Chapter 1

Introduction

Among the definitions for the term **navigation**, many still refer to its early roots in maritime transportation such as the following found in the Oxford English dictionary: “The art or science of directing the movements of a vessel on the sea or other open water; the process of determining and planning a vessel's position and course, by means of geometry, nautical astronomy, instruments”. A more concise interpretation is given by Kaplan and Hegarty (2006) where: “Navigation is defined as the science of getting a craft or person from one place to another”. A common feature among these is the reference to **localisation**, or position fixing which is the process of determining the location of a craft, person or platform[1] with respect to some point of reference. Hence localisation is an integral part of navigation, may it be land, underwater or aerial domains.

The process can be loosely classified as **self-localisation** with respect to one or many fixed beacons (landmarks), *e.g.* trying to find ones own location in an environment using a map based on correlating observed landmarks with those on the map; or **relative localisation** (**mapping**) when locating the position of one or many beacons with unknown positions with respect to your own position. These concepts are not exclusive and can have overlaps depending on the application

---

[1] *The term ‘platform’ is used to refer to a robot or any vehicle performing the task of navigation.*

concerned and sometimes can be simply referred to as **localisation** without the use of prefixes; these are used extensively in robotics especially in the context of navigation and target tracking. The field of simultaneous localisation and mapping (SLAM) research (also known as concurrent mapping and localisation (CML)) use both these concepts extensively. SLAM addresses the process of navigating in unknown environments where the features (landmarks) of the environment are localised based on sensory observations and incorporated in to a map while using the same map as a navigational aid for self-localisation (Thrun, 2002). In the case of known environments, the navigation task becomes somewhat simpler as *a priori* information about the surroundings would be available. This could either be in the form of active landmarks such as navigation beacons (used in aircraft and ship navigation) or passive features in the environment which corresponds to an already available map.

Another localisation scenario is when a communication channel is available between the navigating vehicle and an external observer. The observing platform / station would be actively or passively localising the navigating vehicle (control tower radar observing an aircraft, sonar station observing a submarine) and relaying that information to the navigation/path planning system of the vehicle. In these cases, the localisation task is external to the vehicle and receives the information via the communication channels.

The environments in which autonomous robots operate can be given three independent attributes of varying degree depending on a) the level of *a priori* knowledge about the environment, b) the 'structuredness' or the regularity of environmental features and c) the level of temporal persistence of the aforementioned features. As a guide, navigation in known, structured and static environments is considered relatively easy with currently available methodologies while unknown, unstructured and dynamic environments pose many unaddressed challenges.

1.1 Motivation

Drawing inspiration from nature, the use of multiple robots to solve a task in a collaborative manner, either as swarms or formations, has been growing as a vibrant sub-field in mobile robotics research for over two decades[1]. Some of the attractive features of this concept are, the ability to perform spatially distributed tasks quicker, reliably and more robustly, easier reconfigurability of a system and complex emergent behaviour of a swarm comprising of relatively simple individuals. A large body of accumulating literature had been addressing different aspects of swarming from the higher level swarm behaviours to the lower level sensing strategies. One common thread that emerges throughout the literature is the importance of each vehicle being aware of relative

---

[1] *See Liu (2008), pp. 9- 50 for a survey of swarm robotics research.*

locations of others[1] (at least of the immediate neighbourhood) for the success of the swarming paradigm.

Relative localisation in multi-robot systems

In his work Reynolds (1987) published a set of simple 'swarming rules' aimed at achieving 'life-like' animation of boids[2]. The underlying distributed behavioural model states that individuals' knowledge of the relative positions of neighbouring members is sufficient for sustaining and controlling swarm behaviour. In her seminal work Matarić (1994) postulates that "the ability to distinguish the agents with whom one is interacting from everything else in the environment is a necessary condition for intelligent interaction and group behaviour"[3].

In swarm robotics research, the problem of acquiring information regarding the relative position of other members have been addressed in three broad approaches:

1) Use of explicit on-board exteroceptive sensing to directly detect relative positions of other members (*e.g.* using cameras, acoustic transponders *etc.*).
2) Use of communication channels to exchange / broadcast own localised positions (in reference to a common global frame) with / to the neighbourhood (*e.g.* using radio communication to distribute GPS fixes or self-localised map coordinates obtained via laser range finders, using acoustic communication to distribute self-localised position with respect to an acoustic beacon network *etc.*).
3) Use of external "God's eye view" observations to explicitly sense the position of each individual and relaying that information back to the swarm via communication channels (*e.g.* using over head cameras, acoustic transponders to observer motion and detect positions of individual members *etc.*).

