Chapter 2

Background and related work

Over the last few decades, sensor modalities and methodologies for navigation of robots over land have developed and propagated rapidly when compared to its underwater counterparts. Apart from the challenges presented by the harsher environment for the mechanical aspects of robots, the underwater media itself presents a number of considerable challenges in localisation, navigation and communication for autonomous underwater vehicles (AUVs). This chapter gives a brief insight in to the background of localisation technology, the drivers behind the choice of sensing strategies and methodologies while focusing on work related to the research presented in this thesis.

The following section gives a general introduction to the sensor modalities available for underwater applications with regard to robot localisation. Next, different strategies used to perform relative localisation in multi-robot setups are reviewed with an emphasis on 'real-world' implementations addressing the problem of simultaneous navigation of multiple AUVs. The subsequent section discusses some sensor utilisation strategies for localisation available in the literature and their applicability in the localisation system being developed in this thesis. Finally, drawbacks and benefits of existing underwater localisation systems with respect to the constraints and requirements of the motivating application are discussed.

2.1 Localisation

Localisation, or position fixing, is a topic of interest covered by many diverse research fields. Animal hearing/binaural localisation (Wallach, 1938; Konishi, 1993; Roman and DeLiang, 2003; Stern et al., 2006), Acoustic source localisation/Speaker tracking (Svaizer et al., 1997;Benesty, 2000; Lehmann, 2004), Sensor network localisation (Ajdler et al., 2004; Priyantha, 2005; Mao et al., 2007), target motion analysis (Altes, 1979; Farina, 1999; Arulampalam et al., 2004) and mobile robot navigation (Thrun et al., 2001; Howard et al., 2003; Valin et al., 2003; Kenn and Pfeil, 2004) are some of them. The research presented in this thesis draws insights from many of these areas in designing and developing the relative localisation system. The choice of sensing strategy and methodology mainly depend on the constraints (size, weight, power budget *etc.*), requirements (range, accuracy/precision, update rate *etc.*) of the application and the operating environment (structuredness, degree of clutter, medium *i.e.* in-air, underwater *etc.*). For example, underwater electrolocation schemes such as those presented by Solberg et al. (2008) are only viable for very short ranges (less than 0.5 m). While visual localisation methods (Dellaert et al., 1999b; Huster, 2003) require well lit (or artificially lit) environments, the use of sensors such as laser range finders are more suited to structured environments with laser reflective surfaces.

While localisation schemes involving the electromagnetic spectrum have become ubiquitous for all forms of in-air applications[1], in the underwater domain these are unavailable due to the high rate of attenuation of electromagnetic waves in water. For example, global positioning system (GPS) signals experience an attenuation of well over 50 dB at 1.0 m depth in sea water[2]. Ergo, vehicles operating underwater are deprived of access to navigational aids and wireless communication methods using high frequency electromagnetic waves.

Early localisation systems

During the mid 20th century, maritime and later aircraft navigation used localisation schemes utilising low frequency (LF) and very low frequency (VLF) electromagnetic waves. Among these were the Decca Navigator System, Omega Navigation System and the LORAN system to name a few (Palmer, 1970; Kasper and Hutchinson, 1978; Klepczynski, 1983; Last, 1989). By now these systems (except for the LORAN-c variant which is also in decline) have been superseded by GPS based navigation systems (Kaplan and Hegarty, 2006; Hofmann-Wellenhof et al., 2001). Most of the aforementioned obsolete systems used hyperbolic localisation (also known as multilateration) which operated either by locating the position of a single transmitter using

---

[1] *See Kayton (1988) for a comprehensive survey article on the historical development of navigation technology.*

[2] *GPS signals use a frequency of approximately* $1.5 GHz$ *and electromagnetic loss in sea water is given by* $1400\sqrt{f}\ dBkm^{-1}$ *where* f *is in* kHz *(Waite, 2002).*

multiple receivers (at least three) or by locating the position of a single receiver using at least three synchronised transmitters. Despite the fact that these systems are non-operational, the developed localisation concepts have been adopted by acoustic beacon localisation schemes that are currently used for subsea navigation.

2.2 Underwater localisation

The following sections give a brief evolution of underwater localisation methodologies. While this section concentrates mostly on single vehicle localisation schemes, later sections will elaborate on how some of these methods have been extended to facilitate multi-robot localisation and navigation in the underwater domain.

2.2.1 Sonar based localisation

While the concept of target detection in the underwater environment using acoustics can be traced back to a late 15th century postulate by Leonardo da Vinci[1] (MacCurdy, 1948), practical systems came into use only in the early 20th century. The earliest adoption of this method was sonar (sound navigation and ranging) in its passive and active forms which performed very similar to radar (radio detection and ranging) in target detection, ranging and mapping applications (Altes, 1979). Though sonar was primarily developed for underwater applications, its principles have also been successfully used in air to complement radar systems as well. Sonar sensors have also been successfully used in mobile robotic applications to aid navigation (Elfes, 1987). As the role of electromagnetic waves in air is taken over by acoustics in water, Burdic (1984) gives a comprehensive chapter on the historical developments and technologies that led to modern day sonar systems. In addition see Nielsen (1991),Waite (2002) and Ricker (2003) for more detailed descriptions about these concepts and Etter (2003) regarding further mathematical treatment for both active and passive sonar operation.

In his survey of different localisation and map building methods, Thrun (2002) presents a number of approaches used by land based robots deployed mostly in static and structured environments. As he points out, for the dynamic and unstructured environments in the underwater domain, most of these techniques are inadequate. However, terrain sensing sonar based mapping and localisation has been successfully implemented and demonstrated for AUV navigation. In this regard, Feder, et al. (1998), presents a concurrent mapping and localisation algorithm and tests its long term performance using simulated forward looking sonar data covering an area of 1.2km × 1.2km. Leonard et al. (2001) applies a modified version of this algorithm to sonar data

---

[1] *"If you cause your ship to stop, and place the head of a long tube in the water, and place the other extremity to your ear, you will hear ships at a great distance from you"*

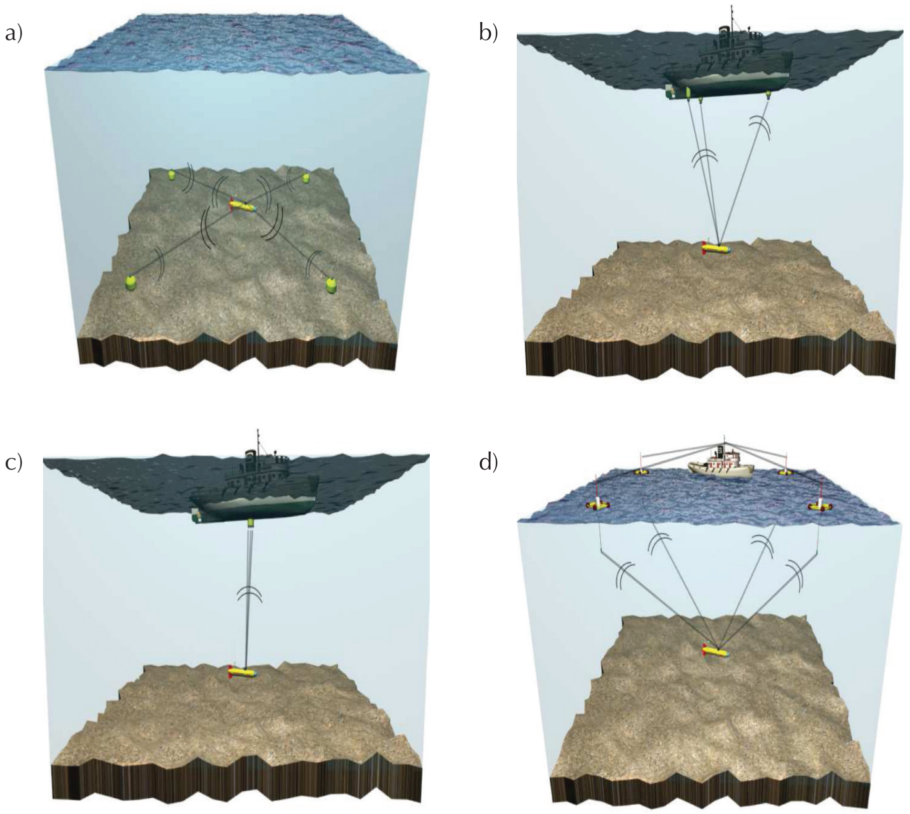

Figure 2.1: Illustrations of acoustic beacon based underwater localisation techniques reproduced from work presented by Alcocer et al. (2006) showing a) Long baseline (LBL), b) Short baseline (SBL), c) Ultra short baseline (USBL) and d) LBL with GPS Intelligent Buoys (GIB).

collected from a test tank experiment and later to forward looking sonar data sets obtained in the ocean by a US Navy vessel equipped with an 87 kHz high resolution array (HRA) sonar. Williams et al. (2001a) presents results of ocean experiments where a simultaneous mapping and localisation algorithm operating on sonar data was deployed on the Oberon AUV with artificial landmarks distributed along a 50 m stretch of Sydney shoreline. Supplementing this scheme, Majumder et al. (2001) presents a framework to fuse vision and sonar data obtained in shallow water environments to perform localisation of an AUV. Williams and Mahon (2004a) further apply and develop this concept in their work to perform simultaneous localisation and mapping in 3D. In addition Newman et al. (2003) presents results of applying simultaneous mapping and localisation schemes for AUVs using a synthetic aperture sonar for terrain sensing and compares its performance to acoustic beacon network based position fixes along with data from an on-board Doppler velocity log (DVL).

