Local Area Positioning of Multiple Moving Objects

A thesis presented in partial fulfilment of the requirements for
the degree of Doctor of Philosophy at
Massey University, Palmerston North,
New Zealand

Timothy Tristan James Brooks
2012

"It doesn't prove anything much except that the awesome splendour of the universe is much easier to deal with if you think of it as a series of small chunks."

"Mort" Terry Pratchett, 1987

Abstract

This research examines a number of aspects of position tracking in complex environments. It postulates that configuring the receivers of the system in a regular fashion does not give optimum results when obstructions are present (especially moving obstructions). The tracking of rugby players is used as a primary example of position tracking in a complex environment. Other applications are considered.

The current literature in the field of position tracking and synchronisation was examined to find the state of the art. This suggests that the method best suited to the research goals is a time-based positioning method using unsynchronised transmitters attached to the players coupled with time-synchronised receivers.

Specifications were generated that defined the minimum operating requirements for a position tracking system that is suitable for use in a sporting application. These specifications are used to create and assess a minimalist tracking system.

A suitable system is characterised and the results are used to derive a model predicting the visibility of a player to the tracking system. For given player positions, and system components, the model determines the visibility of the players. It is designed to be fully configurable to handle any position tracking system and operating environment. The model is used as part of an expert system for performance assessment.

Configuration of the receivers is an important optimisation parameter for the tracking system's performance. A genetic-algorithm optimisation process was tested with several objective functions to find the optimal placement of receivers in both open-field and static-obstruction situations. Calculated optimum solutions are shown to be superior to solutions-by-inspection. Furthermore, the optimisations confirm the premise regarding regular receiver configurations.

Test results show that some loss of visibility is inevitable in a minimalist system and that missing data can be recovered by software, thus increasing system performance while maintaining minimal equipment. A reconstruction algorithm was developed to handle missing data. The known behaviour of objects (specifically rugby players) was characterised into a series of rules to reconstruct missing data from data nearby in both time and space.

Finally, a simulation was performed using game data provided by the Industrial Partner and bringing together the various, disparate threads of research. The system as a whole achieved 93-99.5% player visibility with the use of optimum receiver placements and application of the reconstruction algorithm.

It is concluded that eight is the optimum number of receivers to satisfy the specifications of a minimalist system and that the expert system was successful.

The knowledge created by this research can be applied to any tracking system in order to maximise its efficiency in a given environment. It also demonstrates that the required volume of equipment can be reduced through the use of software tools.

Acknowledgements

Many individuals have helped me during the course of this research. The help was invaluable, and without it I would not have been able to reach this point. It would take far too long to name each of these people individually but I would like to thank the following people and organisations.

- To my supervisors: Huub, Wyatt and Ken. Your supervision and guidance throughout the entire project allowed me to go from a naïve and ignorant graduate to a confident expert in this field. You always pushed me to 'crank the handle' one more time and search for the deeper meaning in the results

- To George and the staff of Verusco Technologies, Palmerston North, New Zealand for teaching me all I now know about the game of rugby and giving me the driving ambition to improve the analysis methods. Also for giving up their time to help with the testing. There is nothing quite as satisfying as ordering a bunch of people to walk around an open field waving sticks above their heads.

- My parents and family for helping with the editing process and interesting discussions on grammar and meaning of the English language. This has been a learning experience for all of us and it goes to show that the process of learning never ends.

- To Red and Woodsy, for their lessons on construction techniques and innovation. Being able to design and build on a shoestring budget allowed this project to be completed.

- To my lovely wife Desiree for pushing me during the final editing.

- And finally, to Heather for her moral support and friendship throughout this project. You have embodied the true definition of a friend. I couldn't have done the testing without you.

Table of Contents

Table of Figures

List of tables

1 Introduction

In which we learn of the context, aim, and research contributions – publications – nomenclature for the thesis – and introduce the Industrial Partner.