This is not a strict exclusive categorisation but rather a classification based on prominent attributes observed among different implementations for the ease of analysis. Except for the first approach, the other two heavily relies on the existence and reliability of communication channels. In addition, the third approach precludes the notion of decentralised control of the swarm. However in practise, these three approaches are often inter-mixed to varying degrees depending on the application, the availability of communication channels and the level of autonomy assigned to each member of the swarm.

---

[1] *In addition to the knowledge about the operating environment (Fenwick et al., 2002; Howard et al., 2003; Pugh and Martinoli, 2006)*

[2] *Simulated bird-like "bird-oid" objects generically called "boids" even when representing other creatures such as schooling fish.*

[3] *In her experiments, members of the 'Nerd Herd' recognised other robots using short range radio broadcasts and their positions were detected using infra red sensors - in the related simulations, precise position information such as distance and direction of nearest neighbours were available to each member.*

With the maturity of the field of mobile robotics and the ubiquity of position sensing and communication methodology (*e.g.* miniature portable GPS receivers, laser range finders, advanced vision processing systems, small low-cost RF transceivers) relative localisation in land based applications has been successfully addressed utilising multitudes of these techniques. Consequently, a substantial portion of the literature assumes an underlying 'localisation sensor' and 'communication channel' when addressing the higher level swarm behaviour/control and navigation problems. For example, Martinelli et al. (2005) assumes the existence of exteroceptive sensors capable of measuring relative bearing, distance and orientation of other robots in their multi-robot localisation work. In addition, the cooperative navigation and localisation work presented by Roumeliotis and Bekey (2000b) as well as Moore et al. (2004) assumes the existence of a fast and reliable communication channel between all participants. While the assumption of accessible localisation sensors and communication channels holds true for most in-air applications, it is not necessarily the case in the underwater environment.

Underwater swarms

As the field of underwater robotics progressed with the advancement of mechanical and electronic technology, the amount of autonomy embodied in the vehicles themselves have considerably increased while making it possible to build and deploy even smaller autonomous underwater vehicles (AUVs). However, the cost of most ocean going scientific AUVs has remained relatively high. The oceanographic research community as well as scientists of many other fields (geological, chemical, biological *etc.*) who need to map and survey large bodies of water rely on AUVs for data collection. Apart from the financial aspects, a loss of an autonomous robot during a mission[1] usually means loss of mission data as well. Under these circumstances, the concept of multiple cooperative underwater robots started to emerge combining expertise from the areas of underwater robotics and distributed sensor networks (Curtin et al., 1993). This was not meant necessarily as a replacement for larger AUVs, but rather to offer a viable alternative for missions which would naturally yield themselves to this approach (Bellingham and Rajan, 2007, p. 1100). Multiple smaller autonomous robots collaborating and sharing information improve reliability and enhance robustness against loss of mission data in the event of an individual vehicle being damaged or lost. Apart from making many surveying tasks much more efficient in terms of area coverage and power requirements, the paradigm of a small school of AUVs instead of a single main robot makes possible certain tasks which were not possible earlier, such as dynamically and simultaneously obtaining spatio-temporal measurements of a body of water at multiple locations, cooperative searching, plume tracking and gradient following (Stojanovic et al., 2002; Kalantar and Zimmer, 2007; Leonard et al., 2007; Ramp et al., 2008).

---

[1] *Kaiko (JAMSTEC) lost in 2003 (Maki et al., 2005), Autosub-1 (NERC) lost in 2005 (Manley, 2007).*

Figure 1.1: Serafina Mk I AUVs [top left], prototype Serafina Mk II AUVs [bottom left] and the author and Jan Zimmer setting up localisation experiments using mock-up hulls [right].

[PHOTOGRAPHY BY UWE R. ZIMMER]

While many researchers are actively engaged in developing cooperative multiple AUV technology, the development of the **Serafina AUV** project (figure 1.1) at The Australian National University (ANU) was initiated with the goal of developing swarming/schooling technology for small submersible robots (Kalantar, 2006; Schill, 2007; Serafina website, 2009).

Relative localisation for underwater swarms

Apart from the many challenges encountered by roboticists when it comes to designing and developing underwater vehicles, the swarming paradigm introduces yet another. As discussed earlier, for multiple vehicles to collaborate, *i.e.* "behave as a swarm", it is essential that each member is aware of the positions of at least its near neighbours. In the existing literature, the problem of underwater localisation in the context of multi-robot systems has been addressed in a number of ways combining the three approaches mentioned earlier in various degrees.

