigma \bullet \mathfrak{h}$ for some $\sigma \in -$, and*
- (ii) *For every $u_1 \in V$, it has an element \mathfrak{h}' that covers u_1 .*

Proof. Take $N = M \ln(S)$ and let H be a subset of $-$ obtained by drawing N permutations of V independently at random with replacement. For any orbit $o \in O$, the probability for a $v \in o$ is not in $\mathcal{R}(\sigma \bullet \mathfrak{h})$ for a randomly drawn $\sigma \in H$ is $1 - 1/|o|$. Let $\mathcal{T} = \{\sigma \bullet \mathfrak{h} \mid \sigma \in H\}$ and $R = \bigcup(\mathcal{R}(\mathcal{T}))$. Drawing N times from $-$ to construct H , the probability that $v \notin R$ is $(1 - 1/|o|)^N$. Consider the random quantity $Z = |V \setminus R|$ with expected value $\mathbb{E}(Z) = \sum_{o \in O} |o|(1 - 1/|o|)^N$. Since $1 - x < e^{-x}$ for $x > 0$, we obtain $\mathbb{E}(Z) < \sum_{o \in O} |o|e^{(-N/|o|)} \leq Se^{-N/M} = Se^{-\ln(S)}$. From $xe^{-\ln(x)} = 1$, it follows that $\mathbb{E}(Z) < 1$. Since $Z \in \mathbb{N}$, then probability $Z = 0$ is more than 0, and thus N transversals of O suffice to cover V . \square

We conjecture that a much smaller number of transversals is actually required, and in all our experimentation have found that the following conjecture is not violated:


Conjecture 1. *Let \dashv be a group acting on a set V (e.g., $\mathcal{D}(\Pi)$). Suppose \triangleright is a transversal of O , a set of orbits induced by the action of \dashv on V . Take $M = \max_{o \in O} |o|$. Where \bullet is composition, there will always be a set of transversals \mathcal{T} with size $\leq M$ such that*

- (i) *Each element $\triangleright' \in \mathcal{T}$ satisfies $\triangleright' = \sigma \bullet \triangleright$ for some $\sigma \in \dashv$, and*
- (ii) *For every $u_1 \in V$, it has an element \triangleright' that covers u_1 .*

We are left to formulate and study a general problem of instantiation, treating it as one of finding graph transversals.

Definition 33 (Consistent Graph Transversals). *For a graph \mathcal{G} and a partition P of $V(\mathcal{G})$, a transversal \triangleright of P is consistent with \mathcal{G} when an edge between p_1 and p_2 in $E(\mathcal{G}/P)$ exists iff there is an edge between $\triangleright(p_1)$ and $\triangleright(p_2)$ in $E(\mathcal{G})$.*

Example 26. *Take \mathcal{G} to be the hexagon that we have illustrated twice below in order to depict partitions $P_1 = \{\{a, b\}, \{c, d\}, \{e, f\}\}$ and $P_2 = \{p' = \{a, b, f\}, p'' = \{c, d, e\}\}$ on the left and right respectively.*



The vertices of \mathcal{G}/P_1 (a 3-clique) and \mathcal{G}/P_2 (a 2-clique) are indicated above by red outlines. There is no consistent transversal of P_1 (LHS) because there is no 3-clique subgraph of \mathcal{G} with one vertex from each set in P_1 . For \mathcal{G}/P_2 (RHS), $\triangleright_1 = \{p' \mapsto f, p'' \mapsto e\}$ is a transversal of P_2 consistent with \mathcal{G} , because the subgraph of \mathcal{G} induced by $\{e, f\}$ is a 2-clique with one vertex from each of p' and p'' .

A transversal of a well-formed \hat{P} consistent with $\mathcal{G}(\Pi)$ is isomorphic to an instantiation of $\Pi/vs_{1..n}$ consistent with Π . Accordingly, Problem 1 is an instance of the following.

Problem 2. *Given \mathcal{G} , a partition P of $V(\mathcal{G})$ and $u_1 \in V(\mathcal{G})$, is there a consistent transversal of P which covers u_1 ?*

We now derive the complexity of Problem 2. We first show that the following problem is NP-complete, and use that result to show that Problem 2 is also NP-complete.

Problem 3. *Given graph \mathcal{G} and a partition P of $V(\mathcal{G})$, is there a transversal of P consistent with \mathcal{G} ?*

Lemma 7. *Problem 3 is NP-complete.*

Proof. Membership in NP is given because a transversal's consistency can clearly be tested in polynomial-time. We then show the problem is NP-hard by demonstrating a polynomial-time reduction from SAT. Consider SAT problems given by formulae φ in conjunctive normal form. Assume every pair of clauses is mutually satisfiable—i.e., for clauses $c_1, c_2 \in \varphi$, for two literals $\ell_1 \in c_1$ and $\ell_2 \in c_2$ we have $\ell_1 \neq \neg \ell_2$ (when this assumption is violated, unsatisfiability can be decided in polynomial-time). Consider the graph $\mathcal{G}(\varphi)$, where $V(\mathcal{G}(\varphi)) = \{u_{1\ell_c} \mid c \in \varphi, \ell \in c\}$,

and $E(\mathcal{G}(\varphi)) = \{\{u_{1\ell_1 c_1}, u_{1\ell_2 c_2}\} \mid c_1, c_2 \in \varphi, \ell_1 \in c_1, \ell_2 \in c_2, \ell_1 \neq \neg \ell_2\}$. Now let $P = \{\{u_{1\ell c} \mid \ell \in c\} \mid c \in \varphi\}$. Note that P is a partition of $V(\mathcal{G}(\varphi))$ and every set in P corresponds to a clause in φ . Because all the clauses in φ are mutually satisfiable, the quotient $\mathcal{G}(\varphi)/P$ is a clique. Now we prove there is a model for φ iff there is a transversal of P consistent with $\mathcal{G}(\varphi)$. (\Rightarrow) A model \mathcal{M} has the property $\forall c \in \varphi. \exists \ell \in c. \mathcal{M} \models \ell$. Due to the correspondence between sets in P and clauses in φ , a transversal \mathfrak{h} of P can be constructed by selecting one satisfied literal from each clause. Based on a model, \mathfrak{h} will never select conflicting literals. All members of $\mathcal{R}(\mathfrak{h})$ are pairwise connected, so \mathfrak{h} is consistent with $\mathcal{G}(\varphi)$ as required. (\Leftarrow) By definition \mathfrak{h} is consistent with $\mathcal{G}(\varphi)$, so the subgraph of \mathcal{G} induced by $\mathcal{R}(\mathfrak{h})$ is a clique. Let \mathcal{L} be literals corresponding to $\mathcal{R}(\mathfrak{h})$, and note that its elements are pairwise consistent. A model \mathcal{M} for φ is constructed by assigning u_1 to \top where u_1 occurs positively in \mathcal{L} , and to \perp otherwise. \square

Theorem 16. *Problem 2 is NP-complete.*

Proof. Consistency of a transversal \mathfrak{h} is clearly polynomial-time testable, as is the coverage condition that $u_1 \in \mathcal{R}(\mathfrak{h})$. Thus, we have membership in NP. NP-hardness follows from the following reduction from Problem 3 (P3). Taking P and \mathcal{G} as inputs to P3, construct the graph \mathcal{G}' , where: $V(\mathcal{G}') = V(\mathcal{G}) \cup \{u_1\}$, u_1 an auxiliary vertex; and $E(\mathcal{G}') = E(\mathcal{G}) \cup \{\{u_1, u_2\} \mid u_2 \in V(\mathcal{G})\}$. There is a solution to P3 with inputs \mathcal{G} and P iff P2 is soluble with inputs \mathcal{G}' , $P \cup \{\{u_1\}\}$ ($= P'$) and u_1 . (\Rightarrow) If \mathfrak{h} is a transversal of P consistent with \mathcal{G} , then $\mathfrak{h} \cup \{\{u_1\} \mapsto u_1\}$ ($= \mathfrak{h}'$) is a transversal of P' . As u_1 is fully connected, \mathfrak{h}' is consistent with \mathcal{G}' . \mathfrak{h}' is a solution to P2 because $u_1 \in \mathcal{R}(\mathfrak{h}')$. (\Leftarrow) If \mathfrak{h}' is a solution to P2 with inputs \mathcal{G}' , P' and u_1 , then \mathfrak{h}' is a transversal of P' , and also P . All edges in $E(\mathcal{G}')$ are in $E(\mathcal{G})$, with the exception of those adjacent u_1 , and since \mathfrak{h}' is consistent with \mathcal{G}' , then \mathfrak{h}' is consistent with \mathcal{G} . Thus, \mathfrak{h}' solves P3. \square

In concluding, we note that the NP-hard canonicalisation problem—the optimisation problem posed at each state encountered by orbit search—is not known to be in NP. Our above results imply that exploitation of symmetry via instantiation poses a decision problem in NP.

4.6 Concrete Plan from Quotient Plan

Having computed isomorphic subproblems Δ that cover the goals of Π , a concrete plan is a concatenation of plans for members of Δ . However, this is not always straightforward.

Example 27. *Take Π_1 , $vs_{1..n}$ and \mathfrak{h} from Example 24. Note $\mathfrak{h}' = \{p_1 \mapsto v_3, p_2 \mapsto v_2, p_3 \mapsto v_5\}$ is also an instantiation of $\Pi_1/vs_{1..n}$ consistent with Π_1 . Observe that $\{\mathfrak{h}(\Pi_1/vs_{1..n}), \mathfrak{h}'(\Pi_1/vs_{1..n})\}$ covers the concrete goal $G = \{\overline{v_4}, \overline{v_5}\}$. A plan for $\Pi_1/vs_{1..n}$ is $\vec{\pi}' = (\{p_1, p_2\}, \{\overline{p_3}, \overline{p_1}\})$ and its two instantiations are $\mathfrak{h}(\vec{\pi}') = (\{v_3, v_1\}, \{\overline{v_4}, \overline{v_3}\})$ and $\mathfrak{h}'(\vec{\pi}') = (\{v_3, v_2\}, \{\overline{v_5}, \overline{v_3}\})$. Concatenating $\mathfrak{h}(\vec{\pi}')$ and $\mathfrak{h}'(\vec{\pi}')$ in any order does not solve Π_1 because both plans require v_3 initially, but do not establish it. To overcome this issue in practice, before we solve $\Pi_1/vs_{1..n}$ we augment its goal with the assignment $p_1 \mapsto \top$.*

We now give conditions under which concatenation is valid, and detail the quotient-goal augmentation step we use to ensure validity in practice. We will use the notion of projection, writing $Y \downarrow^X$ to mean $Y \downarrow_{\mathcal{D}(X)}$, for the common case where we wish to project with respect to all the variables in the domain of an object.

Definition 34 (Needed Assignments, Preceding Problem). *Needed assignments, $\mathcal{N}(\Pi)$, are assignments in the preconditions of actions and goal conditions that also occur in I , i.e., $\mathcal{N}(\Pi) =$*

$(\text{pre}(\delta) \cap I) \cup (G \cap I)$. Problem Π_1 is said to precede Π_2 , written $\Pi_1 \triangleleft \Pi_2$, iff

$$G_1 \downarrow^{\mathcal{N}(\Pi_2)} = (I_2 \downarrow^{\Pi_1}) \downarrow^{\mathcal{N}(\Pi_2)} \wedge G_1 \downarrow^{\Pi_2} = G_2 \downarrow^{\Pi_1}$$

In words,

- (i) The needed assignments of Π_2 which a plan for Π_1 could modify occur in G_1 , and
- (ii) G_2 contains all the assignments in G_1 which a plan for Π_2 might modify.

Example 28. Consider Π_1 and Π_2 from Examples 24 and 25 respectively. $\mathcal{N}(\Pi_1) = \{v_3, v_1, v_2\}$ and $\mathcal{N}(\Pi_2) = G_1 = \{\bar{v}_4, \bar{v}_5\}$. Since $G_1 \downarrow^{\mathcal{N}(\Pi_2)} = G_1$, $(I_2 \downarrow^{\Pi_1}) \downarrow^{\mathcal{N}(\Pi_2)} = G_1$, $G_1 \downarrow^{\Pi_2} = G_1$ and $G_2 \downarrow^{\Pi_1} = G_1$, we have $\Pi_1 \triangleleft \Pi_2$.

Writing $\vec{\pi}_i \frown \vec{\pi}_j$ for concatenation of plans $\vec{\pi}_i, \vec{\pi}_{i+1} \dots \vec{\pi}_j$, a simple inductive argument gives the following:

Lemma 8. Let $\Pi_1 \dots \Pi_N$ be a set of problems satisfying $\Pi_j \triangleleft \Pi_k$ for all $j < k \leq N$. If for all $1 \leq i \leq N$, state x satisfies $I_i \subseteq x$, $\vec{\pi}_i$ solves Π_i and $\text{sat-pre}(\mathcal{N}(\Pi_i), \vec{\pi}_i)$ holds, then $\text{ex}(x, \vec{\pi}_1 \frown \dots \frown \vec{\pi}_N) \downarrow^{G_i} = G_i$.

Proof. By induction on N . First, note $\text{ex}(x, \vec{\pi}_i) \downarrow^{G_i} = \text{ex}(I_i, \vec{\pi}_i) \downarrow^{G_i}$. Our proof is by induction on N , and the base case $N = 1$ is trivial. We now show that if the theorem is true for N , then the theorem also holds for $N + 1$. Below, let $x' = \text{ex}(x, \vec{\pi}_1 \frown \dots \frown \vec{\pi}_N)$.

Note $\mathcal{N}(\Pi_{N+1}) \subseteq x'$, because for each i :

- (i) assignments modified by $\vec{\pi}_i$ whose variables are the subject of a needed assignment from Π_{N+1} are in G_i , because $\Pi_i \triangleleft \Pi_{N+1}$ and thus $G_i \downarrow^{\mathcal{N}(\Pi_{N+1})} = (I_{N+1} \downarrow^{\Pi_i}) \downarrow^{\mathcal{N}(\Pi_{N+1})}$
- (ii) the goals of any Π_i are in x' because the inductive hypothesis gives $x' \downarrow^{G_i} = G_i$
- (iii) any assignments in I_{N+1} not modified by execution $\vec{\pi}_1 \frown \dots \frown \vec{\pi}_N$ from x are in x' , because $I_{N+1} \subseteq x$.

Because $\vec{\pi}_{N+1}$ is a plan for Π_{N+1} , $\text{ex}(x', \vec{\pi}_{N+1}) \downarrow^{G_{N+1}} = G_{N+1}$.

Since $\Pi_i \triangleleft \Pi_{N+1}$, any assignment in x' that is in some G_i is also in $\text{ex}(x', \vec{\pi}_{N+1})$. The inductive hypothesis gives $G_i = x' \downarrow^{G_i}$ thus $G_i = \text{ex}(x', \vec{\pi}_{N+1}) \downarrow^{G_i}$. Because $\text{ex}(x, \vec{\pi}_1 \frown \dots \frown \vec{\pi}_{N+1}) = \text{ex}(x', \vec{\pi}_{N+1})$ we have $\text{ex}(x, \vec{\pi}_1 \frown \dots \frown \vec{\pi}_{N+1}) \downarrow^{G_i} = G_i$. \square

Take problem Π and a set of problems $\mathbf{\Pi}$, we say that $\mathbf{\Pi}$ covers Π iff $\forall g \in G. \exists \Pi' \in \mathbf{\Pi}. g \in G'$ and $\forall \Pi' \in \mathbf{\Pi}. \Pi' \subseteq \Pi$. I.e., every goal from Π is stated in one or more members of $\mathbf{\Pi}$, a set of subproblems of Π .

Theorem 17. Consider a set $\Pi_1 \dots \Pi_N$ of problems that covers Π , satisfying $\Pi_j \triangleleft \Pi_k$ for all $j < k \leq N$. For $1 \leq i \leq N$ let $\vec{\pi}_i$ be a plan for Π_i , then $\text{rem-condless}(\mathcal{N}(\Pi_1), \vec{\pi}_1) \frown \dots \frown \text{rem-condless}(\mathcal{N}(\Pi_N), \vec{\pi}_N)$ is a plan for Π , where for an action sequence $\vec{\pi}$ and a state x , $\text{rem-condless}(x, \vec{\pi})$ denotes $\vec{\pi}$, but without the actions whose preconditions are not satisfied during executing $\vec{\pi}$ from x .

This theorem follows directly from the fact that the set of problems covers Π and from Lemma 8.

We now address the question: When can plans for a set of isomorphic subproblems be concatenated to provide a concrete plan? We provide sufficient conditions in terms of the concepts of *common* and *sustainable* variables.

Definition 35 (Common Variables). For a set of instantiations Δ , the set of common variables, written $\bigcap_v \Delta$, comprises variables in $\bigcup_{\mathfrak{h} \in \Delta} \mathcal{R}(\mathfrak{h})$ that occur in the ranges of more than one member of Δ .

Definition 36 (Sustainable Variables). A set of variables vs is sustainable in a problem Π iff $I \downarrow_{vs} = G \downarrow_{vs}$.

Theorem 18. Take problem Π , partition $vs_{1..n}$ of $\mathcal{D}(\Pi)$ where the quotient $\Pi / vs_{1..n}$ ($= \Pi'$) is well-defined, with solution $\vec{\pi}'$, and consistent instantiations Δ . Suppose $\{\mathfrak{h}(\Pi') \mid \mathfrak{h} \in \Delta\}$ ($= \mathbf{\Pi}$) covers Π , and $\mathcal{Q}(\bigcap_v \Delta) \cap \mathcal{D}(\mathcal{N}(\Pi'))$ —i.e., based on Definition 30, the orbits of common variables from needed assignments—are sustainable in Π' . Then any concatenation of the plans $\{\text{rem-condless}(\mathcal{N}(\mathfrak{h}(\Pi')), \mathfrak{h}(\vec{\pi}')) \mid \mathfrak{h} \in \Delta\}$ solves Π .