The ability to track, with a high degree of accuracy, the positions of multiple moving objects within a given area is very useful, and in some cases necessary to many disciplines. Examples of this would include: tracking the locations of the elderly or the infirm as they move about their home or a care facility, tracking the positions of robots moving about a warehouse and tracking the movements of soldiers on a battle field. There are several choices available when it comes to selecting a position-tracking system for a particular application. The selection will depend on the nature and constraints of the application.

This thesis aims to answer several questions concerning position tracking.

The first question is: *'how best can a minimalist tracking system be implemented?'* A minimalist system is defined as one that uses the least equipment to achieve the desired results. The desired result is further defined as a system that can perform to the specifications of visibility and timeliness (discussed further in the specifications chapter) with sufficient robustness. This question will require summary and evaluation of the different methods of position tracking and the current technologies available including assessment of any deficiencies. Rugby has been identified as a viable commercial application of tracking technology and as such will be used as the primary example in this research and hence, the specification of 'desired results' will be based on this application. Where possible the developments will be kept sufficiently general so that they can be applied to other situations.

By definition, a minimalist system will operate with little to no redundant equipment and thus the configuration of the system must be carefully planned. Hardware characterisation, modelling and optimisation will be performed to maximise the effectiveness of the equipment in order to realise the goal.

The second question is: '*does an irregular configuration of receivers give superior results?'* It can be proven mathematically that the most precise results for any given position-tracking system can be found when the object is in the centre of a regular, n-sided polygon of receivers, where '*n*' is the number of receivers (the measure of the change in precision based on receiver geometry is called geometric dilution of precision[1]) and there are no obstructions. One premise of this research is that this configuration will not provide the best results in situations with complex obstructions (where obstructions are anything that prevents a position measurement). This premise builds on the work of Gillette, *et al.*[2] who discovered that a linear, closed-form solution existed for time distance of arrival (TDOA) calculations that was previously unknown because of the belief that a regular receiver configuration would give the best results; this caused critical matrices to become singular in the calculations which prevented linear calculations. Modelling of the factors affecting tracking system performance and optimisation of receiver con-

figurations will be used to test the premise. It will be shown visually that regular receiver configurations give poor precision in such areas.

Finally we ask: *'can a software model be developed as an expert system for position-tracking systems?'* The term 'expert system' is used here to describe a collection of algorithms that can be run with situation data and will give results that could otherwise only be obtained from consultation with an expert in the field. The premise of this development is that it should be possible to provide an algorithm with data relating to the tracking system (signal reception characteristics, signal penetration, receiver ranges, number of receivers) and the area to be tracked (any obstructions and their shielding effects, tracked areas, etc) and have it calculate the expected performance of the system and suggest the optimum configuration to maximise performance.

1.1 Contributions

This project makes several contributions to expand knowledge in the field of position tracking both locally in New Zealand and internationally.

The Multispectral Solutions, Sapphire tracking system (the hardware used for this research) has been characterised in terms of range and performance. The resulting data are used as a basis for development of visibility models, optimisations and reconstruction algorithms. In the process, a number of key facts about the use of ultra wideband equipment on and around the human body are confirmed.

Significant work has been done on modelling the player visibility that can be expected during a game of rugby. This model focuses on the attenuating effects of the human body on radio frequencies in the 5-7GHz frequency band, and an analysis of the range of the equipment at various angles. This model is applied to the game of rugby; however, the concepts used make it applicable to any sport played on an open field. The model is further adapted to situations where static obstructions exist within the area of interest making it applicable to a wide range of other applications including, but not limited to, the wireless network coverage in an industrial setting or the monitoring and tracking of patients and equipment in a hospital. This model stands by itself as a performance predictor (especially applicable to minimalist systems) and as a component of the expert system.

Several optimisation objective functions have been developed to find the best placement for receivers within a given field of interest to maximise the performance of a particular tracking system in that field. These optimisations, and the validation and verification tests, test the premise that irregular configurations of receivers have superior performance compared with regular configurations in complex environments in terms of availability of positions and precision of said measurements. The optimisations are also key to the expert system.

In situations where signal obstruction has occurred and positions cannot be calculated by the tracking system, a predictive model has been developed that can reconstruct missing positions from the past and present positions of the objects. Reconstruction techniques have been developed for the tracking of general objects and situations specific to rugby. The reconstruction algorithm is shown to be capable of significantly increasing the performance of minimalist tracking systems.