2.2.2 Acoustic beacon based localisation

While there is no equivalent to GPS underwater, considering its availability at almost any location on the surface of the planet, underwater acoustic beacons provide a somewhat similar service for position fixing in a much smaller area where the beacons are deployed. Therefore, the use of acoustic beacon networks has been the localisation method of choice for most underwater robotic applications for many years (Bellingham et al., 1994; Deffenbaugh et al., 1996a). In the context of AUVs, on-board omnidirectional transducers are used to interrogate a transponder beacon using an acoustic signal with a predefined frequency signature. Upon receiving the signal, the transponder responds by transmitting an acoustic signal with a different frequency signature after a predetermined delay. The navigation system on-board the AUV measures the round-trip time for the acoustic signal upon receiving the reply from the transponder; thus estimating the distance between the transponder and the vehicle. However, to accurately obtain a position fix, the AUV needs to interrogate multiple transponder beacons. Depending on the distance between transponders (which could be deployed on the sea floor, attached to floating buoys or mounted on surface vessels), they are categorised as long baseline (LBL) short baseline (SBL) or ultra-short baseline (USBL) techniques. Alcocer et al. (2006) and the references therein gives an introduction to the traditional methods of underwater localisation including LBL, SBL and USBL techniques. Some diagrams illustrating these different acoustic localisation schemes are reproduced in figure 2.1.

A large body of literature exists covering many aspects of acoustic beacon based navigation and localisation, among those; Vaganay et al. (1999) and Matos et al. (1999) discusses the use of dead-reckoning in between position fixes using on-board inertial measurements, while Olson et al. (2004) presents an outlier rejection scheme for localisation using LBL methods. Bingham and Seering (2006) discuss the use of hypothesis grids for improving LBL navigation for AUVs and Larsen (2000) proposes a method called 'synthetic LBL' where a single transponder and dead-reckoned vehicle motion simulates multiple transponders.

A significant challenge faced when using a pre-deployed beacon network (LBL) had been the lack of precise position information of the transponders themselves. Once they are dropped to the bottom of the ocean, careful surveying involving multiple surface vessels is required to accurately calibrate the positions of the transponders. However, these positions can change over time due to shifting sediments and other geological and environmental activity, necessitating periodic re-calibration which makes maintaining such a beacon network extremely costly. Additionally, in the case of sea-floor based, or floating buoy based beacons, the theatre of operation for AUVs is limited by the area serviced by the transponder network. In the case of ship/boat mounted technologies such as SBL and USBL, the area coverage is once again limited to the sensing range of these transponders. However, these techniques have been successfully used for underwater localisation, especially in the context of single AUV missions. The work presented by Rigby et al. (2006) demonstrates accurate geo-referenced underwater navigation in 3-dimensions by fusing measurements from an on-board DVL with localisation information obtained via a boat mounted USBL system.

A slightly different approach is used in the work presented by Liu and Milios (2005) which is later used to track the AQUA robot (Dudek et al., 2007). Instead of a conventional USBL system, a surface floating buoy with an array of four hydrophones are used to localise a sound source located on the robot (using hyperbolic localisation techniques). The localised relative positions are mapped to a global reference frame using the position of the floating buoy acquired via GPS, compass, inclinometers, and inertial sensors. The passive localisation sensor 'raft' used initially had later been developed in to a self-propelled buoy which is capable of positioning itself on the surface to track the AQUA robot using acoustic source localisation.

With the emergence of the field of underwater acoustic sensor networks (UWASN), much effort had been focused on localising static or mobile sensor nodes with respect to a number of 'anchor nodes' whose positions are known (Pompili et al., 2008). These techniques draw insights from traditional acoustic beacon based localisation schemes (Chandrasekhar et al., 2006). Cheng et al. (2007) describes such a strategy and presents simulation results for self-localisation of a mobile (AUV) or stationary underwater sensor node with respect to four anchor nodes with known positions using trilateration. Dive and rise (DNR) beacon networks involve mobile beacons which acquire a GPS position fix by periodical surfacing (Erol et al., 2007). Once the position fix is obtained, they dive becoming anchor nodes which transmit their positions. In 'multi-stage localisation' studies presented by Erol et al. (2008), other nodes (static or mobile) can localise with respect to these anchor nodes by additionally using a communication channel. With results of their simulations, it is concluded that the communication overhead is higher for mobile node localisation compared to static nodes.

In the recent years, traditional long baseline navigation concepts have been extended to accommodate multiple AUV missions with autonomous surface crafts equipped with GPS antennae and underwater acoustic modems implementing 'moving long baseline' (MLBL) concepts (Vaganay et al., 2004; Curcio et al., 2005a). These mitigate some of the drawbacks such as limited area coverage and need for survey and re-calibration of pre-deployed beacon networks. In these MLBL applications, the conventional acoustic transponders are replaced with underwater acoustic modems such as the WHOI micro-modem (Freitag et al., 2005).

Although with the advantage of a much faster speed of propagation of acoustic signals in water (approximately $1500\,\mathrm{ms}^{-1}$) compared to air, sonar ranging systems and transponder beacon localisation schemes operate much slower than their in-air counterparts which utilise electromagnetic waves. This characteristic is shared by communication systems as well. Though underwater acoustic modems are now available with usable bandwidth, range and speed as off-the-shelf products, they are far limited in bandwidth and speed in comparison to other common in-air wireless communication solutions. Apart from the speed of operation, most of the communication and navigation applications involving underwater acoustics are prone to adverse effects such as those introduced by sound speed profiles caused by changing water temperature, multipath propagation, frequency dependent fading, scattering and noise as elaborated by Baggeroer et al. (1993), Collins and Kuperman (1994) and Kilfoyle and Baggeroer (2000).

2.2.3 Vision based localisation

Depending on the application domain, the requirements for localisation varies widely. While the most popular modality is to use variants of acoustic positioning and sonar sensing, a few vision based methods have been successfully used as well. Huster (2003) proposes a localisation system based on monocular vision and an inertial measurement unit for underwater object manipulation while Plotnik and Rock (2005) describes a stereo vision based system used by an AUV for tracking marine organisms in the ocean. Sáez et al. (2006) use a 'trinocular' stereo rig with three grey scale cameras to perform underwater simultaneous localisation and mapping in 3D. Corke et al. (2007) presents experimental results of underwater localisation using visual odometry and scaled optical flow obtained from stereo cameras. The localisation performance is also compared with that of an acoustic localisation system using static acoustical sensor nodes. However, the existing literature suggests that the use of vision underwater is limited to short range sensing in non-turbid, shallow and illuminated (or deep and artificially illuminated) environments.

2.3 Relative localisation in multi-robot setups

Coordination and manoeuvring of multiple vehicles presents additional localisation and navigation requirements. Self-localisation by each vehicle with respect to a common global reference frame is no longer sufficient for most such applications. In the previous chapter, it was established that a minimum requirement for a swarm to operate is for each vehicle to at least be aware of the relative locations of vehicles in its immediate neighbourhood. Different approaches used to achieve this were loosely categorised in section 1.1 as a) direct relative position sensing of other members, b) use of communication channels to exchange own self-localised positions (with respect to a common global reference frame) with the neighbourhood and c) use of external centralised observations to explicitly sense the position of each individual and relaying that information back to the navigation systems of each swarm member via communication channels.