One method is to have a centralised controller which has a "God's eye view" over the swarm using some form of sensing modality to track the positions of each individual vehicle. This task is usually designated to a surface vehicle acting as a base station. Depending on the level of autonomy allocated to the members of the swarm, the controller issues either control commands, way-points or position information via communication channels maintained with each vehicle[1].

---

[1] *E.g. a fleet of gliders surface every 2hrs to obtain a GPS fix and transmits information via satellite phone links to a base station which performs necessary processing and relays back way-points via the same communication channel (Leonard et al., 2007).*

Reliable and fast communication channels underwater are by themselves a challenging problem (Proakis et al., 2001; Lucani et al., 2008). As the number of members in a swarm increase, so does the required bandwidth of the communication. Furthermore, it is preferred to have a decentralised swarm which otherwise impedes some of the attractiveness of this paradigm.

Another approach is for a designated leader vehicle to regularly fix its position with respect to a landmark (acoustic beacon network) and then broadcast this information to the rest of the swarm in the first phase. In the second phase, the members of the swarm individually interrogates the leader vehicle to measure the range between them and the leader. This two phase process leads to each member of the swarm being aware of its position. Once the members are aware of their locations, it is communicated to the rest of the vehicles. This method relies on a designated leader and improves performance over a method where each member of the swarm individually localise with respect to an acoustic beacon network and communicates its position fix to its neighbours. Usually the beacons used can only serve one vehicle at a time where they first need to be interrogated by the vehicle attempting to fix its position. This aspect makes the process of vehicle positions propagating throughout the swarm considerably slow as the number of vehicles increase.

In a variant of the above approach, which does not depend on deployed acoustic beacon networks, the leader vehicle is meant to surface occasionally to get a GPS position fix. The follower vehicles individually interrogate the leader to find its position and uses dead-reckoning in the mean time to update their positions. An underlying feature of these methods apart from the reliance of explicit communication between vehicles for navigation, is that most of the approaches described in the literature are built on a leader-follower paradigm rather than a fully-fledged homogeneous swarm. As a result, the emphasis has been to update the vehicles with their own position and not on each member of the swarm getting regular position updates of vehicles in its local neighbourhood.

It must be emphasised that the existence of communication channels between members is considered essential for the swarming paradigm in light of sharing mission data. However, the speed and bandwidth requirements for such communication tends to be lower compared to those required for multi-robot navigation, especially considering the above mentioned approaches. With the limitations faced by underwater communication techniques at present, it is preferable to make navigation less dependant on inter-vehicle communication, especially in the context of AUV swarms, whose manoeuvrability would otherwise be limited by the speed of the communications network.

Under these circumstances, with the growing research interest in the field of underwater swarm robotics, shortcomings in transferring concepts and methodologies from land based multi-robot systems are emerging. This is especially the case in relative localisation for underwater swarms given the limitations and unavailability of appropriate 'localisation sensors' and communication channels in the underwater medium. With the advent of swarms comprising of small agile vehicles, the conventional schemes such as beacon network based localisation and reliance on slow acoustic communication channels for localisation are becoming inadequate.

The Serafina AUV project was initiated with the goal of developing swarming technology for the small, agile and highly manoeuvrable Serafina class AUVs[1]. Smith et al. (1998) elaborates the attractiveness of using multiple AUVs for synoptic and pseudosynoptic oceanographic data collection while highlighting operational logistics as a key challenge. Using smaller AUVs contribute to greatly reducing the time and cost involved in transport, deployment and recovery while allowing more AUVs to be used in missions with a similar or lower cost and effort compared to the use of larger conventional vehicles.

The use of multiple homogenous vehicles with similar sensing and navigation capabilities further assists in managing operational logistics and yields itself to distributed and decentralised control of the swarm. Additionally, use of such homogenous vehicles can contributes to an overall reduction of navigation errors as pointed out by Roumeliotis and Bekey (2002) in their distributed multi-robot localisation work. The distributed and decentralised control paradigm provides safeguards against mission failures due to loss of or damage to a few individual members of a swarm. By considering these aspects, a decentralised concept of localisation was preferred over strategies requiring a designated leader or centralised controllers.

As pointed out by Leonard et al. (2007), the number of vehicles and their speeds needs to be matched with the spatial and temporal scales of interest when monitoring and sampling time varying spatially distributed fields using swarms of AUVs. This emphasises the need for fast and agile vehicles for applications in environments whose dynamic features vary in short time scales.