Sports-player position analysis has been identified as a valid real world application for this research. The aim of research in this area is to decrease significantly the time required to perform statistical analysis by automating the position tracking part of the analysis process.

The ultra wide band (UWB) tracking system introduced here employs the first low-data-rate, UWB-communications system that has been allowed in New Zealand. Because of this research, the New Zealand regulatory body for radio spectrum – the Radio Spectrum Management (RSM) branch of the Ministry of Economic Development – has had to address the issue of the use of UWB technology in New Zealand and how it is to coexist with contemporary services.

1.2 Publications

The following is a list of the author's publications relating to this research.

- Brooks, T.T.J., et al. "A review of position tracking methods", Proceedings 1st International Conference on Sensing Technology Palmerston North, New Zealand, pp 54-59, (2005)

- Brooks, T.T.J., et al. "A review of synchronization methods in wireless sensor networks", Proceedings First International Conference on Sensing Technology Palmerston North, New Zealand, pp 338-343, (2005)

- Brooks, T.T.J., et al. "Field coverage optimisation using genetic algorithms", Proceedings ENZCon 2006 Christchurch, pp 272-277, (2006).

1.3 Nomenclature

Within this thesis, definitions of symbols in equations will be given following the first use of that equation. A summary of the symbols and their units is given in Appendix D: Nomenclature.

All units will be given adjacent to the measurement to which they apply. SI units are used unless otherwise stated.

Two numbers, separated by a point will label each equation. The number to the left of the point indicates the chapter in which that equation first appeared and the number to the right indicates the placement of that equation within the chapter relative to other equations. No equation will be given two identifications, thus the result of derivations given in the appendices will be given the same identification as the first time it was used in the text.

Figures will be labelled in a similar fashion to equations. Figures may be reproduced in different sections to aid in the understanding of a particular concept. However, all figures will be labelled based on the first appearance in the text.

The technical nature of this research necessitates the use of many acronyms. For reference, a table summarising the acronyms used in this thesis is given in Appendix D: Nomenclature.

1.4 Industrial Partner

The company Verusco provides player statistics and predictions for events occurring in various situations, but mostly involving sporting events, specifically rugby and soccer. Verusco was identified as a leader in the field of sports analysis. They are the Industrial Partner for the real-world application of this research and will be referred to as such hereafter.

The Industrial Partner currently uses video as the source of player data. Trained operators (known as coders) manually track players and input data from the video into proprietary software that performs the statistical analysis. This process requires anything from 50-600 man-hours per event, depending on the level of analysis required by the customer. Because of the high degree of human involvement required in order to input these data, there is a high error rate and significant time must be

dedicated to checking the results, thus adding a further labour and time component.

A particularly labour intensive part of the analysis is deciding where in a given area a player is. This task is made difficult by several factors:

- Image closeness: If the image of the player is too close then there are fewer reference points to identify that player's location. Conversely, if the image is too far away then it is difficult to identify the individual players. Also, given that the coders have multiple visual angles at their disposal for tracking objects, it is hard to reconcile an object between a wide and a narrow angle shot.

- Occlusion: Some of the camera angles used in the Industrial Partner's video analysis are low angle. Thus there is the possibility that one object can occlude another. The best example of this is during a rugby game when an event occurs at the opposite side of the field from the sideline camera that is filming. If a player closer to the camera runs in front of the camera the image of the player of interest will be occluded.

- Out of Frame: There are some situations where one or more players are outside the field-of-view of the current image. There are several potential causes: a sudden change in velocity of a player in a narrow angle shot; cameras not covering the entire field; or television feeds showing something other than the game (replays for example). A significant number of frames can pass before the camera can be retrained on the players. While the players are out of the shot, their positions can only be roughly estimated.

2 Background

In which we learn of triangulation – video tracking – time-based tracking – combined methods – the tracking of multiple objects – a comparative summary of methods is made – we further learn of the need for synchronisation and methods of achieving synchronisation.