The following sections give few of the examples available in the literature which address the problem of relative localisation in multi-robot setups, not limited to underwater environments nor acoustical methods. However, approaches addressing underwater multi-robot localisation are given more emphasis and experimental results showing position estimation performance are also presented where available.

2.3.1 Direct relative position sensing

A few examples of implementations which do not rely on any communication channel to broadcast/exchange self-localised position information while facilitating multi-robot localisation are described in the following sub-sections.

As acoustical navigation system for multi-vehicle operations

A customised LBL method proposed by Atwood et al. (1995) is among the earliest work addressing the problem of multiple AUV navigation. They highlight the need for either the LBL beacons to synchronously emit navigation pings (to implement hyperbolic techniques for localisation) or for the AUVs themselves to schedule their interrogation cycles (to implement spherical techniques for localisation) to avoid confusion caused by overlapping acoustic pings. The latter method is favoured and implemented via a master-slave (leader-follower) approach. The designated master vehicle initiates a cycle by emitting an interrogation ping which is heard by the slave vehicles. Each of them initiate their own interrogation ping after waiting for different preset delays. The master, being aware of the schedule would be able to localise the positions of all slave vehicles, which is identified as an advantage of this approach along with the ability to acoustically monitor the positions of all vehicles if the master interrogation cycle is triggered by a ping from a surface vessel. The main drawback however is the reduction of the 'navigation duty cycle' as the number of slave vehicles are increased. As described by the authors, doubling the number of vehicles more than doubles the time between interrogation cycles.

This technique is demonstrated with two AUVs, one acting as a master and the other as a slave. The ping period of the master was 10 s and the slave emitted an interrogation ping 5 s after hearing the 9 kHz interrogation ping of the master. Apart from the drawbacks associated with conventional LBL beacon networks with regard to deployment and maintenance, the viability of such a system decreases as the number of vehicles increase, due to the long delays between position fixing.

Multi-frequency LBL beacon network for multiple AUV navigation

Cruz et al. (2001) proposes a multiple AUV navigation system with a network of multi-frequency (20-30 kHz) transponders attached to surface buoys with known locations. While each AUV separately interrogates the transponders as in traditional LBL systems, the other AUVs are supposed to listen to these interrogation pings and derive the relative positions of the vehicles attempting localisation. This is facilitated by the different frequency pairs used by each vehicle

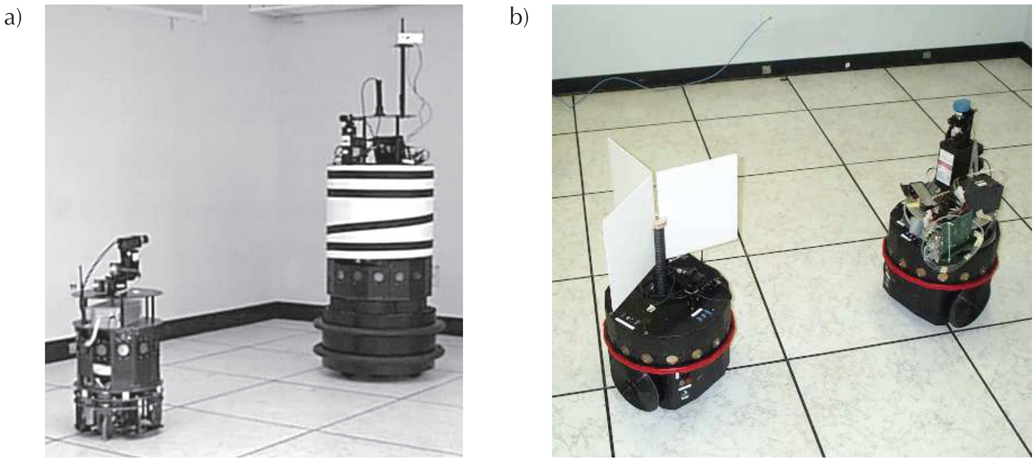

Figure 2.2: Images of the two implementations of the robot tracker sensor, reproduced from work presented by Rekleitis (2003). a) Robot tracker with helical pattern and camera, b) robot tracker with three vertical planes and laser range finder.

to interrogate each transponder. By attaching radio transmitters to these buoys, a land based station can track the position of each of the AUVs while in operation. While the system is purported to facilitate multiple AUV navigation, the presented experiments (in a 8km × 3km effective area) only show remote tracking involving one REMUS class AUV (1.5m in length, 0.2m in diameter). The authors report the externally tracked position was in agreement with the internally logged vehicle position within 5 - 10m.

Robot tracker for indoor multi-robot exploration

The 'robot tracker' sensor described by Rekleitis (2003) is used to localise and track the position of another robot during cooperative localisation in multi-robot exploration. The experiments involving two robots demonstrate how one stationary robot provides a salient landmark for the other robot to navigate, a strategy to compensate for non-salient environmental features interfering with longer term navigation. In the first of the two implementations of the 'robot sensor', a camera on one robot is used to identify a unique helical pattern on the other robot and derive the relative position and orientation from the observed visual pattern. In the second implementation, a laser range finder is used to identify a unique target mounted on the other robot comprising of three vertical planes made of laser reflecting material. The relative position of the target robot is derived using the laser range finder measurements based on reflected intensities. The two robot tracker implementations are depicted in figure 2.2.

Infra-Red relative position sensing for small scale robot formations

Pugh and Martinoli (2006) presents a relative position sensing methodology for small scale indoor robots using an infra-red based system. Here, relative range and bearing of a neighbouring robot is measured using the received signal strength indication (RSSI). This approach has the added advantage of being able to use the infra-red localisation system for low bandwidth communication between the robots. However, relative localisation does not depend on the availability of the communication channel in this approach. Navigation of multiple robots using this scheme is simulated while experiments (with and without communication) involving 4 real robots are also presented. The results suggest that the localisation error is reduced when the communication channels are used to exchange relative position information between localising robots.

2.3.2 Relative position information via communication

Examples of implementations which require some form of position information exchange via communication channels to realise relative position updates to facilitate multi-robot navigation are appraised in the following sub-sections.

Positioning for multiple AUVs using GPS and acoustic communication

This methodology is presented and tested in simulation by Baccou et al. (2001) as a low cost solution for multiple AUV navigation where a flotilla of AUVs consists of a designated leader vehicle and followers. The leader vehicle using dead-reckoning based on inertial data (velocity calculated using the propellor rotation speed), periodically reaches the surface to obtain a GPS position fix. Its displacement with regard to the initial position based on dead-reckoning corrected using the GPS fix is then broadcast to the rest of the followers using acoustic communication modems. Once the information is received, each of the follower vehicles interrogate the leader vehicle which now acts as an acoustic transponder in order to find their distances to the leader. Based on their own dead-reckoned displacements, the leader's displacement and their distances to the leader, the followers update their own positions. This simulation only describes a scenario with one follower, therefore scheduling issues arising from multiple followers trying to interrogate the leader is not addressed. Another aspect of this method is that the leader vehicle is not aware of the positions of the followers. While it is applicable in a single follower scenario, this scheme does not provide a method for followers to be aware of each others positions either, as each localise with respect to the leader only and does not maintain communication between followers.

Localisation and navigation for multiple AUVs using acoustic communication

Freitag et al. (2001) proposes an improved system for multiple vehicles to use LBL transponders efficiently with passive listening and inter-vehicle acoustic communication. This work includes results of tests carried out in a shallow ocean environment using static sensor nodes and a surface vessel mounted mobile node later replaced by a bottom crawling (surf zone crawler) unmanned vehicle. In this scenario, a designated leader vehicle interrogates an LBL transponder network and each of the other vehicles passively listens to the responses by the transponders. This is possible as the other vehicles are equipped with acoustic modems (WHOI utility acoustic modem) which utilises the same frequency band as used by the LBL transponders. The frequency shift keyed (FSK) signals emitted from the transponders are used to calculate the distances to each of the vehicles. The acoustic modems are then used to communicate the position of each of the vehicles to a surface vehicle for monitoring and coordination. The bidirectional modem communication makes it possible for an external controller to provide mission commands to the individual vehicles. Two LBL transponders with a baseline distance of 1500 m, at a depth of 18 m was used during the experiments. The mobile node operated at a depth of 12 - 18 m and the bottom crawling vehicle reached depths as shallow as 1.5 m. The results report an accuracy within 6 m of the corresponding ground truth obtained via GPS fixes (which had an error of the same magnitude). The ranges used during the experiments were up to 2000 m.