With the relatively small size, and scenarios which require a large number of the AUVs to swarm, the localisation system would need to efficiently scale with the number of vehicles involved. In order to make maximum use of the agility of these vehicles and provide the swarm with fast and dynamic manoeuvres, the position update rates needed to be much faster, accurate and precise than those provided by state of the art strategies. In addition, a system which does not depend on pre-deployed acoustic beacons was preferred as well. This eliminates the need for preparing the area of operation and calibrating the beacons and also lifts any constraints placed upon the swarm by a beacon network covering only a limited underwater area. Moreover, the small size of the individual vehicles make it impossible to incorporate existing technology such as underwater

---

[1] *The Serafina class vehicles are just over 50cm in length, 10cm in diameter, actuated by five thrusters giving it five direct degrees of freedom (roll, pitch, yaw, heave and thrust) with sway motion achieved by a combination of roll and heave. The maximum forward speed is* $1.5ms^{-1}$.

acoustic modems, ultra-short baseline transponders and sonar modules due to physical size limitations and power consumption requirements.

Problem statement

The research presented in this thesis addresses the challenging problem of designing, developing and evaluating a decentralised relative localisation system capable of facilitating swarming of autonomous submersibles - a system that scales up with increasing swarm size and does not depend on pre-deployed beacon networks or the speed and bandwidth of the communication channels. The design specifications and constraints are drawn from the requirements of providing relative localisation capability for the small and agile Serafina class AUVs.

1.2 Contributions

In order to design a relative localisation system while considering the aspects discussed earlier, the emphasis was shifted more towards direct position sensing as opposed to explicit and active communication being used to exchange self-localised position information between neighbouring vehicles. This line of thought was further reinforced by observing swarms, flocks and schools of animals in nature. For example, a school of fish would change its swimming direction either by all members reacting after sensing the same stimuli or by individuals reacting after sensing a position change of their local neighbourhood - without explicit communication. Furthermore, the swarming behaviour of the AUVs only required positions of vehicles in the local neighbourhood *relative to the observing vehicle* and not relative to a global frame since there is no global controller. Despite a number of challenges being posed by the requirements and specification of a decentralised swarm consisting of small and fast AUVs, the benefits provided are the relatively shorter distances between vehicles and being able to exploit the existing underlying communication and scheduling scheme developed by Schill (2007).

The relative localisation system

In order to address the requirements of the **Serafina AUV** project, a relative localisation system was designed, developed and evaluated experimentally. This system draws insights from hyperbolic and spherical positioning schemes (Deffenbaugh et al., 1996a) and provides each vehicle with a regularly updated pose vector (consisting of the azimuth, range and heading) of its near neighbours with respect to its own frame of reference. As opposed to most relative localisation schemes described in the literature which treats swarm members as point objects and hence provide no heading direction (apart from integrating multiple position updates), this implementation explicitly estimates the heading direction and is included in the pose vector with each update. The implementation utilises an acoustically transmitted **Maximum Length Sequence**