Proof. Identify the elements in Δ by indices in $\{1..|\Delta|\}$. Let $k \in \{1..|\Delta|\}$ and $\Pi_k = \mathfrak{h}_k(\Pi')$, and note $\mathcal{R}(\mathfrak{h}_k) = \mathcal{D}(\Pi_k)$. Take $\mathfrak{h}_i, \mathfrak{h}_j \in \Delta$, where $i, j \in \{1..|\Delta|\}$ and $i \neq j$. We first show that $\Pi_i \triangleleft \Pi_j$. For any $vs \in vs_{1..n}$, $\mathfrak{h}_i(vs) = \mathfrak{h}_j(vs)$ if $\mathfrak{h}_i(vs) \in \mathcal{R}(\mathfrak{h}_i) \cap \mathcal{R}(\mathfrak{h}_j)$. Therefore, $I_i \downarrow^{\Pi_j} = I_j \downarrow^{\Pi_i}$ and $G_i \downarrow^{\Pi_j} = G_j \downarrow^{\Pi_i}$, providing the right conjunct in Definition 34. For the left conjunct, note $\mathcal{D}(\mathcal{N}(\Pi_j)) \subseteq \mathcal{D}(\Pi_j)$ and $\mathcal{D}(\Pi_i) \cap \mathcal{D}(\Pi_j) \subseteq \bigcap_v \Delta$. The sustainability premise— $\mathcal{Q}(\bigcap_v \Delta) \cap \mathcal{D}(\mathcal{N}(\Pi'))$ is sustainable in Π' —then provides that $\mathcal{D}(\Pi_i) \cap \mathcal{D}(\mathcal{N}(\Pi_j))$ is sustainable in Π_i —i.e., $I_i \downarrow^{\mathcal{N}(\Pi_j)} = G_i \downarrow^{\mathcal{N}(\Pi_j)}$. Thus $G_i \downarrow^{\mathcal{N}(\Pi_j)} = (I_j \downarrow^{\Pi_i}) \downarrow^{\mathcal{N}(\Pi_j)}$, and we conclude $\Pi_i \triangleleft \Pi_j$. Since a plan $\mathfrak{h}_k(\vec{\pi}')$ solves $\mathfrak{h}_k(\Pi')$, $(\mathfrak{h}_k(\Pi')).I \subseteq \Pi.I$, therefore as per Theorem 17 a solution to Π is $\text{rem-condless}(\mathcal{N}(\mathfrak{h}_1(\Pi')), \mathfrak{h}_1(\vec{\pi}')) \frown \text{rem-condless}(\mathcal{N}(\mathfrak{h}_N(\Pi')), \mathfrak{h}_N(\vec{\pi}'))$. \square

In practice we take $vs^* = \mathcal{Q}(\bigcap_v \Delta) \cap \mathcal{D}(\mathcal{N}(\Pi / vs_{1..n}))$, and augment the goal of the quotient $\Pi / vs_{1..n}$ by adding $(\Pi / \mathcal{P}).I \downarrow_{vs^*}$. Call the resulting problem Π^q and its solution $\vec{\pi}^q$. Theorem 18 shows that any concatenation of the plans $\{\text{rem-condless}(\mathfrak{h}(\Pi^q), \mathfrak{h}(\vec{\pi}^q)) \mid \mathfrak{h} \in \Delta\}$ solves Π . Thus, the proposed algorithm is sound.

Example 29. Take Π_1 , $vs_{1..n}$, \mathfrak{h} and \mathfrak{h}' from Example 27. There is one common variable, $\{v_3\} = \bigcap_v \Delta$ from the orbit $\{p_1\} = \mathcal{Q}(\bigcap_v \Delta)$. Here $\mathcal{N}(\Pi_1 / vs_{1..n}) = \{p_1, p_2\}$, and the orbit of the common needed variable is $\{p_1\} = \mathcal{Q}(\bigcap_v \Delta) \cap \mathcal{N}(\Pi_1 / vs_{1..n})$. To solve Π_1 via solving $\Pi_1 / vs_{1..n}$, we augment the goal $(\Pi_1 / vs_{1..n}).G$ with the assignment to p_1 in $(\Pi_1 / vs_{1..n}).I$. The resulting problem, Π_1^q , is equal to $\Pi_1 / vs_{1..n}$ except that it has the goals $\Pi_1^q.G = \{p_1, \bar{p}_3\}$. A plan for Π_1^q is $\vec{\pi}^q = (\{p_1, p_2\}, \{\bar{p}_3, \bar{p}_1\})(\emptyset, \{p_1\})$, and two instantiations of it are $\mathfrak{h}(\vec{\pi}^q) = (\{v_3, v_1\}, \{\bar{v}_4, \bar{v}_3\})(\emptyset, \{v_3\})$ and $\mathfrak{h}'(\vec{\pi}^q) = (\{v_3, v_2\}, \{\bar{v}_5, \bar{v}_3\})(\emptyset, \{v_3\})$. Concatenating $\mathfrak{h}(\vec{\pi}^q)$ and $\mathfrak{h}'(\vec{\pi}^q)$ in any order solves Π_1 .

4.7 Experimental Results

Implemented in C++,¹ our approach uses: NAUTY\TRACES by McKay and Piperno (2014) to calculate symmetries; Z3 by De Moura and Bjørner (2008) to find isomorphic subproblems; and the initial-plan search by LAMA by Richter and Westphal (2010) as the *base planner*. We limit *base planner* runtimes to 30 minutes.

By running our algorithm, we obtained a set of benchmarks with soluble descriptive quotients whose solutions can be instantiated to concrete plans. That set includes 439 problems, from 16 IPC benchmarks and 4 benchmarks from Porco et al. (2013). In all our experimentation we identified 120 instances where we were able to confirm that the descriptive quotient does not have a solution

¹Found at the public repository bitbucket.org/MohammadAbdulaziz/planning.git in the directory codeBase.




Figure 4.4: (a) Scatter-plot comparing the number of *actions* in problem posed by descriptive quotient vs. concrete problem, with red line plotting $f(x) = x$. (b) Scatter-plot comparing the number of *state variables* in problem posed by descriptive quotient vs. concrete problem, with red line plotting $f(x) = x$.

and the concrete problem does. Figure 4.4 plots the sizes in terms of both, the number of actions and the number of state variables, of concrete and quotient problem descriptions. The plotted data shows that the descriptive quotient can be much smaller than its concrete counterpart. In just over 15% of instances the quotient has less than half the number of actions. Here, we also analyzed what aspects of our approach are computationally expensive. In 79% of cases 99% of the runtime of our approach is executing the base planner. In 96% of cases 95% of the runtime of our approach is executing the base planner. Overall, 3% of time is spent in instantiation and finding chromatic numbers.

We examined where our approach is comparatively scalable and fast compared to the base planner. For 430 instances where the base planner and our approach are successful, Figure 4.5a displays the speedup factors where planning *via* the descriptive quotient is comparatively fast. Overall, for 68% of instances our approach is comparatively fast, and in 15% planning *via* the quotient is at least twice as fast. With few exceptions, instances where our approach is at least twice as fast are from GRIPPER, HIKING, MPRIME, MYSTERY, PARCPRINTER, PIPESWORLD, TPP and VISITALL. In 5 problems from HIKING, 2 from PARCPRINTER, 1 from TETRIS and 1 from KCOLOURABILITY, the base planner cannot solve the concrete problem, but can solve the quotient. Planning *via* the quotient can also be slow, primarily due to the extra cost of finding symmetries,² and because LAMA is heuristic and occasionally finds that the quotient poses a more challenging problem. Figure 4.5b provides the dual to Figure 4.5a, showing cases where planning directly for the concrete problem is comparatively fast. This is the case in 32% instances, and indeed in 2% our approach is at least twice as slow.

Finally, it is worth highlighting the difference between searching for a plan in the state space of the descriptive quotient versus approximate orbit search—i.e., searching for a plan in an approximation of the quotient state space—as is done in the state-of-the-art techniques for exploiting symmetries in planning.³ In comparing those approaches, we measure the number of states expanded using a breadth-first search. We consider the IPC instances GRIPPER-20 and MPRIME-21, the largest instances from domains GRIPPER and MPRIME reported solved by Pochter et al. (2011) using breadth-first search. Those authors report the number of states expanded to be 60.5K and 438K, respectively. Using the same search to solve the problem posed by our descriptive quotient,

²If symmetries are given as part of the problem description, or if one resorts to more heuristic methods of discovering them, such as those described by Guere and Alami (2001); Fox and Long (2002), this burden is relieved.

³Such a comparison is admittedly unfair, because the problem posed by a descriptive quotient is not a bisimulation of the concrete problem.

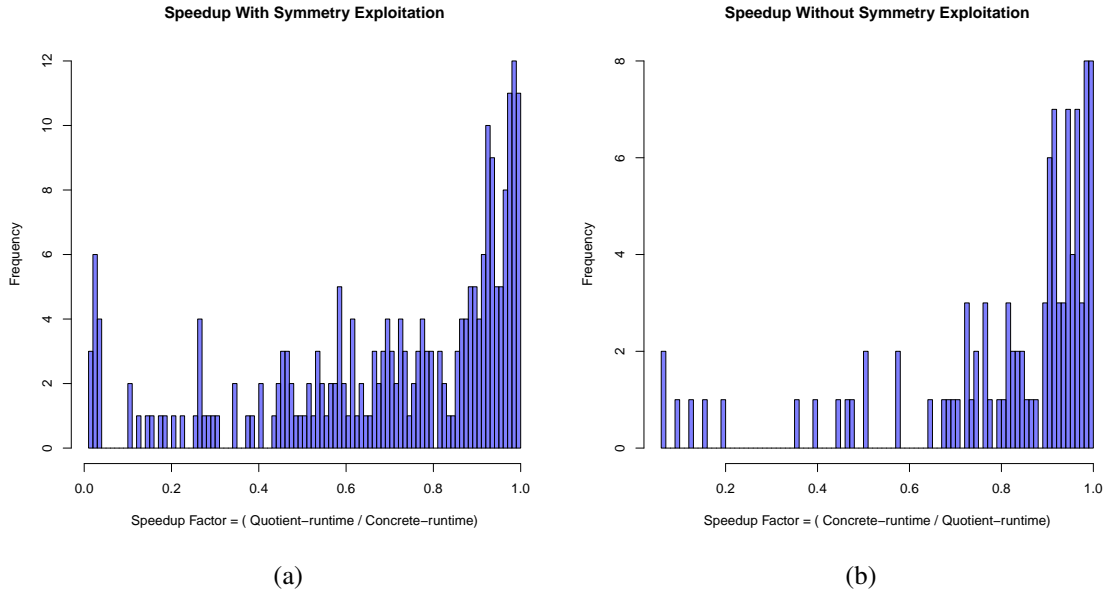


Figure 4.5: (a) Plots histogram of speedup factors experienced when planning **with** symmetry exploitation, reporting only for instances where planning via the quotient is faster, and only for instances where the base planner takes > 5 seconds. (b) Plots histogram of speedup factors experienced when planning **without** symmetry exploitation, reporting only for instances where planning for the concrete problem directly is faster, and only for instances where the base planner takes > 5 seconds.

the base planner expands 6 and 203K states, respectively.

4.8 Conclusions and Future Work

We introduced a new compositional approach for AI planning that is based on a new abstraction, the *descriptive quotient*, as well as exploiting state variable symmetries. A concrete plan is obtained by concatenating instantiations of the plan for the descriptive quotient. Compared to mainstream symmetry breaking techniques this has multiple advantages. Firstly, we do not have to solve the NP-hard constructive orbit problem for every orbit of symmetric states. Secondly, the size of the descriptive quotient’s state space can be exponentially smaller than that of the quotient system, in which a plan is searched for, by other symmetry breaking algorithms. However, effectively this approach breaks a specific type of symmetries: *repetitive symmetries*, which is when the planning problem is constituted of multiple isomorphic subproblems. Also it is incomplete, as the descriptive quotient may not have a plan, while the concrete problem is soluble. Nonetheless, we experimentally show that this approach is successful in 70% of standard international planning competition benchmarks with symmetries, where it brings significant reduction in plan search time.

We suppose a fruitful research direction is exploring descriptive quotients: to develop heuristics, to characterise tractable classes, and to develop bounding methods. Also an interesting question is how much does the quality (in terms of length and cost) of the plans computed by this technique compare to that of plans computed by other symmetry breaking techniques. We suppose that this question is very strongly related to either proving or disproving Conjecture 1

Formalisation

In order to formally verify the correctness of our compositional algorithms we build a library of formal proofs on factored transition systems in the Higher Order Logic (HOL) interactive theorem prover HOL4 by Slind and Norrish (2008). There is a rich body of previous formalisation of concepts related to transition systems in different logics and theorem proving systems. A lot of that work (e.g. Paulson (2015) and Esparza et al. (2013)) focuses on classical textbook results on finite automata and reachability of states within those automata in the setting of model checking LTL formulae. However, we believe that our work is the first to formalise transition systems in their factored representation, and to focus on the topological properties of their state spaces, in addition to reachability within them. In this chapter we discuss our formalisation of factored transition systems and how we used it to formally verify some of the main algorithms we discussed earlier. The general organisation of the library is shown in Figure 5.1 and the size and description of each theory is in Table 5.1.

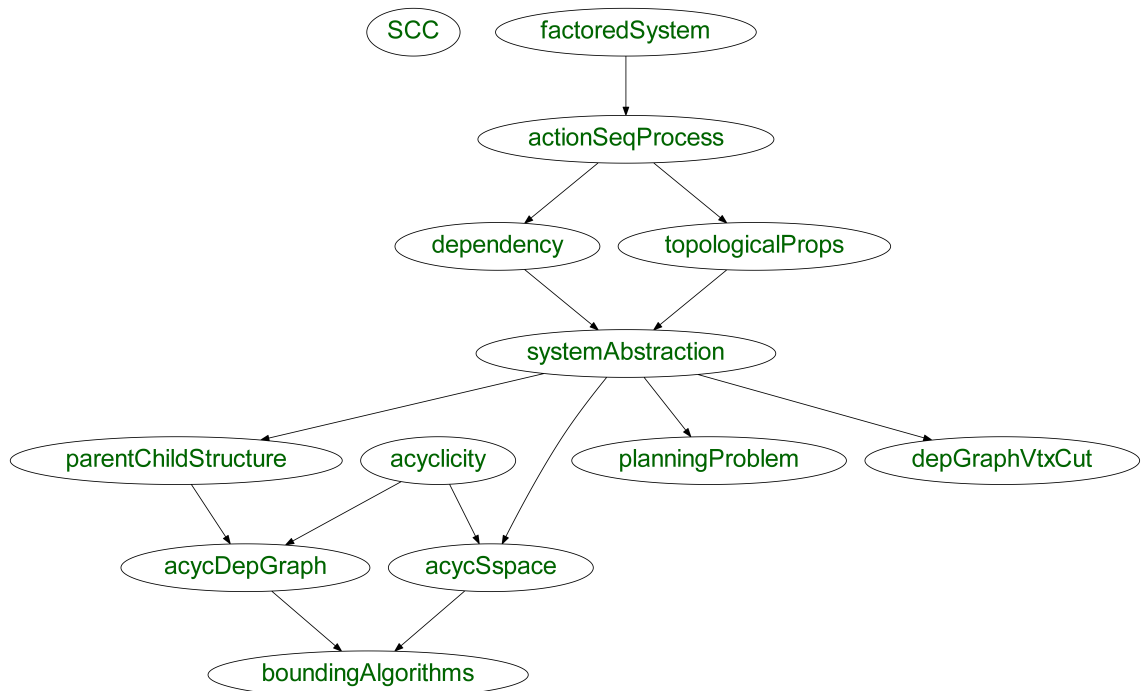


Figure 5.1: The organisation of the different theories concerning factored transition systems. An edge from one theory to another indicates the dependence of the latter on the former.

Theory	Size (LOC)	Description
actionSeqProcess	432	Theory on different functions that process action sequences, e.g. remove actions without effects.
acycDepGraph	1587	Theory related to the top-down algorithm.
acyclicity	130	Theory regarding acyclic digraphs represented as topologically sorted lists.
acycSspace	346	Theory related to the S-algorithm that exploits acyclicity in the state space.
boundingAlgorithms	644	Theory related to the hybrid algorithm.
dependency	83	Basic results related to variable dependency.
depGraphVtxCut	582	Results related to vertex cutting in the dependency graph.
factoredSystem	1335	Basic results and definitions related to factored transition systems.
instantiation	1007	Results related to instantiation of planning problems.
invariantsPlusOne	300	Theory related to the cardinality of sets of states with SAS+ like invariant properties.
invariantStateSpace	121	A bound on the size of a state space of a planning problem that has a SAS+ representation derived invariant property.
parentChildStructure	1304	Formalisation of bound compositionality in the parent-child structure and the stitching function.
planningProblem	994	Formalisation of planning problems on top of factored transition systems.
prodValidtd	451	Verification of the fact that the product of some projections' traversal diameters is an upper bound on the concrete system's traversal diameter.
SCC	46	Basic results related to strongly connected components.
SCCsystemAbstraction	1384	Results related to computing system abstractions based on computing strongly connected components of dependency graphs.
stateSpaceProduct	358	Defining the state space product operator and some basic facts about it.
systemAbstraction	1135	Definition of different abstraction concepts like projection and snapshotting, and verifying some basic facts about them.
topologicalProps	879	Definition of the diameter, sublist diameter, recurrence diameter, and the traversal diameter and verifying basic theory about them.

Table 5.1: A table showing the sizes of different theories and their content.

HOL4 Notation All statements appearing with a turnstile (\vdash) are HOL4 theorems, automatically pretty-printed to \LaTeX .