In any research, it is essential that the researcher understand fully the current state of the art within their field.

This research requires a detailed knowledge of the various position-tracking methods and an understanding of the directions that other researchers are taking towards the next generation of position tracking. A significant portion of this chapter will be given over to the clarification of the characteristics of the major position-tracking methods. The results will be used to find the best architecture for the minimalist system.

A cursory examination of the existing body of research highlighted the fact that all position-tracking implementations require multiple receivers and that the clocks of the receivers must be time synchronised[3-5]. Some time will therefore be taken to outline the different methods of achieving synchronisation and how this affects the final system accuracy.

2.1 Position Tracking

Methods of finding the position of a remote object with respect to a known location have been required for several thousand years. Early applications of these methods included: road surveying, target ranging for artillery, map making and in large-scale construction projects.

Today the applications for accurate positioning are wide and varied. Examples include: personal navigation, smart weapon guidance, location of nodes in a distributed sensor network, tracing the origin of emergency calls made from mobile phones, and locating the positions of emergency beacons.

The most prevalent position-tracking methods found in the literature are summarised here: manual triangulation, video-based tracking, time of flight, time difference of arrival and angle of arrival. However, first it is prudent to discuss a source of uncertainty that applies to all position tracking methods: Geometric dilution of precision (GDOP).

2.1.1 Geometric Dilution of Precision (GDOP)

In the simplest terms, GDOP is the uncertainty introduced to a position measurement caused by the physical arrangement of the position location network. It is without units and higher values denote greater uncertainty.

Levanon[1], Zhu[6] and Yarlagadda, *et al.* [7] all explain the process for calculating GDOP (see section 13.1 for derivations). Yarlagadda cites Lee [8] as one of the original research contributors in the field of GDOP however, earlier research exists such as Marchand [9] who examines error distribution in time-difference hyperbolic networks. A more contemporary examination of the relationship between geometry and position accuracy is

performed by Bronk, *et al.*[10] with the focus being on wireless networks.

Consider the following situation in order to visualise the effects of GDOP. Assume that there are two receivers placed at a known distance from one another. Each receiver is capable of calculating the range; where the resulting range circles intersect is the position of the object. Assume also that each receiver has an inherent ranging uncertainty that cannot be mitigated. Given this uncertainty, there will be an upper and a lower range circle for each receiver and the circles from each receiver will overlap, creating an uncertainty area for the object's location.

Figure 2.1 illustrates the position uncertainty region when the object is directly between the two receivers (blue shaded region). Note that positional uncertainty is increased in the direction perpendicular to the line of the receivers.

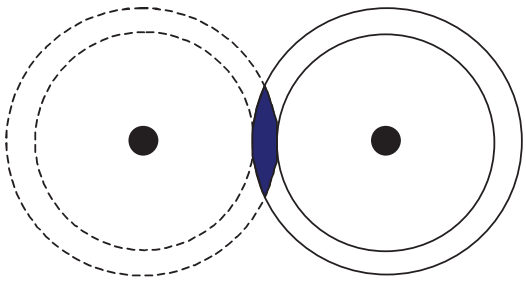

Figure 2.1: Positional uncertainty caused by uncertainty in the range loci for a system of two receivers with the object on the line connecting them.

A complementary situation occurs where the receivers are placed close together and the object is far away in a line perpendicular to the line connecting them (Figure 2.2). The position uncertainty is large parallel to the receivers and low perpendicular to this.

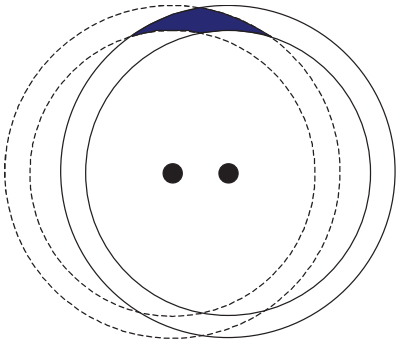

Figure 2.2: Positional uncertainty region of a two receiver time-of-flight where the object is far from the line connecting the receivers.