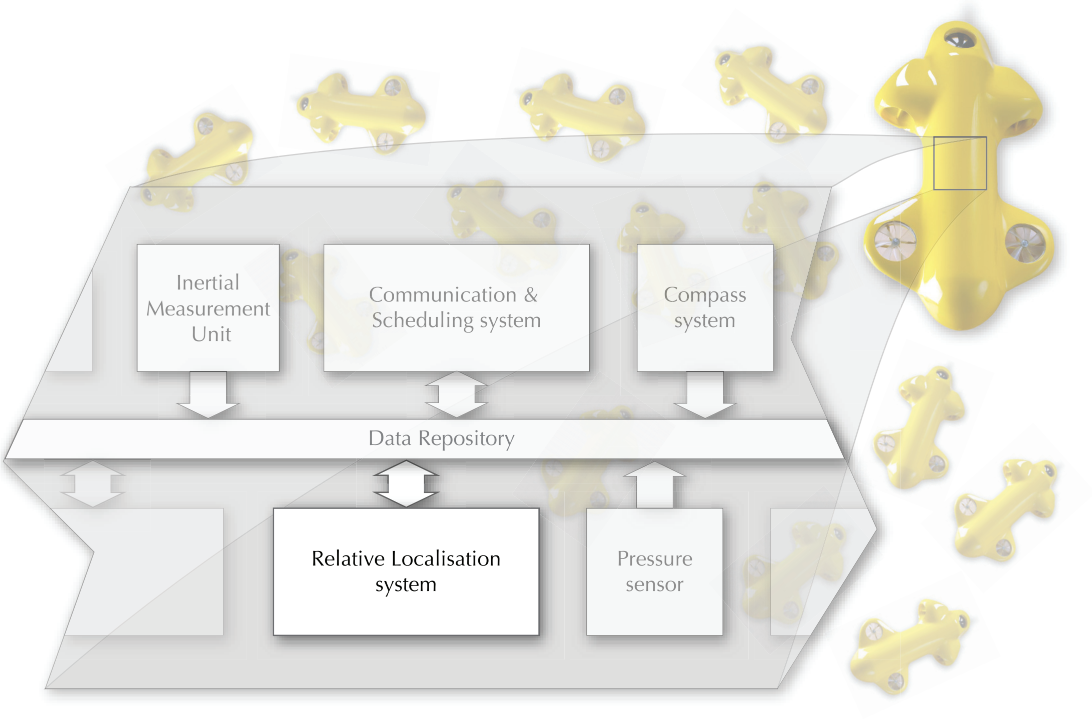

Figure 1.2: Simplified block diagram showing some of the main components that facilitates decentralised swarming capability on a Serafina AUV along with their data flows. This thesis focuses on the relative localisation system.

(MLS) signal from **projectors** on the bow and stern ends of each 'sender' vehicle which in turn is received by a pair of **hydrophones** on the 'observer' vehicles at each update cycle. The statistical properties of the MLS signals provide extremely high robustness against interference by multipath arrivals, cross-talk and noise sources and other inherent detrimental effects in the underwater environment as well as the non-linear characteristics of the transducers used.

While the azimuth is obtained via **hyperbolic positioning** techniques measuring multiple **time-difference-of-arrivals** (TDOA) between the hydrophones, the resolution and accuracy is greatly improved over those achieved by conventional methods using effective outlier handling schemes. The range and heading estimation is performed utilising two independent methods which provide higher reliability and robustness in a fast changing and dynamic environment. One method uses the implicit synchronisation provided by the underlying scheduling system to measure the difference of **time-of-arrivals** (TOA) of the acoustic and electromagnetic (long-wave radio) signals. This gives an equivalent measure to the **time-of-flight** (TOF) of the acoustic signal which is converted to a distance as done in spherical positioning schemes. The second method which only relies on multiple TDOAs, estimates the range without any explicit knowledge of the sending times of the acoustic signals. The **reverse hyperbolic** estimation scheme used here provides a safeguard against erroneous range and heading estimation due to loss or drift of synchronisation between the neighbouring vehicles.

Figure 1.2 shows a simplified block diagram of some of the main components that facilitates decentralised swarming capability on a Serafina AUV. In his thesis Schill (2007) addresses the problem of establishing effective distributed communication in underwater robotic swarms while the work presented in this thesis addresses the problem of achieving localised relative position sensing amongst swarm members. While the implementation characteristics focus on the Serafina AUVs, the relative localisation strategy and innovative methodologies developed in this research can in general be utilised to implement 'localisation sensors' for many other underwater applications (underwater sensor network localisation, tracking of underwater life forms *etc.*) which are not limited to small AUV swarms. The experimental results and analyses presented therein contributes to the growing field of localisation in the context of understanding limitations and opportunities presented by underwater environments. Aspects such as the choice of signal waveform and techniques of handling interference and outliers have potential applications beyond underwater robotics and can be transferred to other application domains with minimal modifications to serve localisation requirements.

Solution synopsis

The work presented in this thesis proposes a novel distributed relative localisation strategy to be used in underwater multi-robot setups with an emphasis on providing swarming capability to small agile AUVs. This strategy is implemented with a relative localisation system comprising of a 'localisation sensor' capable of producing estimates for azimuth, range and heading of neighbouring submersibles. The system is experimentally evaluated and its performance is analysed with regard to aspects such as update rate, sensing range, accuracy and precision of the produced estimates. According to the obtained experimental results, within the required sensing range, the system outperforms state of the art techniques with regard to the speed of updates and the localised position accuracy.

1.3 Thesis outline

The structure of the thesis is outlined in the following sections. A brief introduction is given to each of the chapters which contain the background, methodology developed during the research, experimental results and analyses, extensions and conclusions.