5.1 Related Work

Interactive theorem proving systems (e.g. Gordon and Melham (1993); Paulson (1994); Bertot and Castéran (2004)) have been used in mechanisation and verification in very diverse fields. Fundamental results in different branches of mathematics have been verified. For example Dufourd (2009) formalised a discretised form of the Jordan Curve theorem. In mathematical analysis, Harrison (2007) formalised Cauchy’s integral theorem, and Abdulaziz and Paulson (2016) formalised Green’s theorem. In Bezem and Hendriks (2008) a coherent logic based theorem prover was implemented and used to mechanise Hessenberg’s theorem. Also the work by McKinna and Pollack (1999) shows a mechanisation of a simplified version of beta reduction in Lambda Calculus. Gamboa (2009) presents a mechanisation of the Powerlist Algebra, which is a data structure used for representing concurrent recursive algorithms as well as proving their correctness. Synthetic domain theory was formalised in Taylor (1991).

Many sophisticated algorithms have also been mechanised in the literature. Techniques for theorem proving such as SLD-resolution have been mechanised in Jaume (1999), where the theorem prover Coq was used. The work in Kamareddine and Qiao (2003) presents a formalisation of the termination proof of two styles of substitution calculi. Manolios and Vroon (2005) provided libraries for doing basic tasks in ordinal arithmetic, and mechanised them using ACL2. For propositional satisfiability, Marić (2009); Blanchette et al. (2016) formalised SAT solvers from the basic DPLL solver up to modern solvers which have features such as clause learning or backjumping and other features.

Different work, like Klein et al. (2009); Hales et al. (2011); Gonthier (2008) , show that the scale of mechanisation projects is continuously growing. Bauer et al. (2013) provide a mechanisation of data flow analysis for the purpose of building more secure software systems. In Gonthier (2008) and Hales et al. (2011), the proof Kepler’s conjecture and the four colour theorem, respectively, are verified, which are major mathematical results.

Most relevant to our work are formalisations of automata theory and their reachability properties, in the context of model checking. Textbook results in automata theory were formalised in many approaches. For example, the Myhill-Nerode theorem was formalised in intuitionistic logics by Constable et al. (2000) and Doczkal et al. (2013). It was also formalised by Paulson (2015) in Isabelle/HOL. He used hereditarily finite sets as the type of states, unlike our formalisation which uses finite maps. Furthermore, Wu et al. (2011) formally prove that theorem using a formalisation of regular expressions. Results and algorithms related to reachability in automata were formalised by Sprenger (1998); Schimpf et al. (2009); Esparza et al. (2013), where the main goal of those authors was to obtain formally verified model checking algorithms and implementations.

5.2 Factored Transition Systems in HOL4

We formalise propositionally factored transition systems by first defining states to be finite maps $\alpha \mapsto \beta$ from polymorphic type α to polymorphic type β .¹ An action is a pair of such states $(\alpha \mapsto \beta) \times (\alpha \mapsto \beta)$, and a factored transition system is a set of actions $(\alpha \mapsto \beta) \times (\alpha \mapsto \beta) \rightarrow bool$. We define the set of valid states and valid action sequences as follows:

¹If we would model STRIPS or SMV transition systems, β would be instantiated with *bool*.

HOL4 Definition 1 (Factored System).

$$\mathbb{U}(\delta) = \{x \mid \mathcal{D}(x) = \mathcal{D}(\delta)\}$$

$$\delta^* = \{\vec{\pi} \mid \text{set } \vec{\pi} \subseteq \delta\}$$

Action execution and action sequence execution are defined as follows:

HOL4 Definition 2 (Execution).

$$\text{state-succ } x (p, e) = \text{if } p \sqsubseteq x \text{ then } e \uplus x \text{ else } x$$

$$\text{ex}(x, \pi :: \vec{\pi}) = \text{ex}(\text{state-succ } x \ \pi, \vec{\pi})$$

$$\text{ex}(x, []) = x$$

The result of executing an action (p, e) on a state x depends on whether the preconditions of the action are satisfied by the state or not, which is modelled by the $p \sqsubseteq x$ relation. If the state satisfies the preconditions then the state resulting from the execution is the same as the original state, but amended by the effects of the executed action, otherwise the result of the execution is the same as the original state. The way we model amending the state by an action's effect is by using the finite map union operation $e \uplus x$. We note that this finite map union operation is not commutative, where it gives precedence to the assignments of its first argument. For an action sequence $\vec{\pi}$, the execution semantics are lifted in a straightforward way. A sanity check of our execution semantics is the following theorem, which states that the result of executing a valid action sequence on a valid state is also a valid state.

$$\vdash \vec{\pi} \in \delta^* \wedge x \in \mathbb{U}(\delta) \Rightarrow \text{ex}(x, \vec{\pi}) \in \mathbb{U}(\delta)$$

Note that we do not restrict the codomain of states to be *bool*. This is because a lot of the theory we develop applies to factored systems, regardless of the codomain of the state, so for example, a lot of the theory developed applies to hybrid systems.

5.2.1 Topological Properties

In HOL4, we define the diameter in the following way:

HOL4 Definition 3 (Diameter).

$$d(\delta) = \max \{ \min (\Pi^d (x, \vec{\pi}, \delta)) \mid x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \}$$

where Π^d is defined as

$$\Pi^d (x, \vec{\pi}, \delta) = \{ |\vec{\pi}'| \mid \text{ex}(x, \vec{\pi}') = \text{ex}(x, \vec{\pi}) \wedge \vec{\pi}' \in \delta^* \}$$

In the above definition, we first define the function Π^d which returns, for a state x , an action sequence $\vec{\pi}$, and a factored system δ , the set of lengths of all valid action sequences in δ^* that would yield the same execution result as $\vec{\pi}$, if executed on x . Accordingly, the smallest member of $\Pi^d (x, \vec{\pi}, \delta)$ would be the length of the shortest action sequence equivalent to $\vec{\pi}$ in terms of execution on x . The diameter then would be the largest such length of the shortest equivalent action sequence for all pairs of states and valid action sequences. In a similar approach the sublist diameter is defined as follows:

HOL4 Definition 4 (Sublist Diameter). *We first define the sublist relation between lists as follows:*

$$\begin{aligned}
[] \preceq l_1 &\iff \mathbf{T} \\
h::t \preceq [] &\iff \mathbf{F} \\
x::l_1 \preceq y::l_2 &\iff x = y \wedge l_1 \preceq l_2 \vee x::l_1 \preceq l_2
\end{aligned}$$

Based on that the sublist diameter is defined as:

$$\ell(\delta) = \max \{ \min \Pi^{\preceq}(x, \vec{\pi}) \mid x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \}$$

where Π^{\preceq} is defined as

$$\Pi^{\preceq}(x, \vec{\pi}) = \{ |\vec{\pi}'| \mid \text{ex}(x, \vec{\pi}') = \text{ex}(x, \vec{\pi}) \wedge \vec{\pi}' \preceq \vec{\pi} \}$$

The way we define the sublist diameter resembles the way we defined the diameter. We note however that in Π^{\preceq} we need not add the condition on the equivalent action sequences to be valid action sequences from δ^* since this is implied by the fact that they are also sublists of the given action sequence $\vec{\pi}$.

We define the recurrence diameter and the traversal diameter in a significantly different approach, since they are both best specified in terms of paths of states, and because in both of them we do not need to nest \max and \min . The recurrence diameter is defined as follows:

HOL4 Definition 5 (recurrence Diameter).

$$rd(\delta) = \max \{ |p| - 1 \mid \text{valid-path } \delta \ p \wedge \text{ALL-DISTINCT } p \}$$

where

$$\begin{aligned}
\text{valid-path } \delta \ [] &\iff \mathbf{T} \\
\text{valid-path } \delta \ [x] &\iff x \in \mathbb{U}(\delta) \\
\text{valid-path } \delta \ (x_1::x_2::\text{rest}) &\iff \\
&x_1 \in \mathbb{U}(\delta) \wedge (\exists \pi. \pi \in \delta \wedge \text{ex}(x_1, [\pi]) = x_2) \wedge \\
&\text{valid-path } \delta \ (x_2::\text{rest})
\end{aligned}$$

The predicate `valid-path` indicates that a certain list of states can be an execution trace of a valid action sequence in the given system, and the predicate `ALL-DISTINCT` indicates that all members of a certain list are distinct.

Lastly, we define the traversal diameter as follows:

HOL4 Definition 6 (Traversal Diameter).

$$td(\delta) = \max \{ |\text{ss}(x, \vec{\pi})| - 1 \mid x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \}$$

where

$$\text{ss}(x, \vec{\pi}) = \text{set}(\vec{\text{ss}}(x, \vec{\pi}))$$

and

$$\begin{aligned}
\vec{\text{ss}}(x, \pi::\vec{\pi}) &= x::\vec{\text{ss}}(\text{state-succ } x \ \pi, \vec{\pi}) \\
\vec{\text{ss}}(x, []) &= [x]
\end{aligned}$$

In the above definition the functions $\vec{\text{ss}}$ and ss return the list of states and the set of states, respectively, traversed by executing an action sequence on a state.

Since we are mainly interested in formally verifying the validity of HYB, we only formalised the different upper-bounding relations between different topological properties, leaving out the exponential separations between them. The formalised bounding relations are expressed in the following theorem.

HOL4 Theorem 1.

$$\vdash \text{FINITE } \delta \Rightarrow d(\delta) \leq \ell(\delta) \wedge \ell(\delta) \leq rd(\delta)$$

We note that a sufficient condition for the different topological properties to exist, δ has to be finite and the codomain of the different states has to be *bool*, which is why those conditions exist in the theorem above. This guarantees that the argument sets to the functions *max* and *min* are finite. This is not the only approach to guarantee that the topological properties are well defined. However, it suits our needs since it models the factored systems on which we applied our algorithms. As sanity checks to verify that our definitions of the different topological properties have the desired semantics, we prove the following theorems for the sublist diameter and the traversal diameter. The first theorem says that any valid action sequence has a sublist of it that achieves the same execution outcome, and whose length is bounded by the sublist diameter. The second one says that executing any valid action sequence on a valid state traverses no more states than the traversal diameter.

$$\vdash \text{FINITE } \delta \wedge x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \Rightarrow \\ \exists \vec{\pi}'. \text{ex}(x, \vec{\pi}) = \text{ex}(x, \vec{\pi}') \wedge \vec{\pi}' \preceq \vec{\pi} \wedge |\vec{\pi}'| \leq \ell(\delta)$$

$$\vdash \text{FINITE } \delta \wedge x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \Rightarrow |\text{ss}(x, \vec{\pi})| - 1 \leq td(\delta)$$

Also, we prove the following theorems to aid us in deriving the compositional upper bounds of both the sublist diameter and the traversal diameter. Both theorems state sufficient conditions for upper-bounding the two topological properties. For the sublist diameter, if for some constant k , there is a sublist of every valid action sequence that achieves the same execution outcome whose length is bounded by k , then the sublist diameter is bounded by k . For the traversal diameter, if there is some constant k that bounds the number of states traversed by executing every valid action sequence, then the traversal diameter is bounded by one less than k .

$$\vdash \text{FINITE } \delta \wedge \\ (\forall \vec{\pi} x. \\ x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \Rightarrow \\ \exists \vec{\pi}'. \text{ex}(x, \vec{\pi}) = \text{ex}(x, \vec{\pi}') \wedge \vec{\pi}' \preceq \vec{\pi} \wedge |\vec{\pi}'| \leq k) \Rightarrow \\ \ell(\delta) \leq k$$

$$\vdash \text{FINITE } \delta \wedge \\ (\forall x \vec{\pi}. \\ \text{sat-pre}(x, \vec{\pi}) \wedge x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \Rightarrow \\ |\text{ss}(x, \vec{\pi})| \leq k) \Rightarrow \\ td(\delta) \leq k - 1$$

Note that in the theorem for the traversal diameter, there is a requirement on the action sequence to have all of its preconditions satisfied ($\text{sat-pre}(x, \vec{\pi})$). This is because in our execution semantics, an action whose preconditions are not satisfied is executed, but with no effect. This means that there can be action sequences with unbounded lengths, that contain actions that do not change the state. This condition is formally defined as follows.

$$\begin{aligned} \text{sat-pre}(x, (p, e) :: \vec{\pi}) &\iff \\ p \sqsubseteq x \wedge \text{sat-pre}(\text{state-succ } x (p, e), \vec{\pi}) & \\ \text{sat-pre}(x, []) &\iff \text{T} \end{aligned}$$

5.2.2 Abstraction

An important concept for our work is abstraction of a factored system. The first type of abstraction we consider is projection. To project a state x on a set of variables vs , we use the **DRESTRICT** function that restricts the domain of a finite map to a set of variables vs , and it is pretty printed as $x \downarrow_{vs}$. For actions, action sequences, and factored systems, we defined projection as follows.

HOL4 Definition 7 (Projection). *For an action*

$$(p, e) \downarrow_{vs} = (p \downarrow_{vs}, e \downarrow_{vs})$$

For an action sequence

$$\begin{aligned} ((p, e) :: \vec{\pi}) \downarrow_{vs} &= \text{if } \mathcal{D}(e \downarrow_{vs}) \neq \emptyset \text{ then } (p, e) \downarrow_{vs} :: \vec{\pi} \downarrow_{vs} \text{ else } \vec{\pi} \downarrow_{vs} \\ [] \downarrow_{vs} &= [] \end{aligned}$$

Letting $f \uparrow(t)$ denote the image of a function f on some set t , projection for a factored system is defined as

$$\delta \downarrow_{vs} = (\lambda \pi. \pi \downarrow_{vs}) \uparrow(\delta)$$

For a set of states

$$\vec{x} \downarrow_{vs} = (\lambda x. x \downarrow_{vs}) \uparrow(\vec{x})$$

The following theorems show some of the basic properties of the projection operation that we defined.

$$\begin{aligned} \vdash x \in \mathbb{U}(\delta) &\Rightarrow x \downarrow_{vs} \in \mathbb{U}(\delta \downarrow_{vs}) \\ \vdash \vec{\pi} \in \delta^* &\Rightarrow \vec{\pi} \downarrow_{vs} \in \delta \downarrow_{vs}^* \\ \vdash (\text{ex}(x, \vec{\pi} \downarrow_{vs})) \downarrow_{vs} &= \text{ex}(x \downarrow_{vs}, \vec{\pi} \downarrow_{vs}) \\ \vdash \text{sat-pre } (x, \vec{\pi}) &\Rightarrow \text{ex}(x \downarrow_{vs}, \vec{\pi} \downarrow_{vs}) = (\text{ex}(x, \vec{\pi})) \downarrow_{vs} \end{aligned}$$

Note that in the theorems above that relate executing a concrete action sequence and its projection, there is the condition **sat-pre** $(x, \vec{\pi})$. If this condition does not hold, some actions in the concrete action sequence whose preconditions are not satisfied can have their preconditions satisfied after projection; the projection would then be an incorrect presentation. This could lead to a different execution outcome than the projection of the concrete execution outcome, i.e. $\text{ex}(x \downarrow_{vs}, \vec{\pi} \downarrow_{vs}) = (\text{ex}(x, \vec{\pi})) \downarrow_{vs}$ does not hold unconditionally.

The second type of abstraction that we define is snapshotting.

HOL4 Definition 8 (Snapshotting). *We first define the following relation between states:*

$$\text{agree } x_1 \ x_2 \iff \forall v. v \in \mathcal{D}(x_1) \wedge v \in \mathcal{D}(x_2) \Rightarrow x_1 \uparrow v = x_2 \uparrow v$$

Based on that, a snapshot is defined as

$$\delta \downarrow_x = \{(p, e) \mid (p, e) \in \delta \wedge \text{agree } p \ x \wedge \text{agree } e \ x\}$$

The relation **agree** indicates that two states assign all variables in the intersection of their domains to the same values. A snapshot of a factored system δ on a state x is the set of actions from δ whose preconditions are enabled by x and that if executed they do not change the assignments of variables in the domain of x . The following properties of the agreement relation between states and the snapshot abstraction are sanity checks for the validity of our definitions.

$$\vdash f_1 \sqsubseteq f_2 \Rightarrow \text{agree } f_1 f_2$$

$$\begin{aligned} \vdash \vec{\pi} \in \delta^* \wedge x' \in \mathbb{U}(\delta) \wedge x \in \mathbb{U}(\delta) \wedge \\ (\forall p e. \text{MEM } (p,e) \vec{\pi} \Rightarrow \text{agree } e|_{vs} x|_{vs}) \wedge x'|_{vs} = x|_{vs} \Rightarrow \\ (\text{ex}(x', \vec{\pi}))|_{vs} = x|_{vs} \end{aligned}$$

The second theorem above shows that if all actions in a sequence agree with a projection of an initial state, then the result of executing that action sequence on that state will have the same assignment to the initially agreed upon variables.

We lastly note that the diameter (resp. sublist diameter, recurrence diameter and traversal diameter) of an abstract system obtained using one of the above abstractions is in general incomparable to that of the concrete system.

5.3 Formalising Compositional Upper-Bounding Algorithms

In this section we review our approach to formally verify the validity of different compositional bounding approaches, including those computed by HYB.

5.3.1 Compositional Bounding in the General Case

We now discuss our formalisation of Theorem 3. That theorem states the validity of projection based compositional bounding using the traversal diameter, and it shows that it is always sound regardless of the dependency structure of the system under consideration. To formalise that theorem, we closely follow the pen and paper proof and formalise the propositions on which it depends. Formalising Proposition 1 and Proposition 2 is straightforward, unlike formalising Proposition 3 which warrants more detail. We first define the notion of a state space, which is a set of states that all have the same domain, defined as follows.

HOL4 Definition 9 (State Space).

$$\text{sspc } \vec{x} \text{ vs} \iff \forall x. x \in \vec{x} \Rightarrow \mathcal{D}(x) = \text{vs}$$

We then define the operation of multiplying two state spaces as follows.