Stojanovic et al. (2002) draws upon the work presented above and presents a concurrent mapping and localisation scheme for multiple AUV operation which is based on inter-vehicle distances. This is achieved by measuring inter-vehicle delays using matched filtering which are later refined using Doppler frequency shifts when communicating with each other using acoustic modems (WHOI micro-modems). These modems utilise high rate phase shift keyed (PSK) or quadrature amplitude modulated (QAM) acoustic signals. The authors also presents a slot based communication protocol including an initialization phase such that the network can be built up dynamically. The communication process exchanges local position maps containing the positions of other vehicles in its neighbourhood as measured by the individual vehicle. This enables each vehicle to be aware of the positions of other vehicles. Since the localisation scheme suggested here is tightly coupled with the communication system whose speed and bandwidth is limited, the position update rate can be adversely affected as the number of vehicles increase. Furthermore, the refinement of distance estimation using Doppler frequency shifts is only viable when the vehicles are in motion.

A two hydrophone heading sensor for multiple AUV navigation

A two hydrophone heading sensor presented by Reeder et al. (2004) and a leader-follower navigation algorithm presented by Edwards et al. (2004) is implemented in work presented by Baker et al. (2005b) which facilitates simultaneous navigation of multiple AUVs. The presented scheme requires one designated vehicle (leader) to perform conventional LBL localisation while the others (followers) are equipped with the 'two hydrophone heading sensor' which intercepts acoustic pings emitted by the leader and each of the transponders to derive their relative heading. These are then fused with other information (known geometry of the transponders, inertial heading of vehicle) to ascertain the position of the follower. The presented simulations with five followers assume the leader vehicle position is broadcast to the followers via a parallel acoustic communication channel. While it is purported to be a leader-follower scheme, since the presented experiments only involve one follower, the problems arising due to multiple followers not knowing each others positions are not addressed. In this scheme, as in the previously described approach, the leader vehicle is not aware of the positions of the followers.

The presented experiments involved two fixed transponders at a depth of 12 m with a base distance of 146 m. In the first instance the two hydrophone heading sensor, representing the follower vehicle was mounted below a surface vessel while the acoustic source (projector driven using a WHOI micro-modem) representing the leader vehicle was suspended from a tethered stationary moorage. In the second instance the acoustic source representing the leader vehicle was also mounted below a surface craft which was driven on a straight course at a velocity of $1.05\,\mathrm{ms}^{-1}$ by a human driver. In the earlier instance the follower surface craft was driven past the stationary leader at velocities of 1.0 - $1.8\,\mathrm{ms}^{-1}$. In each case the hydrophones and projectors were at a depth of 2 m and the distances varied between 20 - 40 m.

The authors report an acoustic source level of 183 dB (re 1 µPa @ 1 yd) using an ITC-1032 omnidirectional transducer (resonance at 32 kHz) and the signal comprised of a binary phase shift keyed (BPSK) navigation ping with a carrier frequency of 26 kHz and bandwidth of 4 kHz with a signal duration of 7 ms. Receiving hydrophones were two omnidirectional ITC-8140 transducers (flat frequency response for 1 - 40 kHz) separated by a base distance of 0.457 m. The received signals were sampled at 65 536 Hz with a resolution of 16 bits.

During these experiments the angles were measured using cross-correlation as well as matched filtering and the results were compared. Sub-sample interpolation had been used to enhance the resolution of the bearing estimates. The maximum heading errors are reported as 4° and 9° with cross-correlation providing the greater error. The authors state that the reason for this discrepancy is unknown at the time of publication. Furthermore, it is reported that the first 100 s of the experiment produced only 2 valid bearings out of 7 attempted and the next 100 s produced 11 valid bearings out of 15 attempted. Overall, a precision of 10 m within an area of 500 m × 500 m is reported.

Cooperative localisation for AUVs using moving baseline navigation

Bahr and Leonard (2008) presents a method for multiple AUVs to perform cooperative localisation using WHOI micro-modems for underwater communication. The individual vehicles are meant to localise with respect to a few CNAs (Communication and Navigation Aid AUVs) equipped with sensors to obtain accurate self-localisation information. All clocks on the multiple AUVs and CNAs are assumed to be synchronised via the acoustic modems and externally

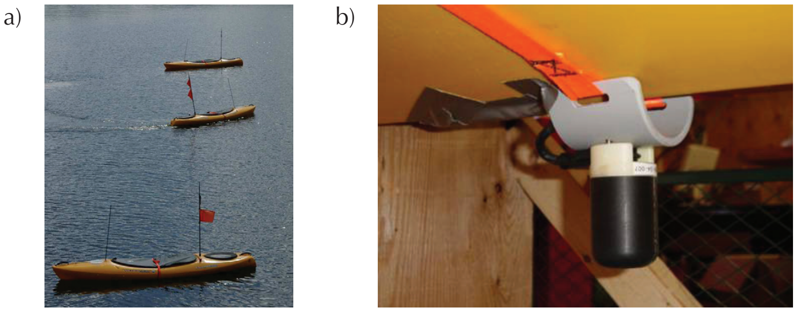

Figure 2.3: Cooperative localisation experiments using autonomous surface crafts (ASCs) with a) an image of three ASCs reproduced from work presented by Bahr and Leonard (2008) and b) an image of a WHOI micro-modem mounted to an ASC reproduced from work presented by Curcio et al. (2005b).

supplied (from GPS receivers) pulse per second (PPS) signals. The presented navigation principle is similar to the moving baseline concept described and demonstrated by Curcio et al. (2005b) where some drawbacks of conventional LBL schemes are mitigated by the use of a mobile transponder network. With each signal transmission by a CNA which includes its absolute position (longitude, latitude and depth), the listening nodes can estimate the distance between themselves and the CNA using one-way range calculations due to synchrony of clocks.

Experiments presented in this work use three autonomous kayaks - ASCs (Autonomous surface crafts) described by Curcio et al., (2005a) in place of AUVs, each equipped with a WHOI micro-modem (mounted to the bottom of the kayaks). Two are designated as CNAs with access to GPS positioning while the third operates as an AUV. The 'AUV kayak' navigates according to a pre-programmed mission using GPS way points while tracking its own position and those of the two CNAs which moved in formation to stay within range of the acoustic modems (figure 2.3). The tracked position is compared to ground truth obtained via the GPS position fixes. Since these experiments only involved one 'AUV kayak', the problem of multiple AUVs trying to localise each other had not been addressed. It can be assumed that some form of communication schedule is to be adopted in such a case where each AUV would make acoustic broadcasts of its depth and position in addition to the CNAs to allow all members of a swarm to localise each other. The authors report a maximum position update rate of 0.1Hz with regard to the localisation algorithm used during the experiments.

Experiments presented by Curcio et al. (2005b) involving three ASCs use round-trip range measurements as well as one-way range measurements using the WHOI micro-modems for localisation and navigation. Among the many experiments conducted with surface crafts emulating AUVs, a maximum range of 400 m is reported while a nominal separation of 25 m to 100 m was maintained between vehicles. Cooperative localisation was performed with inter-vehicle communication using leader-follower and formation keeping configurations during these experiments. The authors report a nominal position error of approximately 1% compared to

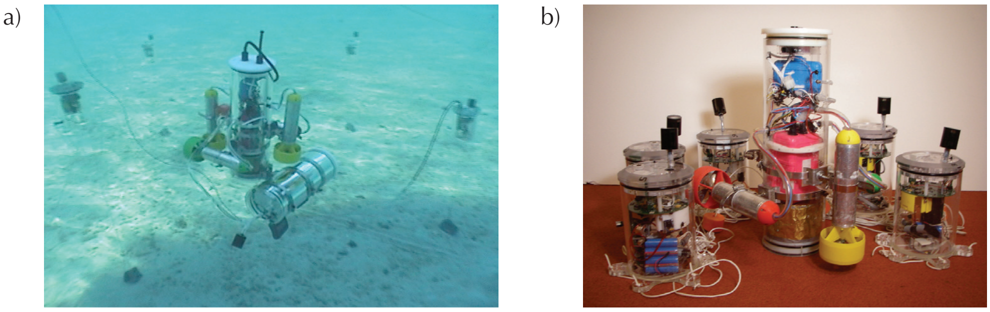

Figure 2.4: Images of the AMOUR AUV and acoustic sensor nodes a) during an underwater localisation experiment reproduced from work presented by Corke et al. (2007) and b) reproduced from work presented by Detweiler et al. (2007).