Chapter 2 - Background and related work

This chapter gives a brief insight in to the background of localisation technology, the drivers behind the choice of sensing strategies and methodologies and focuses on work related to the research presented in this thesis. Underwater localisation methods and modalities are discussed first and different strategies used to perform relative localisation in multi-robot setups are reviewed next with an emphasis on 'real-world' implementations addressing the problem of simultaneous navigation of multiple AUVs. Finally, drawbacks and benefits of existing underwater localisation systems with respect to the constraints and requirements of the motivating application is discussed.

Chapter 3 - Source signals

Description of the time-domain cross-correlation used by the relative localisation system as well as the motivation behind the choice of maximum length sequences (MLS) as the source signal is given in this chapter. This includes performance evaluation of several classes of signals with regard to cross-correlation peak detection and signal to noise ratios (SNR). Also presented in this chapter is an empirical method for overcoming the frequency filtering introduced by the transducers to improve the cross-correlation peak detection performance.

Chapter 4 - Acoustic source localisation

The specific distance and angle measurements and estimations carried out during the process of localisation is explained in this chapter. The methodology and basic measurement schemes are described in detail along with identification of different classes of errors affecting the estimated quantities. An analysis is presented on how the uncertainties associated with the basic measurements propagate towards uncertainties in the estimated quantities which is followed by an explanation of how sub-sample interpolation contributes towards increased precision of the estimation system. Descriptions of how the sub-azimuths and sub-ranges are combined to derive the compound estimates for azimuth, range and heading are given along with theoretical error models associated with each of these quantities.

Chapter 5 - The relative localisation system

An overview of the functional components of the relative localisation system is presented in this chapter and goes on to explains how the relative localisation system can provide a distributed localisation solution for swarming of AUVs. The relationship between the relative localisation system and the underlying communication and scheduling system is elaborated while expanding on how multiple senders and observers are accommodated in the context of a local neighbourhood belonging to a larger swarm of AUVs. This chapter also discusses the interference caused by delayed multipath arrivals, cross-talk due to multiple senders and effects of environmental noise on the system along with strategies to address them. Furthermore, outlier handling schemes are introduced and the performance of the proposed **peak tracking** scheme is experimentally evaluated. In addition, the computational complexity of the system is discussed along with a proposed **range tracking** scheme to overcome the problem of increasing computational cost as the effective sensing range increases.

Chapter 6 - Experiments

The configuration, apparatus and procedure used for the experimental evaluation of the relative localisation system is elaborated in this chapter. The experiments were aimed at gauging the accuracy and precision of the estimates under operational conditions, the angular and radial sensing limits of the system and the overall suitability of the approach to solve the task of relative localisation for small AUVs. Most of the experiments were carried out at the ANU test tank[1] while other experiments were carried out at Lake Burley Griffin[2]. This chapter also explains how ground truth references were established to compare the localisation estimates produced by the system.

Chapter 7 - Results and analysis

The effects of the inverse frequency filtering scheme and the peak tracking scheme have on the estimated quantities are explained and analysed in this chapter, while the results of selected short, medium and long range experiments are presented in detail. The performance of the localisation system is analysed in terms of accuracy, precision while the angular and radial sensing ranges of the system are evaluated with respect to SNR of the received hydrophone channels and position errors resulting from pose vector estimates. Experimental data is also used to demonstrate how the system recovers from degradation of position estimation accuracy.

Chapter 8 - Towards 3D source localisation

The localisation system discussed throughout this thesis focuses on 2-dimensional or planar localisation, where the localised position is expressed in polar coordinates with an azimuth and a range. The source to be localised is assumed to be on the same plane containing the two receivers and their main axes of directivity. When the source leaves this plane, the range estimation remains valid while the estimated angle is no longer contained within the plane for which the azimuth was defined. This chapter proposes several strategies to incorporate additional information about the source position with the estimated range and azimuth quantities to produce the true azimuth and elevation angles. The experimental results are also presented to validate the feasibility of the presented relative localisation system in handling 3-dimensional localisation, either with relative depth information or with additional sensors without further modification of the sensing and processing methodologies.

---

[1] *Cylindrical tank with corrugated metal walls filled with tap water. Diameter 4.2m, depth 1.5m.*

[2] *Lake Burley Griffin has an approximate surface area of 6.64km^2 situated in the centre of Canberra, ACT, Australia.*

Chapter 9 - Conclusions

The performance of the relative localisation system is critically compared against the state of the art in terms of accuracy and precision of the obtained localisation estimates. This chapter concludes the thesis summarising key contributions and drawing insights upon the research conducted. Additional work needed to implement a deployable system are enumerated and future research directions are indicated which could benefit from the outcomes of this thesis.