HOL4 Definition 10 (State Space Multiplication).

$$(\text{ss}_1 \times \text{ss}_2) = (\lambda (x_1, x_2). x_1 \uplus x_2)(\text{ss}_1 \times \text{ss}_2)$$

We then extend that operation to sets of state spaces, following the approach by Nipkow and Paulson (2005). We first define that product as the following inductive relation:

HOL4 Definition 11 (Set of State Spaces Multiplication).

$$\begin{aligned} \text{PROD } \emptyset \vec{x} \vec{x} \\ \vec{x} \notin \text{SS} \wedge \text{PROD } \text{SS } \text{ss}' \text{ss}'' \Rightarrow \text{PROD } (\vec{x} \text{ INSERT } \text{SS}) \text{ss}' (\vec{x} \times \text{ss}'') \end{aligned}$$

Then we prove that this relation is unique, and use that uniqueness to redefine it as the function `PRODF`. The approach in Nipkow and Paulson (2005) requires the operation to be commutative and associative. The state space product as defined above is associative. However, to have commutativity we assumed that the domains of the different state spaces are disjoint.

$$\vdash (\text{ss}_1 \times (\text{ss}_2 \times \text{ss}_3)) = ((\text{ss}_1 \times \text{ss}_2) \times \text{ss}_3)$$

$$\begin{aligned} \vdash \text{sspc } ss_1 \text{ } vs_1 \wedge \text{sspc } ss_2 \text{ } vs_2 \wedge \text{DISJOINT } vs_1 \text{ } vs_2 \Rightarrow \\ (ss_1 \times ss_2) = (ss_2 \times ss_1) \end{aligned}$$

Since the states traversed by executing any valid action sequence will always be a state space, the following formalised version of Proposition 3 follows.

$$\begin{aligned} \vdash \text{FINITE } VS \wedge (\forall vs'. \text{ } vs' \in VS \Rightarrow \text{DISJOINT } vs \text{ } vs') \wedge \\ (\forall vs \text{ } vs'. \\ \text{ } vs \in VS \wedge vs' \in VS \wedge vs \neq vs' \Rightarrow \text{DISJOINT } vs \text{ } vs') \wedge \\ \bigcup (vs \text{ INSERT } VS) = \mathcal{D}(\delta) \wedge x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* \wedge \\ \text{sat-pre } (x, \vec{\pi}) \Rightarrow \\ \text{ss } (x, \vec{\pi}) \subseteq \text{PRODF } (\lambda vs. (\text{ss } (x, \vec{\pi})) \downarrow_{vs})(\|VS\|) (\text{ss } (x, \vec{\pi})) \downarrow_{vs} \end{aligned}$$

Lastly, the formalised version of Theorem 3 is as follows.

HOL4 Theorem 2.

$$\begin{aligned} \vdash \text{FINITE } VS \wedge \text{FINITE } \delta \wedge vs \notin VS \wedge \\ (\forall vs'. \text{ } vs' \in VS \Rightarrow \text{DISJOINT } vs \text{ } vs') \wedge \\ (\forall vs \text{ } vs'. \\ \text{ } vs \in VS \wedge vs' \in VS \wedge vs \neq vs' \Rightarrow \text{DISJOINT } vs \text{ } vs') \wedge \\ \bigcup (vs \text{ INSERT } VS) = \mathcal{D}(\delta) \wedge vs \notin VS \Rightarrow \\ \text{td}(\delta) \leq \Pi (\lambda vs. \text{td}(\delta \downarrow_{vs}) + 1) (vs \text{ INSERT } VS) - 1 \end{aligned}$$

5.3.2 Compositional Bounds in the Presence of Acyclicity

We now review our formalisation of acyclicity, Lemma 3, and Theorem 14. In HOL we model a DAG with the predicate `top-sorted-abs` that means that l is a list of vertices of a digraph topologically sorted w.r.t. to the binary relation R , which is taken to be the edge relation in the digraph. This predicate is defined as follows in HOL:

HOL4 Definition 12 (DAG).

$$\begin{aligned} \text{top-sorted-abs } R \text{ } (u_1 :: A) \iff \\ \text{EVERY } (\lambda u_2. \neg R \text{ } u_2 \text{ } u_1) A \wedge \text{top-sorted-abs } R \text{ } A \\ \text{top-sorted-abs } R \text{ } [] \iff \text{T} \end{aligned}$$

In the definition above, `EVERY` is a high-order predicate that, given a predicate and a list, returns true if every member of the list satisfies the given predicate. The following basic properties hold for DAGs.

$$\begin{aligned} \vdash \text{top-sorted-abs } R \text{ } (u_1 :: A) \wedge u_2 \in \text{set } A \Rightarrow \neg R \text{ } u_2 \text{ } u_1 \\ \vdash \text{top-sorted-abs } R \text{ } (u :: A) \Rightarrow \text{top-sorted-abs } R \text{ } A \end{aligned}$$

Recall that one way to interpret both bounding functions `N` and `S` is as functions that compute some form of weightiest paths in an acyclic graph, where every path is given a weight that depends on the weights of the vertices and the number of edges that it traverses. However, both functions operate on a DAG induced by different relations (dependency for `N` versus successor state for `S`), weigh vertices with different functions (projection's sublist diameter for `N` versus snapshot's sublist diameter for `S`), and combine those weights in different ways (multiply them and add for `N` versus add and take the maximum for `S`). To be able to factor what is common between the two bounding functions and yet accommodate the difference between them, we defined the following function.

HOL4 Definition 13 (Weightiest Longest Path).

$$\begin{aligned}
\text{wp } R \ w \ g \ f \ u_1 \ (u_2 :: A) &= \\
&\text{if } R \ u_1 \ u_2 \ \text{then } g \ (f \ (w \ u_1) \ (\text{wp } R \ w \ g \ f \ u_2 \ A)) \ (\text{wp } R \ w \ g \ f \ u_1 \ A) \\
&\text{else } \text{wp } R \ w \ g \ f \ u_1 \ A \\
\text{wp } R \ w \ g \ f \ u_1 \ [] &= w \ u_1
\end{aligned}$$

This function takes as an argument the relation that induces the DAG R , the weighing function w , and the two functions that combine the weights: g which combines weights of different paths; and f which combines the weights of vertices on one path, as well as the vertex of interest v and the DAG l . In order for wp to capture desired properties about \mathbf{S} and \mathbf{N} , we define the following properties on the weight combination function.

$$\text{geq-arg } f \iff \forall x \ y. \ x \leq f \ x \ y \wedge y \leq f \ x \ y$$

$$\text{increasing } f \iff \forall e \ b \ c \ d. \ e \leq c \wedge b \leq d \Rightarrow f \ e \ b \leq f \ c \ d$$

The previous definitions describe different notions of functions whose value is at least as big as their arguments and functions that are non-decreasing. If the weight combination functions f and g have those properties, the following basic properties hold for the abstract weighted longest path function, and accordingly for \mathbf{N}_{sum} and \mathbf{S}_{max} .

$$\begin{aligned}
&\vdash \text{geq-arg } g \Rightarrow w \ u \leq \text{wp } R \ w \ g \ f \ u \ A \\
&\vdash \text{geq-arg } f \wedge \text{geq-arg } g \wedge (\forall u. \ u \in \text{set } A \Rightarrow \neg R \ u \ u) \wedge R \ u_2 \ u_1 \wedge \\
&\quad u_1 \in \text{set } A \wedge \text{top-sorted-abs } R \ A \Rightarrow \\
&\quad f \ (w \ u_2) \ (\text{wp } R \ w \ g \ f \ u_1 \ A) \leq \text{wp } R \ w \ g \ f \ u_2 \ A \\
&\vdash \text{increasing } f \wedge \text{increasing } g \wedge \\
&\quad (\forall u. \ u \in \text{set } A \Rightarrow w_1 \ u \leq w_2 \ u) \wedge w_1 \ u \leq w_2 \ u \Rightarrow \\
&\quad \text{wp } R \ w_1 \ g \ f \ u \ A \leq \text{wp } R \ w_2 \ g \ f \ u \ A
\end{aligned}$$

Formalising the Validity of \mathbf{N}_{sum}

We now review our formalisation of the compositional bounding of the sublist diameter, for the case of projections induced by an acyclic dependency graph. We defined dependency between variables and dependency between sets of variables as follows.

HOL4 Definition 14 (Dependency). *For two variables v_1 and v_2 , we define dependency as²*

$$\begin{aligned}
v_1 \rightarrow v_2 &\iff \\
&(\exists p \ e. \\
&\quad (p, e) \in \delta \wedge \\
&\quad (v_1 \in \mathcal{D}(p) \wedge v_2 \in \mathcal{D}(e) \vee v_1 \in \mathcal{D}(e) \wedge v_2 \in \mathcal{D}(e))) \vee \\
&\quad v_1 = v_2
\end{aligned}$$

For sets of variables vs_1 and vs_2 , we define

$$\begin{aligned}
vs_1 \rightarrow vs_2 &\iff \\
&\exists v_1 \ v_2. \ v_1 \in vs_1 \wedge v_2 \in vs_2 \wedge \text{DISJOINT } vs_1 \ vs_2 \wedge v_1 \rightarrow v_2
\end{aligned}$$

²The definition of \rightarrow has an implicit δ parameter; however, we hide it using the *ad hoc* overloading ability in HOL.

In the above definition, **DISJOINT** is a predicate that, given two sets, returns true if they share no elements.

Our proof of Theorem 8 depends on Lemma 3 whose proof depends on the following fundamental argument. Consider a set of variables vs with only incoming dependencies from all the state variables in a system δ (a “parent-child” dependency structure). If for some $\vec{\pi} \in \delta^*$ one is to remove actions from the projection $\vec{\pi}|_{vs}$, while preserving the projection’s execution outcome, the restitching function can be used to “stitch” back the shortened projected action sequence into $\vec{\pi}$, producing an action sequence with the same execution outcome as $\vec{\pi}$, but with the redundant actions removed. To formalise that, we formalise the parent-child dependency configuration and the restitching function as follows:

HOL4 Definition 15 (Parent-Child Structure).

$$\text{child-parent-rel } (\delta, vs) \iff vs \not\rightarrow \overline{vs}$$

HOL4 Definition 16 (Restitching Function).

$$\begin{aligned} (\pi' :: \vec{\pi}') \upharpoonright_{vs}^{\rightarrow} (\pi :: \vec{\pi}) &= \\ \text{if } \text{varset-action } (\pi, vs) &\text{ then} \\ \text{if } \pi' = \pi|_{vs} &\text{ then } \pi :: \vec{\pi}' \upharpoonright_{vs}^{\rightarrow} \vec{\pi} \text{ else } (\pi' :: \vec{\pi}') \upharpoonright_{vs}^{\rightarrow} \vec{\pi} \\ \text{else } \pi :: (\pi' :: \vec{\pi}') \upharpoonright_{vs}^{\rightarrow} &\vec{\pi} \\ [] \upharpoonright_{vs}^{\rightarrow} &= \text{FILTER } (\lambda \pi. \neg \text{varset-action } (\pi, vs)) \vec{\pi} \\ (\pi' :: \vec{\pi}') \upharpoonright_{vs}^{\rightarrow} [] &= [] \end{aligned}$$

where **varset-action** is defined as follows

$$\text{varset-action } ((p, e), vs) \iff \mathcal{D}(e) \subseteq vs$$

In words, the stitching function uses the first list as a guide: the second list of actions is filtered, with each action in the first list meant to have a corresponding action in the second list. The formalised theorem stating the aforementioned functionality of the restitching function is as follows. Its proof script is 1300 lines long, with comments, and it is our main tool to formally prove the validity of N_{sum} as a bound.

HOL4 Lemma 1.

$$\begin{aligned} \vdash \text{child-parent-rel } (\delta, vs) \wedge x \in \mathbb{U}(\delta) \wedge \text{no-effectless-act } \vec{\pi} \wedge \\ \vec{\pi}' \preceq \vec{\pi}|_{vs} \wedge \vec{\pi} \in \delta^* \wedge (\text{ex}(x, \vec{\pi}))|_{vs} = \text{ex}(x|_{vs}, \vec{\pi}') \wedge \\ \text{sat-pre } (x, \vec{\pi}) \wedge \text{sat-pre } (x, \vec{\pi}') \Rightarrow \\ \text{ex}(x, \vec{\pi}' \upharpoonright_{vs}^{\rightarrow} \vec{\pi}) = \text{ex}(x, \vec{\pi}) \end{aligned}$$

In the previous statement, the predicate **no-effectless-act** $\vec{\pi}$ asserts that the action sequence has no actions with empty effects. As we sketched earlier in the proof summary of Lemma 3 in Section 3.6, for a lifted dependency DAG A_{VS} , we take every set of variables $p \in A_{VS}$, and remove all redundant actions whose effects are confined to p . As we describe in detail below, we use the stitching function as our main proof tool to perform the action removal. To formalise this, we define the following relation, which generalises the parent-child structure.

HOL4 Definition 17 (Generalised Parent-Child Structure). *For a factored transition system δ and two sets of variables vs_1 and vs_2 , the generalised parent-child relation holds between vs_1 and vs_2 iff (i) $vs_2 \not\rightarrow vs_1$, (ii) $vs_1 \not\rightarrow (vs_1 \cup vs_2)$, and (iii) no bidirectional dependencies exist between any variable in vs_2 and $(vs_1 \cup vs_2)$. Formally:*

$$\begin{aligned} \text{gen-parent-child } (\delta, p, c) &\iff \\ \text{DISJOINT } p \ c \ \wedge \ c \not\rightarrow p \ \wedge \ p \not\rightarrow \overline{p \cup c} \ \wedge \\ \forall v_1 \ v_2. \ v_1 \in c \ \wedge \ v_2 \in \overline{p \cup c} &\Rightarrow v_1 \not\rightarrow v_2 \ \vee \ v_2 \not\rightarrow v_1 \end{aligned}$$

Given the previous definition, the following lemma formally describes the process of removing redundant actions affecting a set of variables $p \in A_{VS}$.

Lemma 9. *Let $n(vs, \vec{\pi})$ be the number of vs -actions contained within $\vec{\pi}$. Consider δ , in which the generalised parent-child relation holds between sets of variables p and c . Then, any action sequence $\vec{\pi}$ has a sublist $\vec{\pi}'$ that reaches the same state as $\vec{\pi}$ starting from any state such that: $n(p, \vec{\pi}') \leq \ell(\delta|_p)(n(c, \vec{\pi}') + 1)$ and $n(\overline{p}, \vec{\pi}') \leq n(\overline{p}, \vec{\pi})$.*

A formal statement of that lemma follows:

$$\begin{aligned} \vdash \text{FINITE } \delta \ \wedge \ x \in \mathbb{U}(\delta) \ \wedge \ \vec{\pi} \in \delta^* \ \wedge \ \text{gen-parent-child } (\delta, p, c) &\Rightarrow \\ \exists \vec{\pi}' & \\ n(p, \vec{\pi}') \leq \ell(\delta|_p) \times (n(c, \vec{\pi}') + 1) \ \wedge \ \vec{\pi}' \preceq \vec{\pi} \ \wedge & \\ n(\mathcal{D}(\delta) \setminus p, \vec{\pi}') \leq n(\mathcal{D}(\delta) \setminus p, \vec{\pi}) \ \wedge \ \text{ex}(x, \vec{\pi}') = \text{ex}(x, \vec{\pi}) & \end{aligned}$$

where

$$n(p, \vec{\pi}) = |\text{FILTER } (\lambda \pi. \text{varset-action } (\pi, p)) \ \vec{\pi}|$$

Proof. The proof of Lemma 9 is constructive. Let $\vec{\pi}_c$ be a contiguous fragment of $\vec{\pi}$ that has no c -actions in it. Then perform the following steps:

- By the definition of ℓ , there must be an action sequence $\vec{\pi}_p$ such that $\text{ex}(s, \vec{\pi}_p) = \text{ex}(s, \vec{\pi}_c|_p)$, and satisfies $|\vec{\pi}_p| \leq \ell(\delta|_p)$ and $\vec{\pi}_p \preceq \vec{\pi}_c|_p$.
- Letting $D \equiv \mathcal{D}(\delta)$, $p \not\rightarrow D \setminus p \setminus c$ holds. Accordingly, from HOL4 Lemma 1, $\vec{\pi}'_c (= \vec{\pi}_p \upharpoonright \vec{\pi}_c|_{D \setminus c})$ achieves the same $D \setminus c$ assignment as $\vec{\pi}_c$ (i.e., $\text{ex}(s, \vec{\pi}'_c)|_{D \setminus c} = \text{ex}(s, \vec{\pi}_c)|_{D \setminus c}$), and it is a sublist of $\vec{\pi}_c$. Also, $n(p, \vec{\pi}'_c) \leq \ell(\delta|_p)$ holds.
- Finally, because $\vec{\pi}_c$ has no c -actions, no c variables change along the execution of $\vec{\pi}_c$ and accordingly any c variables in preconditions of actions in $\vec{\pi}_c$ always have the same assignment. This means that $\vec{\pi}'_c \upharpoonright_{D \setminus c} \vec{\pi}_c$ will achieve the same result as $\vec{\pi}_c$, but with at most $\ell(\delta|_p)$ p -actions.

Repeating the previous steps for each $\vec{\pi}_c$ fragment in $\vec{\pi}$ yields an action sequence $\vec{\pi}'$ that has at most $\ell(\delta|_p)(n(c, \vec{\pi}') + 1)$ p -actions. This is justified by the following lemma, where in this lemma $\text{list-frag } (l_1, l_2)$ means that list l_2 is a contiguous sublist of l_1 .

$$\begin{aligned} \vdash |\text{FILTER } P_1 \ l| \leq k_1 \ \wedge \ (\forall x. \ x \in \text{set } l \Rightarrow P_1 \ x \Rightarrow \neg P_2 \ x) \ \wedge \\ (\forall l'. \\ \text{list-frag } (l, l') \ \wedge \ \text{EVERY } (\lambda x. \ \neg P_1 \ x) \ l' \Rightarrow \\ |\text{FILTER } P_2 \ l'| \leq k_2) \Rightarrow \\ |\text{FILTER } P_2 \ l| \leq (k_1 + 1) \times k_2 \end{aligned}$$

Because $\vec{\pi}'$ is the result of consecutive applications of the stitching function, it is a sublist of $\vec{\pi}$. Lastly, because during the previous steps, only p -actions were removed as necessary, the number of c -actions in $\vec{\pi}'$ is the same as their number in $\vec{\pi}$. \square

We can now describe how we use that result to prove Lemma 3, the main lemma stating the validity of using N_{sum} to compositionally upper-bound the sublist diameter. The main idea is to preform an induction on the acyclic lifted dependency graph, where for every node $p \in A_{VS}$, the redundant actions of p are removed using Lemma 9 after removing the redundant actions of its children c . Then we use the stitching function to reconcile both shortened action sequences. We now describe our formalisation of that proof. Below are the formal definitions of a lifted dependency DAG, and the function N . They are instantiations of `top-sorted-abs` and `wp`, respectively, where the relation that induces the digraph is the variable dependency relation. Also for N , the functions for combining node weights g and f , are instantiated with addition and multiplication, respectively³.