Figure 2.3 illustrates the sum-of-squares positional error in this simple positioning system as the object moves away from a receiver pair in the perpendicular direction. In this particular plot, the receivers have a 20m separation; the object is equidistant to both receivers; and each receiver has a maximum range uncertainty of ±2m[1].

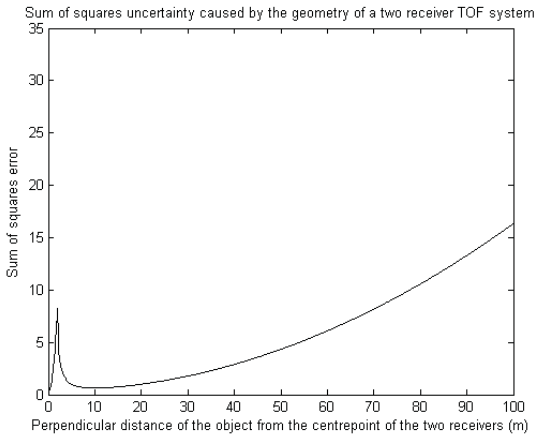

Figure 2.3: The sum of squares position uncertainty perpendicular to a pair of TOF receivers.

It is clear from this figure that there is an exponential decrease in the positional uncertainty as the object initially moves away from the receiver pair then has a gradual in-

[1] This error value was chosen for the example as it makes the trend of the position error easier to see. Changing either the timing uncertainty or the separation of the receiver pair has the effect of scaling the error plot; however, the relative shape will remain constant.

crease in uncertainty as it moves beyond the distance equal to half the separation between the receiver pair.

This confirms that the uncertainty in the perpendicular dimension is dominant when the object is near to the receiver pair and the parallel uncertainty is dominant when the object is far from the receiver pair. It is also interesting to note that GDOP calculated using Levanon's method, for the same situation, will produce exactly the same trend.

GDOP affects all positioning systems in similar ways. Given this common source of uncertainty in all types of positioning system, much work has gone into reducing GDOP as a component of overall uncertainty. Work in this field is a subset of Geomatics Engineering (a discipline focusing on spatial information). For the purposes of mitigating GDOP, Levanon[1] has conducted a number of simulations that show that the lowest GDOP is obtained when the object is at the centre of an N-sided regular polygon (where N is the number of receivers in the system). The results of one of these simulations are shown in Figure 2.4. Here the lowest GDOP is realised in the centre of the receiver formation.

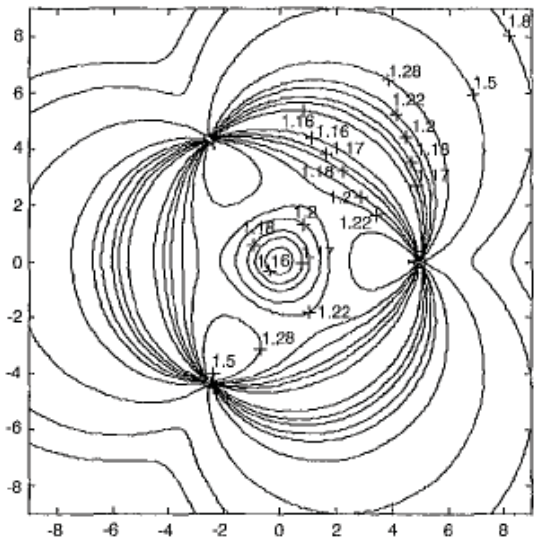

Figure 2.4: GDOP with three receivers in a triangle formation.

Given that GPS is currently one of the most popular implementations of position tracking, it is natural that research is being conducted to increase its precision. Yarlagadda, *et al*.[7] review the GDOP metric as it applies to GPS applications. Zhong, *et al*.[11] have produced a fast algorithm for calculating GDOP in GPS calculations. The use of fast algorithms for obtaining GDOP figures opens the door for systems that can select which receivers to use in a given position calculation to get the highest precision for any given position calculation. Bo[12] proposes such a system to allow a navigation using multiple position location networks (in this case GPS and Galileo).