GPS ground truth while nominal range errors are around 2 m according to the presented plots. While an experiment where one ASC was replaced with an Odyssey III AUV is mentioned, the localisation performance was not presented.

Navigation of multiple AUVs using synchronous clock one way travel time

The work presented by Eustice et al. (2007) uses the same synchronised clock concept described in the previous approach to estimate ranges between nodes using WHOI micro-modems. On board navigational data is broadcast by each vehicle and all receiving vehicles estimate distances between themselves and the broadcasting vehicle based on one-way travel times facilitated by synchronised clocks.

During the experiments presented in this work the node on the surface uses highly accurate GPS based clocks to maintain synchrony while the submerged nodes are equipped with temperature compensated crystal oscillator based clock sources with a drift rate of approximately 1 ms per 14 h. This translates to a maximum drift induced range error of 1.5 m per dive which the authors claim to be similar to errors due to a conventional LBL navigation system. As with the previous cases, while the system is purported to support multiple AUV navigation, the experiments only demonstrate the localisation of one SeaBED AUV with respect to a ship (both equipped with acoustic modems) and does not address the issue of multiple AUVs localising each other. The experiments cover an area of 200 m × 200 m and nominal position errors range from 2 m to 5 m according to the presented plots when compared to ground truth obtained via an LBL system.

Static sensor node networks for AUV localisation

The work presented by Corke et al. (2007) describes the use of a static underwater acoustic sensor node network which can be utilised for localisation and navigation of multiple AUVs. Unlike traditional LBL beacon networks, these nodes are equipped with acoustic modems capable of bidirectional communication. These nodes can localise each other using three methods:

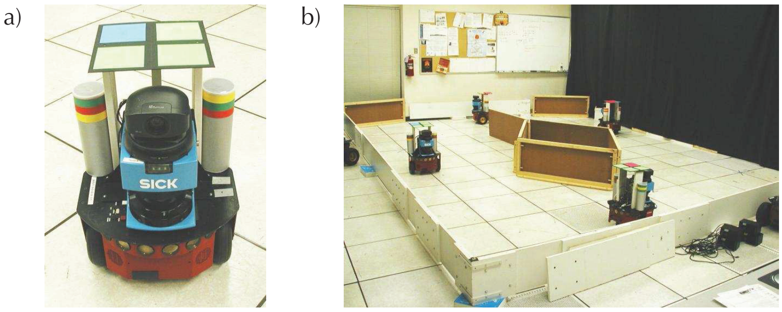

Figure 2.5: Images reproduced from work presented by Howard et al. (2003) showing a) One of the robots used in the localisation experiments equipped with a scanning laser range finder, a pan-tilt-zoom camera, a pair of retro-reflective and colour-coded fiducials, and a second colour-coded fiducial for use with the overhead tracking/ground-truth system and b) the experimental environment.

a) inter-node distance measurement with round-trip time delay when two nodes exchange messages, b) a node broadcasting a range request to which other nodes respond after specific delays allowing inter-node distance measurement with round-trip time delay and c) use of on-board synchronized clocks for nodes to ping at specified intervals allowing listening nodes to compute the distances based on differences in time of arrival of the acoustic signals. The sensor node network performs self calibration (Vasilescu et al., 2007) using a distributed localisation algorithm based on work by Moore et al. (2004) which allows the sensor nodes to be thrown overboard eliminating the need for the precise deployment and survey as in the case of conventional LBL transponders.

The authors claim these sensor networks can be used to localise multiple mobile nodes (AUVs), each equipped with similar acoustic modems, to concurrently perform localisation. However, the presented experiments involve only one mobile node (AMOUR AUV) and four static sensor nodes in one instance and one mobile node (Starbug AUV) and six static sensor nodes in the other (figure 2.4). The mobile node obtained range measurements to some of the static nodes every two seconds which results in a position update rate is 0.5 Hz. The acoustic localisation performance is compared to GPS position fixes (the AUV navigated near the surface of the water allowing it to log GPS positions) and a nominal location error of approximately 2.5 m is reported.

In-air multi-robot localisation with inter-node communication

There are many examples in the literature where relative position sensing coupled with communication channels between robots have been used to perform in-air cooperative multi-robot localisation and navigation. A few of these relative localisation methods are briefly reviewed here.

Fox et al. (2000) introduces a Markov localisation technique for probabilistic multi-robot localisation in known environments (a map is available *a priori*) where position estimates are exchanged between robots. In the physical experiments presented, two Pioneer mobile robots were used in an indoor environment. The 'robot sensor' used to perform relative localisation consisted of a laser range finder and a colour camera mounted on each robot. Each robot was marked with a unique colour marker which can be recognized by the vision system of the robots. Once a robot is recognized, its relative position (angle and distance) was measured using the laser range finder.

In work presented by Howard et al. (2003) Bayesian formalisms are used to perform cooperative relative localisation in multi-robot teams. The Pioneer robots used in the experiments are each equipped with a laser range finder and a camera which constitutes a 'robot sensor' to detect the position of a nearby robot (figure 2.5). Each robot has a unique colour coded fiducial for recognition by the vision system. As in the previous case, once another robot is detected (and recognized - using the unique fiducials, unlike the previous case) using the vision system, the laser range finder measurements are used to obtain the relative position of that robot. At each instance of localising another robot, that information is wirelessly broadcast using the user datagram protocol (UDP) such that the localisation information is shared amongst other robots.

In their work, Mourikis and Roumeliotis (2006) presents a study of various cooperative localisation algorithms used in multi-robot setups. They analyse the dependence of localisation performance of robot teams on factors such as the size of the team and accuracy of the robots' sensors. For this analysis, they perform experiments in a rectangular arena with four Pioneer mobile robots each equipped with a laser range finder for self-localisation. Each robot is mounted with a visual marker which is used to track the position of the robots using an overhead camera based vision system. Since the robots do not have exteroceptive sensors capable of sensing the relative locations of other robots directly, these relative position measurements are synthesized using the individual robot positions tracked via the vision system. The self-localisation information is assumed to be exchanged between robots using a communication channel (the processing was performed offline). The results show that the localisation accuracy of the robot team improves when position information is exchanged between robots while performing relative localisation.

2.3.3 Position sensing by external centralised observer

A few examples of implementations where an external "God's eye view" perspective was used to achieve relative localisation to facilitate multi-robot navigation are presented in the following sub-sections.

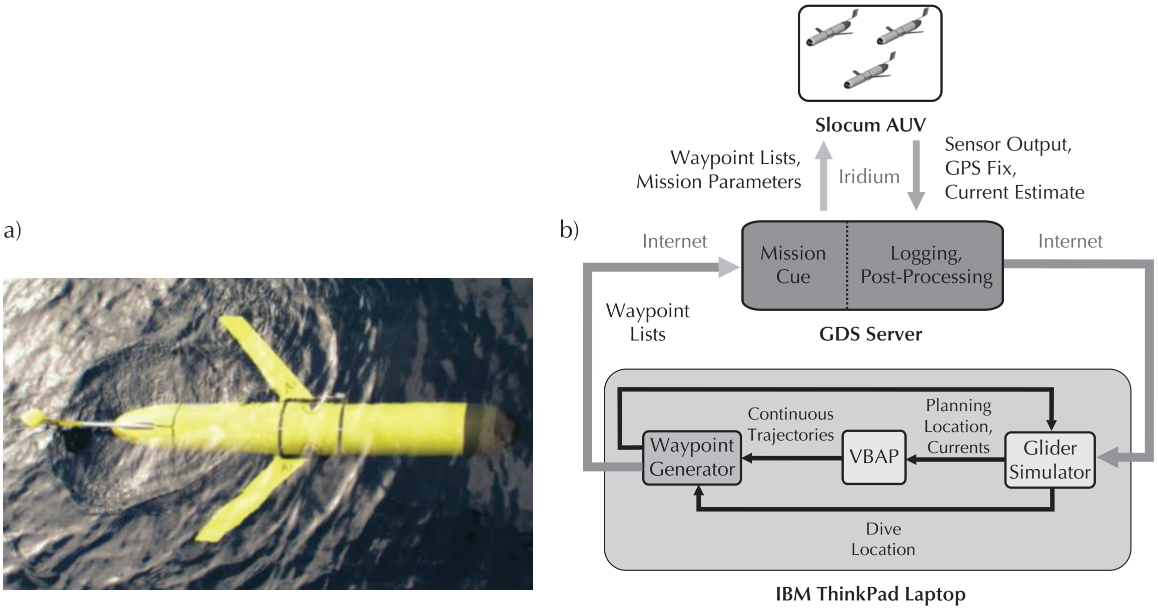

Figure 2.6: The image and diagram reproduced from work presented by Fiorelli et al. (2006) shows a) a Slocum underwater glider and b) the operational configuration and data flow of the system used to coordinate the fleet of gliders.