HOL4 Definition 9 (Lifted Dependency DAG).

$$\begin{aligned} \text{dep-DAG } \delta \ A_{VS} &\iff \\ \mathcal{D}(\delta) &= \bigcup (\text{set } A_{VS}) \wedge \text{ALL-DISTINCT } A_{VS} \wedge \\ &\text{ALL-DISJOINT } A_{VS} \wedge \\ &\text{top-sorted-abs } (\lambda \ v_1 \ v_2. \ v_1 \rightarrow v_2) \ A_{VS} \end{aligned}$$

In the above definition the predicates `ALL-DISTINCT` and `ALL-DISJOINT` return true if all members of a given list are pairwise distinct and pairwise disjoint, respectively.

HOL4 Definition 10 (Acyclic Dependency Compositional Bound).

$$N\langle b \rangle(vs) = \text{wp } (\lambda \ v_1 \ v_2. \ v_1 \rightarrow v_2) (\lambda \ v. \ b(\delta|_{vs})) (+) (\times) \ v \ A_{VS}$$

This is the formal statement of Lemma 3.

HOL4 Theorem 3.

$$\vdash \text{FINITE } \delta \wedge \text{dep-DAG } \delta \ A_{VS} \Rightarrow \ell(\delta) < \text{SUM } (\text{MAP } N\langle \ell \rangle \ A_{VS}) + 1$$

Before we discuss the formal proof, we first introduce the following notation. Let $F(p, c, \vec{\pi})$ be the witness action sequence of Lemma 9. We know that:

- $\text{ex}(s, F(p, c, \vec{\pi})) = \text{ex}(s, \vec{\pi})$,
- $n(p, F(p, c, \vec{\pi})) \leq \ell(\delta|_p)(n(c, \vec{\pi}) + 1)$.
- $F(p, c, \vec{\pi}) \preceq \vec{\pi}$, and
- $n(\bar{p}, F(p, c, \vec{\pi})) \leq n(\bar{p}, \vec{\pi})$.

Proof of Lemma 3. Firstly, for brevity, for a set of variables vs , in the rest of this proof we write $N\langle \ell \rangle(vs)$ as a short-hand for $N\langle \ell \rangle(vs, \delta, A_{VS})$.

Our proof of this lemma follows a constructive approach where we assume we have an action sequence $\vec{\pi} \in \delta^*$ and a state $s \in \mathbb{U}(\delta)$. The goal of the proof is to find a witness sublist, $\vec{\pi}'$, of $\vec{\pi}$ such that $\forall vs \in A_{VS}. n(vs, \vec{\pi}') \leq N\langle \ell \rangle(vs)$ and $\text{ex}(s, \vec{\pi}) = \text{ex}(s, \vec{\pi}')$. We proceed by induction

³ N has δ and A_{VS} parameters, but we hide them with HOL's *ad hoc* overloading ability.

on $V(A_{VS})$ assuming it is topologically sorted in a list l_{VS} (without loss of generality since A_{VS} is a DAG). The base case is the empty list $[],$ in which case $\mathcal{D}(\delta) = \emptyset$ and accordingly $\ell(\delta) = 0.$

In the step case, we assume the result holds for any factored system for which l'_{VS} is a topologically sorted vertices list of one of its lifted dependency graphs. We then show that it also holds for $\delta,$ a factored system whose dependency graph's vertices are topologically sorted into $vs :: l_{VS}.$ Let $\overline{vs} \equiv \bigcup l_{VS}.$ Since l_{VS} is a topologically sorted vertices list of a lifted dependency graph of $\delta|_{\overline{vs}},$ the induction hypothesis applies. Accordingly, there is $\vec{\pi}_{\overline{vs}} \in \delta|_{\overline{vs}}^*$ such that $\text{ex}(s, \vec{\pi}_{\overline{vs}}) = \text{ex}(s, \vec{\pi}|_{\overline{vs}}), \vec{\pi}_{\overline{vs}} \preceq \vec{\pi}|_{\overline{vs}},$ and $\forall vs' \in l_{VS}. \mathfrak{n}(vs', \vec{\pi}') \leq \mathbf{N}\langle\ell\rangle(vs', \delta|_{\overline{vs}}, A_{VS}) \leq \mathbf{N}\langle\ell\rangle(vs').$ Since $vs :: l_{VS}$ is topologically sorted, $\overline{vs} \not\rightarrow vs$ holds. Letting $\vec{\pi}'_{vs} = \vec{\pi}_{\overline{vs}} \upharpoonright_{vs} \vec{\pi},$ from HOL4 Lemma 1 we have $\text{ex}(s, \vec{\pi}'_{vs}) = \text{ex}(s, \vec{\pi}).$ Furthermore, $\forall vs' \in l_{VS}. \mathfrak{n}(vs', \vec{\pi}'_{vs}) \leq \mathbf{N}\langle\ell\rangle(vs')$ and $\vec{\pi}'_{vs} \preceq \vec{\pi}.$ Let $C \equiv \bigcup \text{children}_{A_{VS}}(vs).$ The last step in this proof is to show that $F(vs, C, \vec{\pi}'_{vs})$ is the required witness, which is justified because the generalised parent-child relation holds for δ, vs and $C.$ Since the relations $=, \leq$ and \preceq are transitive, we have

- $\text{ex}(s, \vec{\pi}) = \text{ex}(s, F(vs, C, \vec{\pi}'_{vs})),$
- $\mathfrak{n}(vs, F(vs, C, \vec{\pi}'_{vs})) \leq \ell(\delta|_{vs})(\mathfrak{n}(C, \vec{\pi}'_{vs}) + 1) = \ell(\delta|_{vs})(\sum_{c \in C} \mathfrak{n}(c, \vec{\pi}'_{vs}) + 1),$
- $F(vs, C, \vec{\pi}'_{vs}) \preceq \vec{\pi},$ and
- $\mathfrak{n}(\overline{vs}, F(vs, C, \vec{\pi}'_{vs})) \leq \mathfrak{n}(\overline{vs}, \vec{\pi}'_{vs}).$

Since $\sum_{vs' \in l_{VS}} \mathfrak{n}(vs', \vec{\pi}'_{vs}) = \mathfrak{n}(\overline{vs}, \vec{\pi}'_{vs}),$ then $\forall vs' \in l_{VS}. \mathfrak{n}(vs', \vec{\pi}'_{vs}) \leq \mathbf{N}\langle\ell\rangle(vs')$ and $\mathfrak{n}(vs, F(vs, C, \vec{\pi}'_{vs})) \leq \ell(\delta|_{vs})(\sum_{c \in C} \mathbf{N}\langle\ell\rangle(c))$ hold. Therefore $F(vs, C, \vec{\pi}'_{vs})$ is an action sequence demonstrating the needed bound. \square

Formalising the Validity of \mathbf{S}_{\max}

To formalise compositional upper-bounding under state space acyclicity, we again use `top-sorted-abs` and `wp` to formalise the digraph modelling the state space and the function $\mathbf{S},$ respectively. In this case, the relation that induces the digraph is the successor relation on states. We note that in order for us to use the second property of `wp`, the relation that induces the digraph needs to be irreflexive, which is why the set of successors of a state is defined to not include that state. For $\mathbf{S},$ the vertex weight combination functions g and f are instantiated by a function that chooses the maximum of two arguments and the addition function, respectively. Formally this is defined as follows⁴

HOL4 Definition 11 (Acyclic System State Space).

$$\text{sspace-DAG } \delta \vec{x} \iff \text{set } \vec{x} = \mathbb{U}(\delta) \wedge \text{top-sorted-abs } (\text{succ } \delta) \vec{x}$$

where

$$\text{succ } \delta x = \text{state-succ } x(\delta) \setminus \{x\}$$

HOL4 Definition 12 (Acyclic System Compositional Bound).

$$\mathbf{S}\langle b \rangle(x) = \text{wp } (\text{succ } \delta|_{vs}) (\lambda x. b(\delta \upharpoonright_x)) \text{MAX } (\lambda x y. x + y + 1) x \vec{x}$$

⁴ \mathbf{S} has vs, δ and \vec{x} parameters, but we hide them with HOL's *ad hoc* overloading ability.

Formalising Theorem 14 is a straightforward implementation of the proof discussed in Section 3.6.2. Proposition 5 follows from the properties of **wp** shown earlier as well as HOL4 Definition 12. To formalise the other propositions we define the function ∂ , that gives the number of changes in the assignments of a set of variables vs if an action sequence $\vec{\pi}$ is executed on a state x , as follows

HOL4 Definition 13 (Subsystem Trace).

$$\begin{aligned} \partial (\pi :: \vec{\pi}) \text{ vs } x &= \\ \text{if } (\text{state-succ } x \pi)|_{vs} &\neq x|_{vs} \text{ then} \\ \text{state-succ } x \pi :: \partial \vec{\pi} \text{ vs } &(\text{state-succ } x \pi) \\ \text{else } \partial \vec{\pi} \text{ vs } &(\text{state-succ } x \pi) \\ \partial [] \text{ vs } x &= [] \end{aligned}$$

Proposition 6, Proposition 7 are formalised as follows.

HOL4 Proposition 1.

$$\begin{aligned} \vdash \text{FINITE } \delta \wedge \partial \vec{\pi} \text{ vs } x &= [] \wedge \text{sat-pre } (x, \vec{\pi}) \wedge x \in \mathbb{U}(\delta) \wedge \\ \vec{\pi} \in \delta^* &\Rightarrow \\ \exists \vec{\pi}' . & \\ \text{ex}(x|_{\mathcal{D}(\delta|_{vs})}, \vec{\pi}) &= \text{ex}(x|_{\mathcal{D}(\delta|_{vs})}, \vec{\pi}') \wedge \vec{\pi}' \preceq \vec{\pi} \wedge \\ |\vec{\pi}'| &\leq \ell(\delta|_{vs}) \end{aligned}$$

HOL4 Proposition 2.

$$\begin{aligned} \vdash \partial \vec{\pi} \text{ vs } x = x' :: \vec{x} &\Rightarrow \\ \exists \vec{\pi}_1 \pi \vec{\pi}_2 . & \\ \vec{\pi} = \vec{\pi}_1 \# \pi :: \vec{\pi}_2 \wedge \partial \vec{\pi}_1 \text{ vs } x &= [] \wedge \\ \text{state-succ } (\text{ex}(x, \vec{\pi}_1)) \pi = x' \wedge & \\ \partial \vec{\pi}_2 \text{ vs } (\text{state-succ } (\text{ex}(x, \vec{\pi}_1)) \pi) &= \vec{x} \end{aligned}$$

Note that in the conclusions of the previous formalised statements there is an extra requirement on the witness action sequence: that it is a sublist of the original sequence. This is one difference from the proof discussed of Theorem 13 in Section 3.6.2 and it is because in the earlier proof the aim is to prove the validity of \mathbf{S}_{\max} as a compositional bound on the diameter, while here we prove its validity as an upper bound on the sublist diameter. Other than this extra conclusion, the proofs of Theorem 13 and Theorem 14 are the same. Lemma 6 is then formalised as follows:

HOL4 Lemma 2.

$$\begin{aligned} \vdash \text{FINITE } \delta \wedge \text{sspace-DAG } \delta|_{vs} \vec{x} \wedge x \in \mathbb{U}(\delta) \wedge \vec{\pi} \in \delta^* &\Rightarrow \\ \exists \vec{\pi}' . \text{ex}(x, \vec{\pi}') = \text{ex}(x, \vec{\pi}) \wedge \vec{\pi}' \preceq \vec{\pi} \wedge |\vec{\pi}'| &\leq \mathbf{S}\langle \ell \rangle(x|_{vs}) \end{aligned}$$

Finally, based on HOL4 Lemma 2 the formalised statement of Theorem 14 is as follows:

HOL4 Theorem 4.

$$\begin{aligned} \vdash \text{FINITE } \delta \wedge \text{sspace-DAG } \delta|_{vs} \vec{x} &\Rightarrow \\ \ell(\delta) \leq \max \{ \mathbf{S}\langle \ell \rangle(x') \mid x' \in \mathbb{U}(\delta|_{vs}) \} & \end{aligned}$$

Formalising the Validity of HYB

We formalise HYB as follows:

HOL4 Definition 14 (The Hybrid Algorithm).

$$\begin{aligned}
\text{HYB } f_1 f_2 \delta = & \\
& \text{if FINITE } \delta \text{ then} \\
& \quad \text{if } \forall vs. vs \in \text{set } (f_1 \delta) \Rightarrow vs \subset \mathcal{D}(\delta) \text{ then} \\
& \quad \quad (\text{let} \\
& \quad \quad \quad A_{vs} = f_1 \delta \\
& \quad \quad \text{in} \\
& \quad \quad \quad \text{SUM} \\
& \quad \quad \quad \quad (\text{MAP} \\
& \quad \quad \quad \quad \quad (\text{N} \\
& \quad \quad \quad \quad \quad \quad (\lambda \delta'. \\
& \quad \quad \quad \quad \quad \quad \quad \text{if } \exists vs. vs \in \text{set } A_{vs} \wedge \delta' = \delta|_{vs} \text{ then} \\
& \quad \quad \quad \quad \quad \quad \quad \quad \text{HYB } f_1 f_2 \delta' \\
& \quad \quad \quad \quad \quad \quad \quad \quad \text{else } 0) \delta A_{vs}) A_{vs})) \\
& \quad \text{else if} \\
& \quad \quad f_2 \delta \neq \emptyset \wedge \forall vs \vec{x} x. (vs, \vec{x}) \in f_2 \delta \wedge x \in \text{set } \vec{x} \Rightarrow \delta|_x \subset \delta \\
& \quad \text{then} \\
& \quad \quad (\text{let} \\
& \quad \quad \quad (vs, \vec{x}) = \text{CHOICE } (f_2 \delta) \\
& \quad \quad \text{in} \\
& \quad \quad \quad \text{max} \\
& \quad \quad \quad \quad \{ \text{S-gen} \\
& \quad \quad \quad \quad \quad (\lambda \delta'. \\
& \quad \quad \quad \quad \quad \quad \text{if } \exists x. x \in \text{set } \vec{x} \wedge \delta' = \delta|_x \text{ then HYB } f_1 f_2 \delta' \\
& \quad \quad \quad \quad \quad \quad \text{else } 0) vs \vec{x} \delta x' \mid \\
& \quad \quad \quad \quad \quad \quad \quad x' \in \mathbb{U}(\delta|_{vs}) \} \\
& \quad \quad \text{else } \ell(\delta) \\
& \quad \text{else } 0
\end{aligned}$$

It takes two functions f_1 and f_2 as arguments. The first function is an oracle that given a transition system, returns a lifted dependency DAG of that system. The second function is an oracle that returns a set of pairs, each of which has a strict subset of the state variables of the given system, and the state space of the projection of the system on that subset of the state variables. We note that proving the termination of HYB is not trivial. In order to guarantee termination, we add the conditions that δ is finite as well as that all the lifted dependency DAGs computed by f_1 and the subsets of the domain of δ computed by f_2 provide non trivial decompositions of δ , i.e. they are strict subsets. Also, in order to guide HOL4 to extract the right termination conditions, we add the redundant if-then-else statements in the functional parameters passed to N and S. After proving termination, we are able to obtain the following characterisation that resembles Algorithm 5.

HOL4 Lemma 3 (The Hybrid Algorithm).

$$\begin{aligned}
& \text{FINITE } \delta \wedge (\forall vs \vec{x}. (vs, \vec{x}) \in f_2 \delta \Rightarrow \text{set } \vec{x} = \mathbb{U}(\delta|_{vs})) \Rightarrow \\
& \text{HYB } f_1 f_2 \delta = \\
& \quad \text{if } \forall vs. vs \in \text{set } (f_1 \delta) \Rightarrow vs \subset \mathcal{D}(\delta) \text{ then} \\
& \quad \quad (\text{let } A_{vs} = f_1 \delta \text{ in SUM (MAP N(HYB) } A_{vs}))
\end{aligned}$$

```

else if
   $f_2 \delta \neq \emptyset \wedge \forall vs \vec{x} x. (vs, \vec{x}) \in f_2 \delta \wedge x \in \mathbf{set} \vec{x} \Rightarrow \delta \uparrow_x \subset \delta$ 
then
   $(\text{let } (vs, \vec{x}) = \mathbf{CHOICE} (f_2 \delta) \text{ in } \mathbf{max} \{ \mathbf{S}_{\langle \text{HYB} \rangle}(x') \mid x' \in \mathbf{U}(\delta \downarrow_{vs}) \})$ 
else  $\ell(\delta)$ 

```

Lastly, proving its validity as an upper bound on the sublist diameter follows directly from HOL4 Theorem 3 and HOL4 Theorem 4.