Moving away from GPS and into pure mathematics, Chaffee, *et al*.[13] link the GDOP matrix to the Cramer-Rao bound (CRB or sometimes CRLB for Cramer-Rao lower bound). The CRB is a part of estimation theory in statistics that can be used to measure the 'efficiency' of a measurement. A solution that is efficient and hence approaches the CRLB achieves the lowest possible mean squared error. Reference to the CRB appears often in the field of position location.

2.1.2 Basic Triangulation

At the heart of many positioning methods is the idea of triangulation: the use of the properties of triangles to calculate distances and hence find the positions of objects. Credit for the discovery of the underlying theory of triangulation goes to the Ancient Greek philosopher Thales (Lahanas [14]). Because of the ubiquitous nature of Thales' theorems in modern positioning methods, a basic description of triangulation is given here.

Given any two reference points, knowledge of the angles between both references and the object, and the distance between the reference points, it is possible to calculate the distance from one reference point to the object. Figure 2.5 illustrates this.

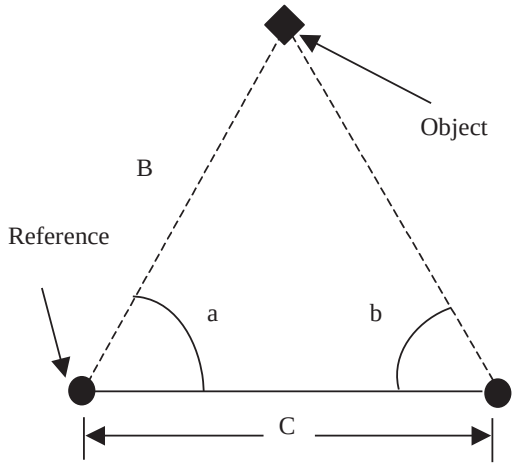

Figure 2.5: Mathematical approach to triangulation.

The distance *B* can be found by using the sine rule for triangles[15]. Equation (2.1) shows the formula rearranged to find this distance. Note that *a*, *b*, and *C* must all be known.

$$B = \frac{C \sin b}{\sin(180 - a - b)} \tag{2.1}$$

Where:

- *B* and *C* are distances shown in Figure 2.5 measured in metres.
- *a* and *b* are the angles shown in Figure 2.5 measured in degrees.

Originally, triangulation was used for navigation, surveying and civil engineering purposes, such as cartography, and later for finding the range of targets for artillery strikes. Tools of the time include the astrolabe, sextant and theodolite.

In modern times, the applications of triangulation remain the same but the methods of calculating positions have evolved to take advantage of better technology to increase the accuracy or speed of measurements.

2.1.3 Video position tracking

In general, the term video position tracking can be applied to any system in which the positions of objects are located using video footage of the area of interest. An example of such a system is presented by Charlton, *et al*.[16] who present a system that provides tracking information to a two-degree-of-freedom actuator by means of a monochrome video image.

If placed appropriately, a video tracking system requires only one calibrated camera to calculate positions in two dimensions. The ideal location for placement of the camera is orthogonal to the area of interest. For example, in a game of rugby the best location for the camera would be in the zenith direction above the field looking down, as this will minimise the distortion and obstruction in the image. Unfortunately, it is seldom possible to mount a camera in this location. In these situations, it is desirable to place the camera as close to the zenith as possible. As the camera moves away from the zenith the distortion present in the video image increases.

All automated video-tracking systems rely on the objects being identifiable in each of the video frames. So long as the objects to be tracked can be seen by the camera, then the positions of those objects can be calculated. If an object is occluded, then a model-based approximation can be applied to give the object a valid position.

When considering multiple moving objects, the probability of one or more objects being occluded becomes significant. Here multiple cameras can be used to decrease the probability of occlusion.

Identification of objects in an image is a critical issue for video-based tracking systems. The objects must have some identifying feature that is visible in enough images for the analysis algorithm to be able to lock on. In the game of rugby, unique identification of each

player would be very difficult. It is a contact sport where groups of as many as 16 players will be in close physical contact. Since teams of players wear uniforms whose most prominent distinguishing feature is the number on the back, expecting any one player to be identifiable by one or more cameras at all times is unrealistic. A level of human intervention is required to positively identify players.