Coordination and control of an underwater glider fleet

Underwater gliders (Bachmayer et al., 2004) have drawn much attention in recent years for long duration wide area coverage missions due to their superior endurance compared to other AUVs. Many sea trials have been performed with these including participation in the autonomous oceanographic sampling network (AOSN) experiments (Bellingham and Zhang, 2005; Ramp et al., 2008).

Over a number of years Naomi Leonard and colleagues have developed strategies and methods for coordination and navigation of a fleet of underwater gliders for adaptive oceanographic sampling (Paley et al., 2008; Davis et al., 2008). The central concept of 'virtual bodies and artificial potentials' (VBAP) used for navigation of these gliders is explained by Ögren et al. (2004). In order to maintain a stable formation of the reference points (gliders) constituting the virtual body and to perform adaptive sampling based on artificial potential gradients, each vehicle needs to know the position of at least its near neighbours (Leonard et al., 2007, p.52). However, the Slocum underwater gliders do not have a facility for inter-vehicle communication, especially in the context of the large inter-vehicle distances maintained during typical glider fleet missions. The sea trial experiments presented by Fiorelli et al. (2006) involve inter-vehicle distances of 6 km and 3 km for three Slocum gliders attempting to maintain a formation at the vertices of an equilateral triangle. Given these large inter-vehicle distances and the slow effective speed of the gliders ($0.35\,\mathrm{ms}^{-1}$), a relatively slow position update can be tolerated by the navigation system. During operation, each vehicle reach the ocean surface periodically (every two hours) to obtain a GPS position fix and asynchronously transmit its position to an on-land base station via satellite phone links. The navigation coordination system located on the base station transmit waypoints back to the gliders via the same satellite phone links (figure 2.6). It must be noted that while the vehicles were operating in the underwater medium, the self-localisation (using GPS) and communication (Iridium satellite phone links) were all conducted in air. Relative positions between vehicles are measured using an external centralised system and relayed back to the vehicles. In addition, unlike other examples, this implementation does not use acoustical methods for localisation or communication.

Experiments with cooperative aerobots to simulate underwater swarms

In a novel approach, Honary et al. (2009) uses aerobots (automated blimps) to simulate a swarm of underwater robots in their experiments to test cooperative navigation algorithms. For the formation flying and cooperative area coverage missions, each member of the swarm needed position information of other members in the neighbourhood. Since the aerobots lacked an appropriate exteroceptive sensor capable of performing relative localisation, this capability was simulated using external sensors. During the experiments conducted inside a large auditorium, each of the blimps (three were used) were affixed with custom reflective markers which were tracked using a network of 12 infrared cameras. The positions of individual vehicles were then broadcast via wireless local area network. Each vehicle then derived relative positions of other swarm members using this information.

Air traffic control, take off and landing approaches of aircraft

While it is arguable if modern aircraft can be classified as ‘robots’, the role played by human air traffic controllers at airports is an example of localisation and navigation using a “God’s eye view” perspective. Even though most aircraft are equipped with sensors to detect other aircraft in the vicinity, these are meant to be used for collision avoidance rather than relative localisation. During the ‘cruising’ phase of navigation, aircraft perform self-localisation with respect to a suite of sensors ranging from precise inertial navigation systems, GPS and radio navigation beacons. However, when approaching or leaving airports where a large number of aircraft converge, localisation is usually taken over by ground based air traffic controllers. Powerful radar, radio beacons and radio communication channels are used to localise and track individual members among the ‘swarm’ of aircraft in the vicinity of most major airports. Navigation waypoints, landing approach and take off patterns are relayed back to the aircraft from the ground based control centres based on this localisation information.

2.3.4 Comparison of performance

The earlier sections gave examples of relative localisation strategies implemented in multi-robot setups with an emphasis on those operating in the underwater environment. All presented underwater implementations except one, utilise underwater acoustics for either localisation, communication or both. Despite the fact that the discussed multi-robot implementations all performed either explicit or implicit forms of position sensing of other members in the 'swarm', the objectives were varied. In some cases it was motivated by maintaining rigid formations, to perform cooperative localisation aiming to minimise navigation errors or to facilitate adaptive spatio-temporal sampling.

Description	Number of nodes (Type)	Used comms?	Internode distance or survey area	Position errors (Reference)
Navigation of multiple AUVs using synchronous clock one way travel time (Eustice et al., 2007)	2 (1 AUV, 1 ship)	yes	200 m × 200 m	~2 m - 5 m (LBL)
Static sensor node networks for AUV localisation (Corke et al., 2007)	5 (4 static, 1 AUV)	yes	80 m × 80 m	~2.5 m (GPS)
Coordination and control of an underwater glider fleet (Fiorelli et al., 2006)	3 (Gliders)	yes *	3000 m - 6000 m	255 m - 623 m (GPS)
Cooperative localisation for AUVs using moving baseline navigation (Curcio et al., 2005b; Bahr and Leonard, 2008)	3 (Autonomous Kayaks)	yes	25 m - 100 m	~1% of range (GPS)
A two hydrophone heading sensor for multiple AUV navigation (Baker et al., 2005b)	2 (1 static, 1 boat)	no**	500 m × 500 m	~10 m (LBL)
Localisation and navigation for multiple AUVs using acoustic communication (Freitag et al., 2001; Stojanovic et al., 2002)	3 (2 static, 1 boat / surf zone crawler)	yes	up to 2000 m	~6 m (GPS)
Multi-frequency LBL beacon network for multiple AUV navigation (Cruz et al., 2001)	3 (2 static, 1 AUV)	no	4000 m × 8000 m	~5 m - 10 m (LBL/ Inertial)

Table 2.1: Summery of experimental setups and results extracted from the literature where the problem of simultaneous navigation of multiple AUVs was addressed. The position errors are either nominal or average errors explicitly reported by the authors or derived from the provided plots.
* Communication with external base station via satellite phone links once at the water surface.
** The proposed strategy requires inter-vehicle communication but not used during experiment.

Given the various experimental platforms, spatial scales of operation and performance metrics used by the different research groups it is extremely difficult to qualitatively compare the performance of different localisation methods and strategies against each other. Summarised experimental results extracted from the literature (where available) associated with seven of the 'real-world' underwater multi-robot implementations discussed earlier are tabulated in table 2.1 in reverse chronological order. While each method attempt to address the problem of simultaneous navigation of multiple AUVs, it must be noted that most of these strategies are highly specialised and specifically adapted to the experimental platforms and the application. Therefore, attempting to derive an overall performance ranking is not realistic.

2.4 Sensor utilisation strategies for localisation

Configuration of localisation sensors as well as processing techniques have been extensively studied in fields of target motion analysis (active and passive detection/tracking), acoustic source localisation (ASL) and wireless sensor network research. Additionally, it also has overlap in the field of wireless communication with regard to receiver/transmitter configuration and processing. The following sections give brief overviews of techniques which have relevance to the problem of relative localisation in multi-robot setups.

Multiple Input Multiple Output (MIMO) methods

In most cases, any system which involves multiple receivers and multiple transmitters can be classified as a MIMO system while in some cases the specific signal processing strategy involving multiple inputs and outputs is a pre-requisite for a system to be identified as a MIMO system. In wireless communication systems, use of multiple receiver and transmitter antennae had led to remarkable improvements in overcoming problems caused by multipath propagation, interference and behaviour of time-varying channel characteristics such as fading. Instead of treating it as a problem, multipath propagation is in fact exploited by MIMO techniques to improve reliability and throughput of the communication channels without any further increase of channel bandwidth or transmission power (Gesbert et al., 2000; Goldsmith et al., 2003). These aspects such as channel reliability and capacity which attract a lot of attention in wireless communication research is of limited relevance for localisation. However, researchers in the field of acoustic source localisation have adopted some MIMO strategies to implement target tracking in noisy reverberant environments (Huang et al., 2006; Fallon and Godsill, 2008).