HOL4 Theorem 5.

```

 $\vdash \mathbf{FINITE} \delta \wedge$ 
 $(\forall \delta'.$ 
   $\mathbf{dep-DAG} \delta' (f_1 \delta') \wedge$ 
   $\forall vs \vec{x}. (vs, \vec{x}) \in f_2 \delta' \Rightarrow \mathbf{sspace-DAG} \delta' \downarrow_{vs} \vec{x} \Rightarrow$ 
   $\ell(\delta) \leq \mathbf{HYB} f_1 f_2 \delta$ 

```

5.4 Formalising the Compositional Reachability Algorithm

In this section we review our formalisation of our algorithm to compute plans via planning for the descriptive quotient. Precisely, we formalise Theorem 17, Theorem 18 and the validity of concatenating instantiations of the descriptive quotient's plan after its goal is augmented. However, in our approach we prove the validity of the algorithm after generalising it in two different ways. Firstly, we prove its validity for a more general class of abstractions: abstractions that are isomorphic to subproblems of the concrete problem, which are more general than the descriptive quotients. Secondly, we prove that the algorithm works for planning problems whose state characterising variables are of arbitrary types, and not necessarily Boolean. This shows that the algorithm has potential application to numerical planning as well as verification of hybrid systems.

An important point to state is that as a result of the formalisation, we found some subtle but non-trivial mistakes in our original algorithm and framework. We will point out those mistakes within our description of our formalisation. We also found that to verify that algorithm, there were other arguments that are easier to formalise than the ones we presented in our earlier proofs, arguments that we shall elaborate on in the next sections.

We first review how we formalise the basic concepts related to planning problems. First we formalise a planning problem as a record type in HOL4 as follows:

```

 $(\alpha, \beta) \mathbf{planningProblem} =$ 
 $\langle \mid \mathbf{l} : (\alpha \mapsto \beta);$ 
 $\delta : ((\alpha \mapsto \beta) \times (\alpha \mapsto \beta) \rightarrow \mathbf{bool});$ 
 $A : (\alpha \mapsto \beta) \mid \rangle$ 

```

Note that this record type extends the sets of actions that we used in our treatment of upper-bounding topological properties. In the above definition $\Pi.\mathbf{l}$ is the initial state, $\Pi.\delta$ is the set of actions a problem has, and $\Pi.A$ is the goal, which is a partial state. Note that this formalisation of planning problems does not require that state characterising variables to be Boolean.

An action sequence is a plan for a planning problem if it satisfies the following predicate:

HOL4 Definition 15 (Plan).

```

 $\Pi \mathbf{solved\_by} \vec{\pi} \iff \vec{\pi} \in \Pi.\delta^* \wedge \Pi.A \sqsubseteq \mathbf{ex}(\Pi.\mathbf{l}, \vec{\pi})$ 

```

In some cases we require the planning problem to be “valid”, which is defined as follows:

HOL4 Definition 16 (Planning Problem).

$$\text{planning_problem } \Pi \iff \Pi.l \in \mathbb{U}(\Pi.\delta) \wedge \mathcal{D}(\Pi.A) \subseteq \mathcal{D}(\Pi)$$

Also we define the following union operation on planning problems and a lifted union operation for lists of planning problems.

HOL4 Definition 17 (Planning Problem Union).

$$(\Pi_1 \cup \Pi_2) = \langle |l := \Pi_1.l \uplus \Pi_2.l; \delta := \Pi_1.\delta \cup \Pi_2.\delta; A := \Pi_1.A \uplus \Pi_2.A| \rangle$$

$$\bigcup \Pi_l = \text{FOLDER planning_prob_union } \Pi_\emptyset \Pi_l$$

Π_\emptyset is the empty problem defined as follows

$$\Pi_\emptyset = \langle |l := x_\emptyset; \delta := \emptyset; A := x_\emptyset| \rangle$$

where x_\emptyset denotes a state with an empty domain, i.e. a state that maps nothing to nothing.

The following theorems show that the semantics of the planning problem union operations are as intended.

- ⊢ $\text{planning_problem } \Pi_1 \wedge \text{planning_problem } \Pi_2 \Rightarrow \text{planning_problem } (\Pi_1 \cup \Pi_2)$
- ⊢ $(\forall \Pi. \text{MEM } \Pi \Pi_l \Rightarrow \text{planning_problem } \Pi) \Rightarrow \text{planning_problem } (\bigcup \Pi_l)$
- ⊢ $\mathcal{D}(\Pi_1 \cup \Pi_2) = \mathcal{D}(\Pi_1) \cup \mathcal{D}(\Pi_2)$

5.4.1 Formalising Soundness of Sub-solution Concatenation (Theorem 17)

We now discuss our formalisation of Lemma 8 from page 60. Before we go into the details of our formalisation, we would like to highlight that in the process of formalisation, we found a mistake in the statement of that lemma as we formulated and published it our International Joint Conference on Artificial Intelligence paper. In the original formulation, we left out the assumption that for every problem Π_i , the corresponding plan $\vec{\pi}_i$ needs to satisfy the condition $\text{sat-pre}(\mathcal{N}(\Pi_i), \vec{\pi}_i)$, i.e. all the actions in $\vec{\pi}_i$ have to have their preconditions satisfied during the execution of $\vec{\pi}_i$ from the needed assignments of Π_i . Without having this assumption, concatenating the plans for the different problems in the problem list does not necessarily satisfy all the goals for all the problems in the list. This is because there may be some action in a plan $\vec{\pi}_i$, whose conditions are not satisfied by $\mathcal{N}(\Pi_i)$, that gets activated in the concatenated plan, thus jeopardising the outcome of the concatenated plan. Implications of missing this assumption in our original published treatment propagated to the theorems that depend on Lemma 8 and eventually to the main algorithm. In some cases this mistake could have led to the computation of buggy plans, i.e. plans that do not reach the goals they are supposed to reach. Interestingly, this case never showed up in any of the thousands of standard planning benchmarks on which we conducted our experiments. Bugs in such corner cases cannot be afforded if the AI algorithm was to be deployed in a safety sensitive application, and the existence of such bugs further strengthens the argument for using formal verification for AI algorithms.

We now go to the details of the formalisation. To formalise the concept of needed assignments, we took a different approach from the one we took in our pen and paper proof. We first define the needed variables as follows, which are equivalent to $\mathcal{D}(\mathcal{N}(\Pi))$ in the original formulation, i.e. the domain of the needed assignments.

HOL4 Definition 18. *Needed Variables*

$$\begin{aligned} \mathcal{N}_{\mathcal{D}}(\Pi) = & \\ \{v \mid & \\ v \in \mathcal{D}(\Pi.I) \wedge & \\ ((\exists p \ e. (p,e) \in \Pi.\delta \wedge v \in \mathcal{D}(p) \wedge p \cdot v = \Pi.I \cdot v) \vee & \\ v \in \mathcal{D}(\Pi.A) \wedge \Pi.I \cdot v = \Pi.A \cdot v)\} & \end{aligned}$$

We then formalise needed assignments as follows:

HOL4 Definition 19 (Needed Assignments).

$$\mathcal{N}(\Pi) = \Pi.I \upharpoonright_{\mathcal{N}_{\mathcal{D}}(\Pi)}$$

The following theorem shows that the definition of the needed assignments has the right semantics. It effectively says that for a problem, a plan will work from any state x that provides the needed assignments of that problem, even if x contradicts the initial state of the problem on the assignments to some state variables. Note the assumption **sat-pre** $(\mathcal{N}(\Pi), \vec{\pi})$: it guarantees that the state x does not unintentionally activate actions that are not active when $\vec{\pi}$ is executed from $\Pi.I$.

HOL4 Lemma 4.

$$\begin{aligned} \vdash \text{planning_problem } \Pi \wedge \mathcal{N}(\Pi) \sqsubseteq x \wedge \text{sat-pre } (\mathcal{N}(\Pi), \vec{\pi}) \Rightarrow \\ \Pi \text{ solved_by } \vec{\pi} \Rightarrow \Pi.A \sqsubseteq \text{ex}(x, \vec{\pi}) \end{aligned}$$

Based on the concept of needed assignments and needed variables, defining the concept of a planning problem that “precedes” another planning problem is straight forward, as follows.

HOL4 Definition 20 (Preceding Problems).

$$\Pi_1 \triangleleft \Pi_2 \iff \Pi_1.A \upharpoonright_{\mathcal{N}_{\mathcal{D}}(\Pi_2)} = \mathcal{N}(\Pi_2) \upharpoonright_{\mathcal{D}(\Pi_1)} \wedge \Pi_1.A \upharpoonright_{\mathcal{D}(\Pi_2)} = \Pi_2.A \upharpoonright_{\mathcal{D}(\Pi_1)}$$

The following HOL4 lemmas show that our definition of the precede relation has the right semantics. It guarantees that if a planning problem Π_1 precedes another planning problem Π_2 , then the a plan for Π_1 always preserves the needed assignments for Π_2 , and a plan for Π_2 does not invalidate a goal for Π_1 . This is crucial to show that the concatenation of the plans of the problems will solve the union of the two problems.

HOL4 Lemma 5.

$$\begin{aligned} \vdash \Pi_1 \triangleleft \Pi_2 \wedge \Pi_1.A \sqsubseteq \text{ex}(x, \vec{\pi}) \wedge \vec{\pi} \in \Pi_1.\delta^* \wedge \mathcal{N}(\Pi_2) \sqsubseteq x \wedge \\ \text{planning_problem } \Pi_1 \Rightarrow \\ \mathcal{N}(\Pi_2) \sqsubseteq \text{ex}(x, \vec{\pi}) \end{aligned}$$

HOL4 Lemma 6.

$$\begin{aligned} \vdash \Pi_1 \triangleleft \Pi_2 \wedge \Pi_2.A \sqsubseteq \text{ex}(x, \vec{\pi}) \wedge \vec{\pi} \in \Pi_2.\delta^* \wedge \text{planning_problem } \Pi_2 \wedge \\ \Pi_1.A \sqsubseteq x \Rightarrow \\ \Pi_1.A \sqsubseteq \text{ex}(x, \vec{\pi}) \end{aligned}$$

Another way the formalisation departs from our pen and paper proof, is in the way we specify the situation of a list of planning problems which all precede each other in the order of the list. We do that in the following recursive way.

HOL4 Definition 21 (List of Preceding Problems).

$$\begin{aligned} \triangleleft_l (\Pi :: \Pi_l) &\iff (\forall \Pi'. \text{MEM } \Pi' \Pi_l \Rightarrow \Pi \triangleleft \Pi') \wedge \triangleleft_l \Pi_l \\ \triangleleft_l [] &\iff \top \end{aligned}$$

We are now ready to state our formalised statement of Lemma 8 from page 60, which states sufficient conditions for the concatenation of sub-problem plans to solve the goals for each of the sub-problems.

HOL4 Lemma 7.

$$\begin{aligned} \vdash \triangleleft_l \Pi_l \Rightarrow \\ \forall x. \\ (\forall \Pi. \\ \text{MEM } \Pi \Pi_l \Rightarrow \\ \text{planning_problem } \Pi \wedge \Pi.l \sqsubseteq x \wedge \Pi \text{ solved_by } f \Pi \wedge \\ \text{sat_pre } (\mathcal{N}(\Pi), f \Pi)) \Rightarrow \\ \forall \Pi. \text{MEM } \Pi \Pi_l \Rightarrow \Pi.A \sqsubseteq \text{ex}(x, \text{FLAT } (\text{MAP } f \Pi_l)) \end{aligned}$$

Note that in the statement above f is a function that maps every planning problem to a plan that solves it. A significant difference between the formalisation of this lemma and the pen and paper treatment is the proof. An informal version of the proof that we formalised follows.

Proof. In the pen and paper proof we proceed by induction on the size of the list of problems, and effectively add a new member to the end of list of problems, which is equivalent to induction on the reverse of Π_l . In the formalisation we proceed in the opposite direction, where the induction adds a new planning problem as a head to the list of problems Π_l . The base case is trivial. In the step case we have the theorem for list of problems Π_l , and we need to show that it applies to Π_l with the problem Π pre-pended to it. The key idea of the proof is to deal with $\bigcup \Pi_l$ as one planning problem. Since Π precedes every problem in Π_l , we have that Π precedes $\bigcup \Pi_l$. From this, the inductive hypothesis, HOL4 Lemma 4, HOL4 Lemma 5, and HOL4 Lemma 6, the result follows. \square

The formalisation of Theorem 17 follows directly from HOL4 Lemma 7. However, we review our formalisation of concepts necessary to state Theorem 17. First we formalise the concepts of a subproblem and based on it the concept of coverage, whose formalisations are straightforward as follows:

HOL4 Definition 22 (Sub-Problem).

$$\Pi_1 \subseteq \Pi_2 \iff \Pi_1.l \sqsubseteq \Pi_2.l \wedge \Pi_1.\delta \subseteq \Pi_2.\delta$$

HOL4 Definition 23. *Covering Problems*

$$\begin{aligned} \text{covers } \Pi_l \Pi &\iff \\ (\forall x. \\ x \in \mathcal{D}(\Pi.A) \Rightarrow \\ \exists \Pi'. \text{MEM } \Pi' \Pi_l \wedge x \in \mathcal{D}(\Pi'.A) \wedge \Pi.A \text{ ' } x = \Pi'.A \text{ ' } x) \wedge \\ \forall \Pi'. \text{MEM } \Pi' \Pi_l \Rightarrow \Pi' \subseteq \Pi \end{aligned}$$

During formalisation, we discovered a mistake in our definition of a subproblem: in the original definition we omitted the requirement that $\Pi_1 . l \sqsubseteq \Pi_2 . l$ holds, and instead required that $\mathcal{D}(\Pi_1) \subseteq \mathcal{D}(\Pi_2)$ holds.

Secondly we formalise the function **rem-condless** as follows. This function removes the actions whose preconditions are not satisfied during execution from a certain state

HOL4 Definition 24 (Remove Condition-less Actions).

$$\begin{aligned} \text{rem-condless } (x, \text{pfx_a}, (p, e) :: \vec{\pi}) &= \\ \text{if } p \sqsubseteq \text{ex}(x, \text{pfx_a}) \text{ then } \text{rem-condless } (x, \text{pfx_a} \# [(p, e)], \vec{\pi}) & \\ \text{else } \text{rem-condless } (x, \text{pfx_a}, \vec{\pi}) & \\ \text{rem-condless } (x, \text{pfx_a}, []) &= \text{pfx_a} \end{aligned}$$

The following theorems show that our formalisation of that function has the desired properties, where the preconditions of all the actions in the resulting action sequence are satisfied, and the resulting action sequence produces the same execution outcome.

$$\begin{aligned} \vdash \text{sat-pre } (x, \text{rem-condless } (x, [], \vec{\pi})) & \\ \vdash \text{ex}(x, \vec{\pi}) &= \text{ex}(x, \text{rem-condless } (x, [], \vec{\pi})) \end{aligned}$$

Given those definitions, the following formalised statement of Theorem 17 follows directly from HOL4 Lemma 7, and it formalises sufficient conditions for the concatenation of sub-problem plans to solve the union of those sub-problems.

HOL4 Theorem 6.

$$\begin{aligned} \vdash \text{covers } \Pi_l \Pi \wedge \triangleleft_l \Pi_l \Rightarrow & \\ (\forall \Pi. \text{MEM } \Pi \Pi_l \Rightarrow \text{planning_problem } \Pi \wedge \Pi \text{ solved_by } f \Pi) \Rightarrow & \\ \Pi \text{ solved_by} & \\ \text{FLAT (MAP } (\lambda \Pi'. \text{rem-condless } (\mathcal{N}(\Pi'), [], f \Pi')) \Pi_l) & \end{aligned}$$

5.4.2 Formalising Instantiation

The most challenging part of formalising Theorem 18 and the augmentation algorithm is in capturing the instantiation of states and building on top of the instantiation of actions, planning problems, and action sequences. Because we formalise states as finite maps, the instantiation $\mathfrak{h}(x)$ of a state x corresponds to applying an image of the instantiation \mathfrak{h} to the domain of the finite map, which is why we denote the instantiation operation as a function image application. For example, for a state $\{o_1 \mapsto T, o_2 \mapsto F\}$ and an instantiation function \mathfrak{h} , the instantiation of that state using that function is the state $\{\mathfrak{h}(o_1) \mapsto T, \mathfrak{h}(o_2) \mapsto F\}$. This is equivalent to composing the inverse of the instantiation function with the state. The first step in formalising this is to define the inverse of a function.

HOL4 Definition 25 (Function Inverse).

$$f^{-1} x = \text{CHOICE } \{y \mid f y = x\}$$

In order for the inverse of a function to behave as we desire, that function needs to be a bijection, in which case the following theorems follow.

$$\vdash \text{BIJ } f \mathcal{U}(:\alpha) \mathcal{U}(:\beta) \Rightarrow f (f^{-1} x) = x$$

$$\vdash \text{BIJ } f \ \mathcal{U}(:\alpha) \ \mathcal{U}(:\beta) \Rightarrow f^{-1} (f \ x) = x$$

$$\vdash \text{BIJ } f \ \mathcal{U}(:\alpha) \ \mathcal{U}(:\beta) \Rightarrow f^{-1-1} = f$$

Note that this bijectivity condition holds for instantiations that are transversals of variable orbits which we considered in Chapter 4. This is because the set of orbits forms a partition of the domain of the planning problem, and transversals map every orbit to one of its members.

Based on this characterisation of an inverse function, the way we formalise state instantiation in HOL4 is as follows:

HOL4 Definition 26 (State Instantiation).

$$\mathfrak{h}(x) = x \ \mathfrak{f_o} \ \mathfrak{h}^{-1}$$

In HOL4 the term $\mathfrak{f}m \ \mathfrak{f_o} \ f$ is the composition of a finite map $\mathfrak{f}m$ with a function f . Based on that, we define the instantiation operation as follows for an action, a factored system, a planning problem and an action sequence.