The accuracy of the position measurements from this type of system is dependent on factors such as the number of cameras, the accuracy of their placement and the ability of the processing algorithm to correct for the distortion present in the images and calculate a position at any point in the entire field of interest. If more than one camera is used then the streams from each must be synchronised so that the processing algorithm can combine data from each stream.

Research in the area of vision-based tracking has wide ranging applications that highlight the flexibility of this kind of tracking. Vision-based tracking appears to be popular in applications where non-invasive tracking is required and where there is little to no known information regarding the environment. Some of the general areas of research and the projects being conducted are summarised below.

The use of vision type systems to give autonomy to robots or production machinery is a popular research area. The following paragraphs give examples of this type of research.

Work by Kais, *et al.* [17] on using Geographic Information System to allow mobile robots to locate themselves in an urban environment.

Cevher, *et al.*[18] and Chellappa, *et al.*[19] present systems that use a combination of video and acoustic sensors to track objects.

There has been much research into the security applications of video tracking. Yang, *et al.*[20] propose a new system for automatically moving a gimballed camera to keep an object in the centre of the image. The aim of this research is to increase the effectiveness of security systems. As an extension to this, Kakiuchi, *et al*[21] discuss a system for tracking an object as it moves across multiple sensors (they refer to these as neighbour nodes) to allow a tracking system to automatically pursue the object. Similarly, Li, *et al.*[22] propose to use an interval recursive least-squares filter to produce a tracking algorithm for objects in an image. Li, *et al.*[23] discuss an efficient way to track an object moving through a crowd in real time. They deal with the problem of occlusion by employing what they call a dominant colour histogram.

Head tracking is another important area of research. As the name implies, head tracking is concerned with the location of a person's head relative to their body. Crowley, *et al.*[24] identify the need for accurate object tracking when performing facial recognition and encoding operations; the person's face must be held steady in the image if critical features are to be extracted. Yu, *et al.*[25] also use the concept of head tracking in order to detect when elderly or infirm people fall down. This is a non-invasive way to monitor someone's health.

Combining data from multiple sensors is a challenge in video tracking and research continues to progress in this area. Snidaro, *et al*[26] have successfully performed experiments that fuse the data from multiple video streams. This has allowed them to better identify and classify individual objects in the area of interest. Snidaro, *et al.*[27] extended this work two years later to include automated classification of objects where the classification is a learned response from analysis of the heterogeneous features of the tracked objects. Combining multiple video streams can be a complex task, Akazawa, *et al.*[28] suggest a simplified video tracking that uses a maximum of two cameras and can generate position information in real time.

Biological observation is another application of video-based tracking. It is particularly useful because it is non-invasive thus allowing the organism to interact with its environment in the most natural way. A video tracking system is proposed by Wang, *et al.*[29] that can be used to measure the body deformation of a swimming fish. Structured light is used to allow measurements to be made from the images and a control module orients the vision system to track the movements of the fish. Brooks Zurn, *et al.*[30] describe a system that is similarly non-invasive for tracking rodents under near-infrared illumination. The system can distinguish rodents from their background (which is a similar colour in near infrared) and classify their basic actions.

Vision-based tracking can be used to track, identify and classify events occurring in pre-recorded situations as well as be used in real-time applications. The advent of the 'visual database' can save time on classification of film or images and increase the efficiency and accuracy of searches for particular events. There are a number of projects going on in this area. Schonfeld, *et al.*[31] created a visual search engine called Vortex that can retrieve video from a compressed multimedia database. The search engine takes a template image as input and finds video sequences based on the objects in the input. Mokhtarian, *et al.*[32] propose a similar system that they refer to as content-based video retrieval in which an object, specified by the user, is tracked both forwards and backwards in the video via the extraction of corners from each frame. Jiang, *et al.*[33] take the novel approach of examining any text (such as closed captions) embedded within a video to highlight events occurring within that image.