Time delay estimation is essential in a number of signal processing techniques used for source localisation where two or more input channels are involved. The methods used by many researchers can be classified as adaptations of the generalised cross-correlation (GCC) approach introduced by Knapp and Carter (1976). However, these methods have limitations when applied in the presence of multipath arrivals in highly reverberant environments. Influenced by MIMO techniques, the adaptive eigenvalue decomposition (AED) approach proposed by Benesty (2000) makes use of the additional information presented by multipath arrivals to perform passive acoustic source localisation. These techniques have been used in applications such as acoustic echo cancellation, time delay estimation, cross-talk cancellation and speech de-reverberation as described by the extensive survey presented by Huang et al. (2006).

Another area which benefits from MIMO techniques is when multiple source signals need to be separated from a mixture without explicit knowledge about the source signals, which is referred to as blind source separation (BSS). This is used in areas such as multiple speaker tracking and speech recovery using at least as many microphones as the number of sources to be separated. Buchner et al. (2005) presents an approach for simultaneous estimation of multiple time difference of arrivals (TDOAs) based on blind adaptive MIMO filtering using a microphone array to track multiple speakers in a reverberant environment. Lombard et al. (2008) presents experimental evaluation of a BSS-based real-time demonstrator for the localisation of two sound sources using MIMO processing. The authors report that the system is capable of accurately localising two speech sources in two dimensional space within a few seconds and with a precision better than one degree. The experimental implementation which used four microphones (two pairs) did not rely on any prior knowledge of the source positions.

Beamforming methods using sensor arrays

Beamforming with receiver arrays (usually three or more sensors) is used in sonar detection, source localisation and target tracking to improve performance. Time domain beamforming is achieved by setting delays and gain factors for each of the array elements appropriately such that a larger more sensitive sensor can be simulated using multiple smaller less sensitive sensors. In frequency domain beamforming, the received signals are separated in to different frequencies using a fast Fourier transform (FFT) either across time or across different array elements and gain factors are set to each separated frequency[1]. Additionally, by dynamically changing the delays and gain factors, the 'beam' can be steered in an arbitrary direction. Adaptive beamforming strategies are used in localisation where the beam is dynamically steered to point to the signal source to maximise the signal to noise ratio. These techniques are also used to locate and track multiple simultaneous signal sources. Chen et al. (2003) demonstrates such methods in acoustic source localisation. Valin et al. (2004) presents an experimental evaluation of a mobile robot mounted frequency-domain steered beamformer approach capable of localising up to two simultaneous moving sound sources using an array of eight microphones.

---

[1] *See Van Veen and Buckley (1988) for an in-depth description of beamforming techniques.*

Applicability of MIMO and beamforming techniques

As mentioned before, advantages presented by MIMO processing techniques such as improvement in channel reliability and capacity are of interest for communication applications but not necessarily for localisation. However performance gains are demonstrated in acoustic source localisation applications with regard to multiple source tracking and speech separation in reverberant environments when MIMO techniques are applied. The approach used in this thesis to address the localisation problem stated in chapter 1 differs from these applications as explained below.

The use of broadband source signal pings with known statistical properties and relatively short durations can mitigate detrimental effects caused by delayed multipath arrivals when operating in reverberant environments. For localisation purposes, the direct path arrival of the source signal ping is sufficient and this can be recovered using time-domain channel windowing coupled with cross-correlation and matched filtering techniques. However, since the duration of the signal can be considered continuous in speaker tracking and speech separation applications, multipath arrivals can cause echoes and interference. Channel windowing techniques can still be used for source tracking but not in the context of speech recovery as the source signal needs to be recovered in its entirety.

Both in MIMO processing and beamforming techniques, for accurate simultaneous localisation of multiple signal sources a relatively large number of receivers need to be used. The approach used in this thesis minimises the chance of multiple sources from emitting pings simultaneously[1] within the range of a receiver. This is achieved by a) temporally separating the multiple pings emitted by a single vehicle, b) exploiting time division multiple access (TDMA) scheduling within a local neighbourhood and c) implicitly synchronised signal transmission across the swarm[2]. Under these circumstances, the utility of a technique for simultaneous localisation of multiple sources becomes redundant. Additionally, the approach presented in this thesis aims to perform localisation with minimum possible hardware, space, power and processing requirements. In this context, while MIMO and beamforming techniques might well achieve the same localisation performance, the additional processing and sensors required for the effective implementation of these methods place a strain on limited resources available on the small Serafina class AUVs.

In summary, the strong points of MIMO and beamforming techniques are a) being able to accurately localise and extract continuous signal sources in reverberant environments and b) being able to effectively separate multiple simultaneous signal sources. In order to realise the full

---

[1] *Handling of such occurrences is addressed in section 5.3 of chapter 5.*

[2] *The relationship between the localisation system and the scheduling system is discussed in section 5.2 of chapter 5. See work by Schill (2007) for an in depth explanation of the communication scheduling system developed for Serafina class AUV swarms.*

potential of these techniques, it is desirable to have spatially distributed arrays with many sensor elements. Since a) and b) stated above are not high priority requirements of the approach used in this thesis to address the localisation problem, the additional hardware, space, power and processing requirements imposed by MIMO and beamforming techniques are difficult to justify.

2.5 Discussion

Localisation is an integral part of mobile robot navigation regardless of the operating environment. Due to the non-availability and limitations of sensor modalities, robots operating in the underwater environment face a number of obstacles as discussed by Leonard et al. (1998) and Loebis et al. (2002). Unlike in operation of a single AUV, the multi-robot paradigm presents many additional challenges for localisation and navigation along with the bounty of new applications it makes possible. Smith et al. (1998) identifies navigation, synchronisation techniques and logistics as key problems in realising the full potential of multiple AUV missions for synoptic and pseudosynoptic data collection.

LBL navigation

It was established in the previous chapter that each member of a multi-robot setup having access to position information of at least their near neighbours is a minimal requirement for the successful operation of a swarm or a formation. In principle, it is possible for multiple AUVs to be deployed with pre-programmed navigation waypoints to conduct 'formation flying' with only self-localisation information obtained via LBL beacon networks without relative position awareness or inter-vehicle communication. However, this approach precludes many of the synergetic benefits championed by the swarm robotics research community and cannot be used for applications such as adaptive oceanographic sampling (Martins et al., 2003; Bhatta et al., 2005). The strategy used in many in-air applications had been to establish communication links between the vehicles to exchange absolute position information which is then used to derive relative position of other members of the swarm. In the face of limitations in speed and bandwidth of underwater communication channels, this strategy has had limited success when applied to multiple AUV navigation. The other problem of LBL navigation is that the operational area is limited to a pre-instrumented segment of the ocean. While USBL navigation provides independence of pre-deployed beacon networks, this method imposes harsher limits on maximum transponder-vehicle distances (Smith and Kronen, 1997). In both cases (LBL and USBL navigation) beacons need to be individually interrogated by each vehicle to update its own position causing the position update rate to decreases as the number of vehicles increase. As Eustice et al. (2007) points out, this limits LBL and USBL navigation to multi-robot groups of only few members.

Underwater acoustic modems with ranging capability

In the backdrop of the circumstances mentioned above, the WHOI micro-modem for underwater acoustic communication and navigation and its availability as an 'off the shelf' package was considered a significant innovation and the oceanographic community had been quick to adopt this technology. Despite its relatively low data rate of 80 bps, the use of frequency hopping FSK (Frequency Shift Keying) modulation technology provide reliable communication links up to about 4 km even in shallow water environments with multipath propagation. In addition to providing a communication link, the micro-modem has the capability of supporting LBL type navigation by acting as a transponder and up to four vehicles can share the acoustic navigation beacons, by using different broadband interrogation codes. The modem network protocol also supports up to 15 nodes to facilitate underwater acoustic networks (Freitag et al., 2005). Another feature which had been exploited for localisation is the ability to do ranging between modem nodes using round-trip travel time or one-way travel time when externally provided pulse per second (PPS) reference clock signals are available. Apart from the examples described in earlier sections (Curcio et al., 2005b; Eustice et al., 2007; Bahr and Leonard, 2008) a number of additional underwater acoustic navigation applications demonstrating the versatility of the WHOI micro-modem are given by Singh et al. (2006). With the availability of inter-vehicle communication channels and flexible and versatile transponder beacons which can easily be mounted on AUVs or ASCs (autonomous surface crafts) such as the MIT SCOUT, the concept of 'moving long baseline' (MLBL) navigation for multiple AUVs have been proposed (Vaganay et al., 2004; Curcio et al., 2005a).