HOL4 Definition 27 (Action Instantiation).

$$\mathfrak{h}((p, e)) = (\mathfrak{h}(p), \mathfrak{h}(e))$$

HOL4 Definition 28 (System Instantiation).

$$\mathfrak{h}(\delta) = (\lambda \pi. \ \mathfrak{h}(\pi))(\delta)$$

HOL4 Definition 29 (Planning Problem Instantiation).

$$\mathfrak{h}(\Pi) = \Pi \ \text{with} \ \langle |I := \mathfrak{h}(\Pi.I); \ \delta := \mathfrak{h}(\Pi.\delta); \ A := \mathfrak{h}(\Pi.A) | \rangle$$

HOL4 Definition 30 (Action Sequence Instantiation).

$$\mathfrak{h}(\vec{\pi}) = \text{MAP} \ (\lambda \pi. \ \mathfrak{h}(\pi)) \ \vec{\pi}$$

Below are some HOL4 theorems that show some of the properties of instantiations, where $\text{valid_inst } f$ means that f is a bijection.

$$\vdash \text{valid_inst } \mathfrak{h} \Rightarrow \mathfrak{h}(x \ \downarrow_{vs}) = \mathfrak{h}(x) \ \downarrow_{\mathfrak{h}(vs)}$$

$$\vdash \text{valid_inst } \mathfrak{h} \Rightarrow \text{ex}(\mathfrak{h}(x), \mathfrak{h}(\vec{\pi})) = \mathfrak{h}(\text{ex}(x, \vec{\pi}))$$

$$\vdash \text{valid_inst } \mathfrak{h} \wedge \text{planning_problem } \Pi \Rightarrow \text{planning_problem } \mathfrak{h}(\Pi)$$

$$\vdash \text{valid_inst } \mathfrak{h} \wedge \Pi \ \text{solved_by} \ \vec{\pi} \Rightarrow \mathfrak{h}(\Pi) \ \text{solved_by} \ \mathfrak{h}(\vec{\pi})$$

$$\vdash \text{valid_inst } \mathfrak{h} \Rightarrow \mathcal{N}_{\mathcal{D}}(\mathfrak{h}(\Pi)) = \mathfrak{h}(\mathcal{N}_{\mathcal{D}}(\Pi))$$

In addition to being bijections, an additional property needs to hold for a set of instantiations of a planning problem in order for the algorithm of planning via the quotient to work. That property is that two different instantiations from a set of instantiations Δ should not map different variables from the domain of the quotient to the same variable in their range. This is formally stated as follows.

HOL4 Definition 31 (Valid Instantiation).

$$\begin{aligned} \text{valid_instantiations } \Delta \text{ } vs &\iff \\ \forall \mathfrak{h}_1 \mathfrak{h}_2 v_1 v_2. & \\ \mathfrak{h}_1 \in \Delta \wedge \mathfrak{h}_2 \in \Delta \wedge v_1 \in vs \wedge v_2 \in vs \wedge v_1 \neq v_2 \Rightarrow & \\ \mathfrak{h}_1 v_1 \neq \mathfrak{h}_2 v_2 & \end{aligned}$$

This condition guarantees that different instantiations of the same state are consistent with each other, i.e. a certain variable is mapped to the same value in all the instantiated states, which is formally stated in the following theorem.

$$\begin{aligned} \vdash \text{valid_inst } \mathfrak{h}_1 \wedge \text{valid_inst } \mathfrak{h}_2 \wedge \\ \text{valid_instantiations } \{\mathfrak{h}_1; \mathfrak{h}_2\} \mathcal{D}(vs) \Rightarrow \\ \text{agree } \mathfrak{h}_1(vs) \mathfrak{h}_2(vs) \end{aligned}$$

Note that this condition also holds for instantiations that are transversals of variable orbits, again, since the set of orbits forms a partition of the domain of the planning problem.

5.4.3 Formalising Theorem 18 and the Validity of Goal Augmentation

Now that we formalised instantiations, we describe how we formalised Theorem 18. Below are the formalisations of the concepts of common variables between instantiations and the sustainability relation between planning problems and sets of state variables.

HOL4 Definition 32 (Sustainable Variables).

$$\text{sustainable_vars } \Pi \text{ } vs \iff \Pi \cdot I|_{vs} = \Pi \cdot A|_{vs}$$

HOL4 Definition 33 (Common Variables).

$$\begin{aligned} \bigcap_v \Delta \text{ } vs = \\ \{v \mid \exists \mathfrak{h}_1 \mathfrak{h}_2. \mathfrak{h}_1 \in \Delta \wedge \mathfrak{h}_2 \in \Delta \wedge \mathfrak{h}_1 \neq \mathfrak{h}_2 \wedge v \in vs \wedge \mathfrak{h}_1 v = \mathfrak{h}_2 v\} \end{aligned}$$

The following HOL4 lemma summarises the main argument in the proof of Theorem 18, and it shows that the semantics of the definitions above are as intended. It shows that if the needed variables of the quotient are sustained then any two instantiations of that quotient precede each other.

$$\begin{aligned} \vdash \text{valid_inst } \mathfrak{h}_1 \wedge \text{valid_inst } \mathfrak{h}_2 \wedge \\ \text{valid_instantiations } \Delta \mathcal{D}(\Pi) \wedge \text{planning_problem } \Pi \wedge \mathfrak{h}_1 \in \Delta \wedge \\ \mathfrak{h}_2 \in \Delta \wedge \mathfrak{h}_1 \neq \mathfrak{h}_2 \wedge \text{sustainable_vars } \Pi (\bigcap_v \Delta \mathcal{D}(\Pi) \cap \mathcal{N}_{\mathcal{D}}(\Pi)) \Rightarrow \\ \mathfrak{h}_1(\Pi) \triangleleft \mathfrak{h}_2(\Pi) \end{aligned}$$

That lemma can then be generalised to a list of instantiations as follows, where similar sufficient conditions on lists of planning problems that precede each other are formalised.

$$\begin{aligned} \vdash \text{ALL-DISTINCT } \Delta \wedge (\forall \mathfrak{h}. \text{MEM } \mathfrak{h} \Delta \Rightarrow \text{valid_inst } \mathfrak{h}) \wedge \\ \text{valid_instantiations } (\text{set } \Delta) \mathcal{D}(\Pi) \wedge \text{planning_problem } \Pi \wedge \\ \text{sustainable_vars } \Pi (\bigcap_v (\text{set } \Delta) \mathcal{D}(\Pi) \cap \mathcal{N}_{\mathcal{D}}(\Pi)) \Rightarrow \\ \triangleleft_l (\text{MAP } (\lambda \mathfrak{h}. \mathfrak{h}(\Pi)) \Delta) \end{aligned}$$

From this and from HOL4 Theorem 6, we can derive the following theorem. It is a formalisation of Theorem 18, the main theorem that states sufficient soundness conditions for planning via the descriptive quotient.

HOL4 Theorem 7.

$$\begin{aligned}
&\vdash \text{ALL-DISTINCT } \Delta \wedge \\
&\text{INJ } (\lambda \hbar. \hbar(\Pi')) (\text{set } \Delta) \mathcal{U}(:(\alpha, \beta) \text{ planningProblem}) \wedge \\
&(\forall \hbar. \text{MEM } \hbar \Delta \Rightarrow \text{valid_inst } \hbar) \wedge \\
&\text{valid_instantiations } (\text{set } \Delta) \mathcal{D}(\Pi') \wedge \text{planning_problem } \Pi' \wedge \\
&\text{sustainable_vars } \Pi' (\bigcap_v (\text{set } \Delta) \mathcal{D}(\Pi') \cap \mathcal{N}_{\mathcal{D}}(\Pi')) \wedge \\
&\text{covers } (\text{MAP } (\lambda \hbar. \hbar(\Pi')) \Delta) \Pi \wedge \Pi' \text{ solved_by } \vec{\pi}' \Rightarrow \\
&\quad \Pi \text{ solved_by} \\
&\text{FLAT } (\text{MAP } (\lambda \hbar. \text{rem-condless } (\mathcal{N}(\hbar(\Pi')), [], \hbar(\vec{\pi}')))) \Delta)
\end{aligned}$$

The formalised version of Theorem 18 is more general in that it does not assume that the quotient problem is necessarily a descriptive quotient obtained through symmetry analysis. It rather assumes that the quotient can be instantiated into subproblems of the concrete planning problem that cover its goals, i.e. the quotient is isomorphic to subproblems of the given planning problem. It also should be clear that the concrete planning problem and the quotient are not necessarily propositionally factored systems, i.e. the state variables are not necessarily Boolean and are not required to have finite or countable domains. This makes the planning via descriptive quotients a candidate approach to be used with hybrid systems.

The last step in our formalisation is to verify that instantiations of a plan for a quotient whose goal is augmented, as per our approach described in Section 4.6, can be concatenated into a plan for the concrete planning problem. The first step in doing so is to show that augmenting the quotient's goal with the common needed assignments guarantees that it sustains those assignments. This is shown in the following theorem.

$$\begin{aligned}
&\vdash \text{let} \\
&\quad \Pi^q = \Pi' \text{ with } A := \Pi'.I|_{\bigcap_v (\text{set } \Delta) \mathcal{D}(\Pi') \cap \mathcal{N}_{\mathcal{D}}(\Pi')} \uplus \Pi'.A \\
&\text{in} \\
&\quad \text{sustainable_vars } \Pi^q (\bigcap_v (\text{set } \Delta) \mathcal{D}(\Pi^q) \cap \mathcal{N}_{\mathcal{D}}(\Pi^q))
\end{aligned}$$

It is worth noting that HOL4 is able to prove the above goal entirely automatically.

Now that we proved that goal augmentation works as intended, the following theorem follows, which verifies the validity of our approach of planning using the descriptive quotient. The main idea in that theorem is that, since a planning problem with an augmented goal sustains the common needed variables, as we show in the lemma above, then the same algorithm as in HOL4 Theorem 7 can always be used the augmented quotient problem.

$$\begin{aligned}
&\vdash \text{let} \\
&\quad \Pi^q = \Pi' \text{ with } A := \Pi'.I|_{\bigcap_v (\text{set } \Delta) \mathcal{D}(\Pi') \cap \mathcal{N}_{\mathcal{D}}(\Pi')} \uplus \Pi'.A \\
&\text{in} \\
&\quad \text{ALL-DISTINCT } \Delta \wedge (\forall \hbar. \text{MEM } \hbar \Delta \Rightarrow \text{valid_inst } \hbar) \wedge \\
&\quad \text{planning_problem } \Pi' \Rightarrow \\
&\quad \text{INJ } (\lambda \hbar. \hbar(\Pi^q)) (\text{set } \Delta) \mathcal{U}(:(\alpha, \beta) \text{ planningProblem}) \wedge \\
&\quad \text{valid_instantiations } (\text{set } \Delta) \mathcal{D}(\Pi^q) \wedge \\
&\quad \text{covers } (\text{MAP } (\lambda \hbar. \hbar(\Pi^q)) \Delta) \Pi \wedge \Pi^q \text{ solved_by } \vec{\pi}^q \Rightarrow \\
&\quad \quad \Pi \text{ solved_by} \\
&\quad \text{FLAT } (\text{MAP } (\lambda \hbar. \text{rem-condless } (\mathcal{N}(\hbar(\Pi^q)), [], \hbar(\vec{\pi}^q)))) \Delta)
\end{aligned}$$

5.5 Concluding Remarks

Much of the previous work on formalised algorithms related to transition systems, like Esparza et al. (2013), focus on algorithms that target model checking applications. Although the algorithms

that we introduce and formally verify here apply to *propositionally factored transition systems* in their full generality, we focus more on AI planning applications. Thus, with this work, we believe we have launched a fruitful inter-disciplinary collaboration between the fields of AI planning and mechanised mathematics. From the interactive theorem-proving community’s point of view, it is gratifying to be able to find and fix errors in the modern research literature. For example, the insights which led us to develop the sublist diameter and the top-down algorithm followed from attempting to formalise the results in Rintanen and Gretton (2013). That effort helped us find a bug in their algorithm, where they incorrectly theorise that the diameter can be compositionally upper-bounded by projections’ diameters, if the projections are induced by a partition of state variables whose members are closed under mutual dependency.

Also during formalisation, we found mistakes in our own algorithms that are subtle but non-trivial. For instance, in our published version and implementation of the algorithm for planning via the descriptive quotient, we missed an important step in the algorithm. That step is removing actions from the quotient’s plan whose preconditions are not satisfied, before the instantiation and concatenation steps of the quotient’s plan, which can possibly produce an invalid plan. Interestingly, this error never showed up during experimentation, although we ran our algorithm on thousands of diverse standard planning benchmarks. Although such errors can be practically rare, they cannot be tolerated in safety critical applications, thus, their existence makes a strong case for the utility of mechanical verification.

It is vital that we give AI researchers assurances that their algorithms, theory and systems are correct. For AI planning systems to be deployed in safety critical applications and for autonomous exploration of space, they must not only be efficient, and provably conservative in their resource consumption, but also correct. The upper-bounding algorithms like the ones we formalise here underpin fixed-horizon planning; indeed, they provide the fixed horizon past which an algorithm need not search. If an autonomous vehicle exploring outer space implemented diameter-based compositional bounding suggested by Rintanen and Gretton (2013) to do its plan-search, that system could incorrectly conclude that no plan exists.

Another point we would like to raise is one concerning the difference between the formal proofs and the informal proofs. We oftentimes found that some proofs were more intuitive and accordingly easier to explain to deliver the gist of how a proof works. In contrast, we sometimes found it very hard to formalise those proofs and found that substantially different approaches were more natural and more fit for formalisation. An example of that is the proof of Lemma 8 and its formalised equivalent HOL4 Lemma 7. We found that to deliver the argument, doing the induction on the tail of the list is more natural and clearer, while it was substantially much easier to formalise the proof by doing the induction on the head of the list.

Furthermore, our formalisation made it much easier to generalise some of our results and algorithms from propositionally factored systems to more general systems, systems in which the codomains of states are not necessarily Boolean, finite or even countable. This raises the possibility of applying the algorithms that we developed to hybrid systems. Also, we would like to note that for our formalisation we developed a large formalised library on factored transition systems. Since a lot of theory in that library applies to factored systems that are not propositionally factored, it can be used for verifying algorithms on hybrid systems, for instance. In the case of hybrid systems, an interesting challenge would be extending the theory we developed to be capable of representing actions whose preconditions and effects are functions in state variables, versus assignments to state variables.

Our experience mechanising results in AI planning allows us to provide insights regarding the scalability of formalising AI planning algorithms. To formalise the compositional algorithms we developed a library of HOL4 proof scripts that is around 14k lines long, including comments.

Interestingly, the size of developed code does not necessarily scale linearly with the number of algorithms that we formalised. The first algorithm that we formalised was N_{sum} , and to do so we developed around $10k$ lines of proof script. That script was developed in approximately six months. Our subsequent formalisation of S_{max} and HYB required an additional $2k$ lines of proof scripts each. Each of those algorithms took around two and half weeks to formalise. This productivity improvement follows because when we formalised N_{sum} , we developed the majority of the needed formal background theory. That theory is leveraged in our formalisation of other algorithms on factored transition systems.

We made a number of observations in our efforts that we believe provide insight into how HOL4 can be improved. The feature of HOL4 that we would cite as the most positive, is the ability to quickly modify existing tactics, or add new tactics, since the entire system is completely implemented in SML. Also, automation tactics in general are reasonable. Nonetheless, we think that other aspects of automation can still be improved. Tasks that can be automated which we found cumbersome include: (i) searching for theorems in the library, (ii) the generation of termination conditions, (iii) the definition of relations, (iv) proving the uniqueness of relations, and (v) deriving the function form of a relation. Another more general issue that we faced is the absence of a mechanism akin to classes in Isabelle, which could have allowed us to reduce the repetition of theorem hypotheses.

In terms of formally verifying AI planning algorithms, we have only scratched the surface here. For instance, when a tight bound for a planning problem is known, one effective technique for finding a plan is to reduce that problem to SAT (Rintanen, 2012). Proposed reductions are constructive, in the sense that a plan can be constructed in linear time from a satisfying assignment to a formula. A key recent advance in the setting of planning-via-SAT has been the development of compact SAT-representations of planning problems. Such representations facilitate more efficient search (Rintanen, 2012; Robinson et al., 2009). In future work, we would like to verify the correctness of both the reductions to SAT, and the algorithms that subsequently construct plans from a satisfying assignment.

Bibliography

- Amir Abboud, Virginia Vassilevska Williams, and Joshua Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete Algorithms*, pages 377–391. SIAM, 2016.
- Mohammad Abdulaziz and Lawrence C Paulson. An Isabelle/HOL formalisation of Green’s theorem. In *International Conference on Interactive Theorem Proving*, pages 3–19. Springer, 2016.
- Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing*, 28(4): 1167–1181, 1999.
- Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *Journal of the ACM (JACM)*, 42(4): 844–856, 1995.
- Noga Alon, Zvi Galil, and Oded Margalit. On the exponent of the all pairs shortest path problem. *Journal of Computer and System Sciences*, 54(2):255–262, 1997.
- Rajeev Alur, Robert K Brayton, Thomas A Henzinger, Shaz Qadeer, and Sriram K Rajamani. Partial-order reduction in symbolic state space exploration. In *International Conference on Computer Aided Verification*, pages 340–351. Springer, 1997.
- Eyal Amir and Barbara Engelhardt. Factored planning. In *IJCAI*, volume 3, pages 929–935. Citeseer, 2003.
- Andreas Bauer, Peter Baumgartner, Martin Diller, and Michael Norrish. Tableaux for verification of data-centric processes. In *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 28–43. Springer, 2013.
- Jason Baumgartner, Andreas Kuehlmann, and Jacob Abraham. Property checking via structural analysis. In *Computer Aided Verification*, pages 151–165. Springer, 2002.
- Richard Bellman. On a routing problem. *Quarterly of applied mathematics*, pages 87–90, 1958.
- Sergey Berezin, Sérgio Campos, and Edmund M Clarke. Compositional reasoning in model checking. In *Compositionality: The Significant Difference*, pages 81–102. Springer, 1998.
- Yves Bertot and Pierre Castéran. *Interactive theorem proving and program development: Coq’Art: the calculus of inductive constructions*. springer, 2004.
- Marc Bezem and Dimitri Hendriks. On the mechanization of the proof of Hessenberg’s theorem in coherent logic. *Journal of Automated Reasoning*, 40(1):61–85, 2008.
- Armin Biere, Alessandro Cimatti, Edmund M. Clarke, and Yunshan Zhu. Symbolic model checking without BDDs. In *TACAS*, pages 193–207, 1999.