Building on the success of the WHOI micro-modem, novel underwater acoustic sensor nodes with modem and ranging capability have been developed and presented by Vasilescu et al. (2007). This implementation also uses FSK modulation and reports a data rate of 300 bps verified up to 300 m in freshwater and ocean environments. These nodes also support a time division multiple access (TDMA) scheduling protocol which can be used to perform self-synchronisation and self-calibration of the network. Given their small size and versatility, as proposed in work presented by Corke et al. (2007), these sensor node networks can be used to facilitate localisation for multi-vehicle AUV missions.

Moving away from the 'self-localise and communicate' strategy for relative localisation, the ranging capability of acoustic modems and the possibility to mount these modems on AUVs makes other cooperative localisation approaches possible. These methods based on sensor node localisation research presented by Moore et al. (2004) attempts to directly measure distances to others vehicles without relying on beacons or anchor nodes. Proposed multi-AUV localisation strategies reviewed previously (Bahr and Leonard, 2008; Corke et al., 2007) are hybrid approaches which use relative range measurements, self-localisation with respect to beacons (static sensor nodes or mobile CNAs) and communication channels to propagate position information among other members in the group. A fully decentralised 'real-world' implementation of an AUV swarm (more than two AUVs instead of leader-follower schemes) is yet to be deployed at the time of writing as the references in the current literature all point to either software simulations or hardware simulations using maximally two AUVs and additional surface crafts.

Localisation for small agile AUVs in dynamic unstructured environments

To facilitate the swarming paradigm, the rate at which relative position information needs to be updated depends on the application and the robotic platforms being used. For larger inter-vehicle distances and slower vehicle speeds the position update rate can afford to be relatively low[1] while on the other hand, a swarm consisting of agile vehicles operating in highly dynamic and unstructured environments with relatively short inter-vehicle distances would require a much higher position update rate for successful operation. The Serafina class AUVs (Serafina website, 2009) are highly agile[2] with their small size and five thruster actuation compared to both traditional AUVs of 'torpedo style' and 'crate style'[3]designs. They are less than half the length of the Starbug AUVs which are similarly actuated (Dunbabin et al., 2004). This unique design, the level of agility and small size will allow the Serafina AUVs to be deployed in large dense swarms with relatively short inter-vehicle distances (up to 20 m) operating in dynamic and unstructured environments. This is an unprecedented prospect in underwater robotics and opens up many application possibilities in scales which were previously infeasible.

Since fast and reliable communication links are essential for the operation of a swarm of Serafina AUVs as described above, a specialised short range inter-vehicle communication system using a 122 880 Hz long-wave radio carrier signal with a maximum data rate of 8 192 bps have been developed and tested (Schill, 2007, pp. 66-69). The projected range of this system is up to 30 m (at full drive voltage) while the current state of experimentation verifies ranges of over 15 m. Due to the small size of the AUVs, available space inside the hull is extremely limited. This precludes possibilities of using existing acoustic modem technology. Moreover, the data rate achievable by the long-wave radio communication system is far superior than those provided by state of the art underwater acoustic modems. As further elaborated by Schill (2007), the communication system is reinforced with fully decentralised spatially distributed TDMA routing schedules which allow dynamic re-configuration, addition and deletion of communication nodes. The two 'flavours' of scheduling regimes enforced by the **distributed dynamical omnicast routing** (DDOR)

---

[1] *E.g. for underwater glider fleets with large inter-vehicle distances of 3 km to 6 km and low effective speeds 0.35 ms^{-1}, a relatively slow position update rate of once every 2 hours can be tolerated (Fiorelli et al., 2006).*

[2] *Serafina class AUVs are 50 cm in length, 10 cm in diameter with a maximum forward speed of 1.5 ms^{-1} and roll, pitch and yaw rates of 360° s^{-1}, 180° s^{-1} and 90° s^{-1} respectively.*

[3] *REMUS class AUVs (Allen et al., 2000) are examples of 'torpedo style' design while Kambara and Oberon AUVs are examples of 'crate style' designs (Wettergreen et al., 1999; Williams et al., 2001a).*

algorithm (Schill and Zimmer, 2006b) and the **pruned distributed omnicast routing** (PDOR) algorithm (Schill and Zimmer, 2007) are both fully scalable and places no upper bounds to the number of participating swarm members. By facilitating simultaneous broadcasts within spatially distributed local neighbourhoods, these routing schedules implement ‘every node to every node’ (omnicast) communication within the swarm in a fast and efficient manner within one scheduling run.

Despite the availability of a fast and reliable communication system, the problem of relative localisation was not yet addressed for the realisation of a swarm of Serafina AUVs. With shorter inter-vehicle distances and small fast moving members, the accuracy and update rate requirements placed on the localisation system is high. Self-localisation methods such as individually interrogating transponder beacons was not possible due to limitations imposed on update rates by swarm size when using such methodologies as elucidated earlier. Inter-vehicle ranging methods such as those used by applications using state of the art underwater acoustic modems offer certain possibilities in this regard. Nevertheless, in the absence of externally supplied precise reference clock signals and synchronised clocks on all swarm members, these methods have to rely on round-trip travel time for range measurement. This feature coupled with the space constraints within the Serafina AUVs mentioned earlier makes this solution undesirable for the problem at hand[1].

The relative localisation system

‘Robot sensors’ or ‘localisation sensors’ such as those mentioned in the literature (Fox et al., 2000; Rekleitis, 2003; Howard et al., 2003) with regard to land based multi-robot setups which provides relative estimates for range (distance), azimuth (bearing) and heading (rotation) of another robot based on observation and sensing does not have equivalents in underwater robotics. The nearest related work in this regard is the two hydrophone heading sensor proposed and tested by Baker et al. (2005b) but fails to deliver a fully decentralised and accurate position estimation methodology.

The research presented in this thesis addresses the problem of designing, developing and evaluating a fully decentralised relative localisation system comprising of a ‘localisation sensor’ capable of producing relative estimates for range, azimuth and heading of other nodes aimed at facilitating swarming of small and agile Serafina class AUVs. Drawing insights from strengths and weaknesses of underwater localisation methods referred earlier, the proposed strategy exploits the available communication scheduling system developed by Schill (2007). While they communicate, the vehicles synchronously sends out short acoustic pings with long-wave radio broadcasts

---

[1] *In the future, maturity of technology and miniaturization coupled with availability of precise clocks, this method might be a feasible alternative.*

within a local neighbourhood. In addition to measuring the azimuth of the acoustic source using hyperbolic techniques, the rest of the members of the neighbourhood can each measure ranges to the sender using the arrival time difference between the electromagnetic and acoustic signals using matched filtering techniques. By sequentially sending two acoustic pings from the two ends of each vehicle, it also enables the observers to measure the relative heading of the sender. Additionally, the sensor geometry used in the implementation allows for additional 'reverse hyperbolic' techniques to be used to derive alternate heading and range estimates which do not rely on the synchrony provided by the communication and scheduling system. This adds redundancy to the position estimation thus improving reliability.

The availability of such a low level 'localisation sensor' makes it possible to explore the abundant swarm robotics literature to find many candidate high level cooperative navigation and localisation strategies that have been developed and studied over the years which assumes the availability of such sensors. In addition, using the communication system to exchange the relative localisation system leads to each swarm member being aware of positions of all other swarm members (at least within the local neighbourhood). Not only does this facilitate a platform for robust swarming applications, it also improves overall localisation accuracy as evident from studies presented in available literature (Mourikis and Roumeliotis, 2006; Pugh and Martinoli, 2006). Moreover, by using the communication channel to broadcast depth measurements derived from on-board pressure sensors with each acoustic sending event, the observing submersibles can all calculate relative depth information based on their own pressure sensor reading and incorporate this information with the 2-dimensional position given by the acoustic localisation system to estimate the 3-dimensional position of the sending submersible. This approach has been extensively used to derive 3-dimensional position information by reducing the underwater localisation problem to 2-dimensions (Bellingham et al., 1992; Vaganay et al, 2000; Baccou et al., 2001; Cheng et al., 2008; Bahr and Leonard, 2008).

The next few chapters give in-depth descriptions of the enabling technologies and methodology used to develop the relative localisation system. This is followed by experimental evaluations and performance analyses of the setup to discern its suitability to address the stated problem of facilitating swarming of small agile autonomous submersibles with an accurate, precise and robust localisation system that scales up with increasing swarm size and independent of pre-deployed beacon networks.