-
- Armin Biere, Alessandro Cimatti, Edmund M. Clarke, Ofer Strichman, and Yunshan Zhu. Bounded model checking. *Advances in Computers*, 58:117–148, 2003.
- Andreas Björklund and Thore Husfeldt. Finding a path of superlogarithmic length. *SIAM Journal on Computing*, 32(6):1395–1402, 2003.
- Andreas Björklund, Thore Husfeldt, and Sanjeev Khanna. Approximating longest directed paths and cycles. In *International Colloquium on Automata, Languages, and Programming*, pages 222–233. Springer, 2004.
- Jasmin Christian Blanchette and Tobias Nipkow. Nitpick: A counterexample generator for higher-order logic based on a relational model finder. In *Interactive Theorem Proving, First International Conference, ITP 2010*, pages 131–146, 2010. doi: 10.1007/978-3-642-14052-5_11.
- Jasmin Christian Blanchette, Mathias Fleury, and Christoph Weidenbach. A verified SAT solver framework with learn, forget, restart, and incrementality. In *International Joint Conference on Automated Reasoning*, pages 25–44. Springer, 2016.
- Avrim L Blum and Merrick L Furst. Fast planning through planning graph analysis. *Artificial intelligence*, 90(1):281–300, 1997.
- Blai Bonet and Hector Geffner. Planning as heuristic search. *Artificial Intelligence*, 129(1-2): 5–33, 2001.
- R. I. Brafman and C. Domshlak. Structure and complexity in planning with unary operators. *Journal of Artificial Intelligence Research*, 18:315–349, 2003.
- Ronen I Brafman and Carmel Domshlak. Factored planning: How, when, and when not. In *AAAI*, volume 6, pages 809–814, 2006.
- Cynthia A. Brown, Larry Finkelstein, and Paul Walton Purdom Jr. Backtrack searching in the presence of symmetry. *Nord. J. Comput.*, 3(3):203–219, 1996.
- Daniel Bundala, Joël Ouaknine, and James Worrell. On the magnitude of completeness thresholds in bounded model checking. In *Proceedings of the 2012 27th Annual IEEE/ACM Symposium on Logic in Computer Science*, pages 155–164. IEEE Computer Society, 2012.
- T. Bylander. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69(1-2):165–204, 1994.
- Michael L. Case, Hari Mony, Jason Baumgartner, and Robert Kanzelman. Enhanced verification by temporal decomposition. In *FMCAD 2009, 15-18 November 2009, Austin, Texas, USA*, pages 17–24, 2009.
- Timothy M Chan. More algorithms for all-pairs shortest paths in weighted graphs. *SIAM Journal on Computing*, 39(5):2075–2089, 2010.
- Shiri Chechik, Daniel H Larkin, Liam Roditty, Grant Schoenebeck, Robert E Tarjan, and Virginia Vassilevska Williams. Better approximation algorithms for the graph diameter. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1041–1052. Society for Industrial and Applied Mathematics, 2014.
- Yixin Chen and Guohui Yao. Completeness and optimality preserving reduction for planning. In *IJCAI*, pages 1659–1664, 2009.

-
- Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Progress on the state explosion problem in model checking. In *Informatics*, pages 176–194. Springer, 2001.
- Edmund Clarke, Daniel Kroening, Joël Ouaknine, and Ofer Strichman. Completeness and complexity of bounded model checking. In *International Workshop on Verification, Model Checking, and Abstract Interpretation*, pages 85–96. Springer, 2004.
- Edmund M Clarke, Orna Grumberg, and David E Long. Model checking and abstraction. *ACM transactions on Programming Languages and Systems (TOPLAS)*, 16(5):1512–1542, 1994.
- Edmund M Clarke, Reinhard Enders, Thomas Filkorn, and Somesh Jha. Exploiting symmetry in temporal logic model checking. *Formal methods in system design*, 9(1-2):77–104, 1996.
- Edmund M Clarke, E Allen Emerson, Somesh Jha, and A Prasad Sistla. Symmetry reductions in model checking. In *Computer Aided Verification*, pages 147–158. Springer, 1998.
- Edmund M Clarke, E Allen Emerson, and Joseph Sifakis. Turing lecture: model checking—algorithmic verification and debugging. *Communications of the ACM*, 52(11):74–84, 2009.
- Robert L Constable, Paul B Jackson, Pavel Naumov, and Juan C Uribe. Constructively formalizing automata theory. In *Proof, language, and interaction*, pages 213–238, 2000.
- James Crawford, Matthew Ginsberg, Eugene Luks, and Amitabha Roy. Symmetry-breaking predicates for search problems. *KR*, 96:148–159, 1996.
- Peter Dankelmann. The diameter of directed graphs. *Journal of Combinatorial Theory, Series B*, 94(1):183–186, 2005.
- Peter Dankelmann and Michael Dorfling. Diameter and maximum degree in Eulerian digraphs. *Discrete Mathematics*, 339(4):1355–1361, 2016.
- Peter Dankelmann and Lutz Volkmann. The diameter of almost Eulerian digraphs. *the electronic journal of combinatorics*, 17(1):R157, 2010.
- Leonardo De Moura and Nikolaj Bjørner. Z3: An Efficient SMT Solver. In *Proc. TACAS’08/ETAPS’08*, pages 337–340, 2008.
- Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische mathematik*, 1(1):269–271, 1959.
- Christian Doczkal, Jan-Oliver Kaiser, and Gert Smolka. A constructive theory of regular languages in Coq. In *International Conference on Certified Programs and Proofs*, pages 82–97. Springer, 2013.
- Carmel Domshlak, Michael Katz, and Alexander Shleyfman. Enhanced symmetry breaking in cost-optimal planning as forward search. In *ICAPS*, 2012.
- Carmel Domshlak, Michael Katz, and Alexander Shleyfman. Symmetry breaking: Satisficing planning and landmark heuristics. In *ICAPS*, 2013.
- Jean-François Dufourd. An intuitionistic proof of a discrete form of the Jordan Curve Theorem formalized in Coq with combinatorial hypermaps. *Journal of Automated Reasoning*, 43(1): 19–51, 2009.

-
- E Allen Emerson and A Prasad Sistla. Symmetry and model checking. *Formal methods in system design*, 9(1):105–131, 1996.
- Paul Erdős, Janos Pach, Richard Pollack, and Zsolt Tuza. Radius, diameter, and minimum degree. *Journal of Combinatorial Theory, Series B*, 47(1):73–79, 1989.
- Javier Esparza, Peter Lammich, René Neumann, Tobias Nipkow, Alexander Schimpf, and Jan-Georg Smaus. A fully verified executable LTL model checker. In *International Conference on Computer Aided Verification*, pages 463–478. Springer, 2013.
- M Eugene. Permutation groups and polynomial-time computation. In *Groups and Computation: Workshop on Groups and Computation, October 7-10, 1991*, volume 11, page 139. American Mathematical Soc., 1993.
- Richard E Fikes and Nils J Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4):189–208, 1971.
- Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin. Antichains and compositional algorithms for LTL synthesis. *Formal Methods in System Design*, 39(3):261–296, 2011.
- Maria Fox and Derek Long. The detection and exploitation of symmetry in planning problems. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI 99, Stockholm, Sweden, July 31 - August 6, 1999. 2 Volumes, 1450 pages*, pages 956–961, 1999.
- Maria Fox and Derek Long. Extending the exploitation of symmetries in planning. In *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems, April 23-27, 2002, Toulouse, France*, pages 83–91, 2002.
- Michael L Fredman. New bounds on the complexity of the shortest path problem. *SIAM Journal on Computing*, 5(1):83–89, 1976.
- Michael L Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM (JACM)*, 34(3):596–615, 1987.
- Ruben A Gamboa. A formalization of powerlist algebra in ACL2. *Journal of Automated Reasoning*, 43(2):139–172, 2009.
- Rob Gerth, Ruurd Kuiper, Doron Peled, and Wojciech Penczek. A partial order approach to branching time logic model checking. In *Theory of Computing and Systems, 1995. Proceedings., Third Israel Symposium on the*, pages 130–139. IEEE, 1995.
- Patrice Godefroid. Using partial orders to improve automatic verification methods. In *International Conference on Computer Aided Verification*, pages 176–185. Springer, 1990.
- Georges Gonthier. Formal proof—the four-color theorem. *Notices of the AMS*, 55(11):1382–1393, 2008.
- Michael JC Gordon and Tom F Melham. *Introduction to HOL: a theorem proving environment for higher order logic*. Cambridge University Press, 1993.
- Emmanuel Guere and Rachid Alami. One action is enough to plan. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001, Seattle, Washington, USA, August 4-10, 2001*, pages 439–444, 2001.

-
- Thomas C Hales, John Harrison, Sean McLaughlin, Tobias Nipkow, Steven Obua, and Roland Zumkeller. A revision of the proof of the Kepler conjecture. In *The Kepler Conjecture*, pages 341–376. Springer, 2011.
- John Harrison. Formalizing basic complex analysis. *From Insight to Proof: Festschrift in Honour of Andrzej Trybulec. Studies in Logic, Grammar and Rhetoric*, 10(23):151–165, 2007.
- Malte Helmert. The Fast Downward planning system. *Journal of Artificial Intelligence Research*, 26:191–246, 2006a.
- Malte Helmert, Patrik Haslum, Jörg Hoffmann, and Raz Nissim. Merge-and-shrink abstraction: A method for generating lower bounds in factored state spaces. *Journal of the ACM (JACM)*, 61(3):16, 2014.
- Jörg Hoffmann and Bernhard Nebel. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.
- Daniel Jackson. *Software Abstractions: Logic, Language, and Analysis*. MIT Press, 2006. ISBN 978-0-262-10114-1.
- Mathieu Jaume. A full formalization of SLD-resolution in the calculus of inductive constructions. *Journal of Automated Reasoning*, 23(3):347–371, 1999.
- David Joslin and Amitabha Roy. Exploiting symmetry in lifted CSPs. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Innovative Applications of Artificial Intelligence Conference, AAAI 97, IAAI 97, July 27-31, 1997, Providence, Rhode Island.*, pages 197–202, 1997.
- Fairouz Kamareddine and Haiyan Qiao. Formalizing strong normalization proofs of explicit substitution calculi in ALF. *Journal of Automated Reasoning*, 30(1):59–98, 2003.
- H. A. Kautz and B. Selman. Planning as satisfiability. In *ECAI*, pages 359–363, 1992.
- Elena Kelareva, Olivier Buffet, Jinbo Huang, and Sylvie Thiébaux. Factored planning using decomposition trees. In *IJCAI*, pages 1942–1947, 2007.
- Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, et al. seL4: Formal verification of an OS kernel. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, pages 207–220. ACM, 2009.
- C. A. Knoblock. Automatically generating abstractions for planning. *Artificial Intelligence*, 68(2):243–302, 1994.
- AV Knyazev. Diameters of pseudosymmetric graphs. *Mathematical Notes*, 41(6):473–482, 1987.
- Daniel Kroening. Computing over-approximations with bounded model checking. *Electronic Notes in Theoretical Computer Science*, 144(1):79–92, 2006.
- Daniel Kroening and Ofer Strichman. Efficient computation of recurrence diameters. In *VMCAI*, pages 298–309, 2003.
- Daniel Kroening, Joël Ouaknine, Ofer Strichman, Thomas Wahl, and James Worrell. Linear completeness thresholds for bounded model checking. In *Computer Aided Verification*, pages 557–572. Springer, 2011.

-
- Panagiotis Manolios and Daron Vroon. Ordinal arithmetic: Algorithms and mechanization. *Journal of Automated Reasoning*, 34(4):387–423, 2005.
- Filip Marić. Formalization and implementation of modern SAT solvers. *Journal of Automated Reasoning*, 43(1), 2009.
- Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL: The Planning Domain Definition Language. Technical report, CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.
- Drew V McDermott. A heuristic estimator for means-ends analysis in planning. In *AIPS*, volume 96, pages 142–149, 1996.
- Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, {II}. *Journal of Symbolic Computation*, 60(0):94 – 112, 2014. ISSN 0747-7171. doi: <http://dx.doi.org/10.1016/j.jsc.2013.09.003>. URL <http://www.sciencedirect.com/science/article/pii/S0747717113001193>.
- James McKinna and Robert Pollack. Some lambda calculus and type theory formalized. *Journal of Automated Reasoning*, 23(3):373–409, 1999.
- Kenneth L McMillan. Symbolic model checking. In *Symbolic Model Checking*, pages 25–60. Springer, 1993.
- Ian Miguel. Symmetry-breaking in planning: Schematic constraints. In *Proceedings of the CP'01 Workshop on Symmetry in Constraints*, pages 17–24, 2001.
- John W Moon et al. On the diameter of a graph. *The Michigan Mathematical Journal*, 12(3): 349–351, 1965.
- Tobias Nipkow. Verifying a hotel key card system. In K. Barkaoui, A. Cavalcanti, and A. Cerone, editors, *Theoretical Aspects of Computing (ICTAC 2006)*, volume 4281 of *Lecture Notes in Computer Science*. Springer, 2006. Invited paper.
- Tobias Nipkow and Lawrence C Paulson. Proof pearl: Defining functions over finite sets. In *International Conference on Theorem Proving in Higher Order Logics*, pages 385–396. Springer, 2005.
- Panos M Pardalos and Athanasios Migdalas. A note on the complexity of longest path problems related to graph coloring. *Applied mathematics letters*, 17(1):13–15, 2004.
- Lawrence C Paulson. *Isabelle: A generic theorem prover*, volume 828. Springer, 1994.
- Lawrence C Paulson. A formalisation of finite automata using hereditarily finite sets. In *International Conference on Automated Deduction*, pages 231–245. Springer, 2015.
- Nir Pochter, Aviv Zohar, and Jeffrey S Rosenschein. Exploiting problem symmetries in state-based planners. In *AAAI*, 2011.
- Aldo Porco, Alejandro Machado, and Blai Bonet. Automatic reductions from PH into STRIPS or how to generate short problems with very long solutions. In *ICAPS*, 2013.
- Silvia Richter and Matthias Westphal. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research*, 39(1):127–177, 2010.

-
- Jussi Rintanen. Symmetry reduction for SAT representations of transition systems. In *ICAPS*, pages 32–41, 2003.
- Jussi Rintanen. Planning as satisfiability: Heuristics. *Artificial Intelligence*, 193:45–86, 2012.
- Jussi Rintanen and Charles Orgill Gretton. Computing upper bounds on lengths of transition sequences. In *International Joint Conference on Artificial Intelligence*, 2013.
- Nathan Robinson, Charles Gretton, Duc Nghia Pham, and Abdul Sattar. SAT-based parallel planning using a split representation of actions. In *ICAPS*, 2009.
- Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 515–524. ACM, 2013.
- Alexander Schimpf, Stephan Merz, and Jan-Georg Smaus. Construction of Büchi automata for LTL model checking verified in Isabelle/HOL. In *International Conference on Theorem Proving in Higher Order Logics*, pages 424–439. Springer, 2009.
- Jendrik Seipp, Florian Pommerening, Silvan Sievers, Martin Wehrle, Chris Fawcett, and Yusra Alkhazraji. Fast Downward Aidos. *International Planning Competition*, 2014.
- Mary Sheeran, Satnam Singh, and Gunnar Stålmarmark. Checking safety properties using induction and a SAT-solver. In *Formal Methods in Computer-Aided Design, Third International Conference, FMCAD 2000, Austin, Texas, USA, November 1-3, 2000, Proceedings*, pages 108–125, 2000. doi: 10.1007/3-540-40922-X_8.
- Alexander Shleyfman, Michael Katz, Malte Helmert, Silvan Sievers, and Martin Wehrle. Heuristics and symmetries in classical planning. In *AAAI*, 2015.
- Silvan Sievers, Martin Wehrle, Malte Helmert, Alexander Shleyfman, and Michael Katz. Factored symmetries for merge-and-shrink abstractions. In *Proc. 29th National Conf. on Artificial Intelligence*. AAAI Press, 2015.
- A Prasad Sistla and Edmund M Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM (JACM)*, 32(3):733–749, 1985.
- Konrad Slind and Michael Norrish. A brief overview of HOL4. In *Theorem Proving in Higher Order Logics*, volume 5170 of *LNCS*, pages 28–32. Springer, 2008.
- José Soares. Maximum diameter of regular digraphs. *Journal of Graph Theory*, 16(5):437–450, 1992.
- Christoph Sprenger. A verified model checker for the modal μ -calculus in Coq. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 167–183. Springer, 1998.
- Paul Taylor. The fixed point property in synthetic domain theory. In *Logic in Computer Science, 1991. LICS'91., Proceedings of Sixth Annual IEEE Symposium on*, pages 152–160. IEEE, 1991.
- Mikkel Thorup. Integer priority queues with decrease key in constant time and the single source shortest paths problem. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 149–158. ACM, 2003.

- Thomas Wahl and Alastair F. Donaldson. Replication and abstraction: Symmetry in automated formal verification. *Symmetry*, 2(2):799–847, 2010. doi: 10.3390/sym2020799. URL <http://dx.doi.org/10.3390/sym2020799>.
- Brian C. Williams and P. Pandurang Nayak. A reactive planner for a model-based executive. In *International Joint Conference on Artificial Intelligence*, pages 1178–1185. Morgan Kaufmann Publishers, 1997.
- Chunhan Wu, Xingyuan Zhang, and Christian Urban. A formalisation of the Myhill-Nerode theorem based on regular expressions (proof pearl). In *International Conference on Interactive Theorem Proving*, pages 341–356. Springer, 2011.
- You Xu, Yixin Chen, Qiang Lu, and Ruoyun Huang. Theory and algorithms for partial order based reduction in planning. *arXiv preprint arXiv:1106.5427*, 2011.
- Raphael Yuster. Computing the diameter polynomially faster than APSP. *arXiv preprint arXiv:1011.6181*, 2010.