In addition to research relating to the applications of video-based tracking there are projects aimed at improving the sensor systems and controls. The mechanics of producing a higher resolution sensor are beyond the scope of this thesis but the research these sensors encourage is not. As stated previously, mounting a sensor at a low azimuth angle with respect to the area of interest leads to distortion in the image. The method for dealing with this distortion is to warp the image using a matrix transformation which compresses the near part of the image, and expands the far part and the depth of the image. Choe, *et al.*[34] propose a new and more efficient way of warping high resolution images to allow the required transformation of video in real time.

An obstacle to position tracking through video is low frame rate. Many systems use the mean-shift algorithm which relies on tracked objects not moving too far between frames. Mean shift does not perform well in low-frame-rate environments. Li, *et al.*[35] address this problem by enhancing the mean shift-algorithm to use prediction about where the object will appear next. The idea of tracking using prediction is not uncommon and is also suggested by Bouttefroy, *et al.*[36] in their particle filtering approach to object tracking. Here the object trajectory is transformed from the real world to the camera plane to increase the accuracy of position measurements.

Work is also being done to allow single sensors to perceive depth within an image. This further increases the effectiveness of video-based tracking system. Seal, *et al.*[37] propose a compact vision system for calculating the depth of an object using a single camera and two sets of planar mirrors to provide an offset image. The intended application for this system is on an asparagus canning line where the size of each stalk must be classified to correctly fill each can.

Orchard[38] examines the location uncertainty in images and shows that linear decompositions such as transformations cannot fully mitigate this uncertainty. Based on these observations, Orchard proposes improved image processing tools for minimising uncertainty. Modelling of such uncertainty is what has al-

lowed advances in the areas of software correction of distortion and adaptive optics for telescopes as reported by Bikkannavar[39].

Video-based tracking is also used in various sporting applications where knowing the positions of objects on the field is advantageous without interfering with them.

Various robot sports leagues now exist as a spur to further research in this area. One of the most prevalent is the Federation of International Robot-soccer Association (FIRA). In MiroSot, a game type supported by FIRA, a single video sensor is mounted at the zenith of the pitch and it is the job of the competitors/researchers to use the information that can be gleaned from the video. Jiang *et al.*[40] describes a particular algorithm for locating the ball and the robots within the video stream.

There is similar interest in video-based tracking for sports played by humans, especially soccer (the likely reason for this is soccer's level of financial backing). There is a large body of research on tracking both the players and the ball in both real-time and previously recorded games.

The work of Zhang, *et al.*[41] and Utsumi, *et al.*[42] plumb the depths of recoded soccer games to extract positions and trajectories of players and balls. Each is specifically designed to deal with the shortfalls of broadcast video (such as blur, camera motion and noise). Mya, *et al.*[43] and Zhu, *et al.*[44] both describe systems with similar aims of tracking multiple players. Intille, *et al.*[45] discuss a new approach for tracking what they refer to as 'weakly modelled objects'. They have the aim of semi-automatic annotation of sports videos.

The following two research groups have combined aspects of previously discussed research and applied them to tracking players during a soccer match. Han, *et al.*[46] have developed a sports-video analyser to extract key events from a video stream and use these as markers for a video database. Saito, *et al.*[47] use a multi-angle, multi-camera system. Data fusion techniques are used to combine the streams and track the players.

In a sport other than soccer, Li, *et al.*[48] present a system that can extract pertinent details of a competitor from video streams in which there is a moving background. The system was tested on competitors in diving and jump games.

As an offshoot of sports tracking, research has been conducted into the generation of three-dimensional scenarios where a consumer can view a sports game from any angle, including the players' point of view. While this application is not strictly video tracking, it does use many of the same techniques such as vector predictions and some form of player tracking. Kasuya, *et al.*[49] propose a system that can automatically generate a player's view based on an actual game of soccer. The system described by Saito, *et al.*[47] above is also capable of generating a player's view.

2.1.4 Time of Flight (TOF)

This method, interchangeably known as Time of Arrival (TOA), uses the time taken for a signal to travel between the object and a receiver to calculate the range of the object from the receiver. With multiple receivers and knowledge of their locations, multilateration can be used to find the position of the object.

The range is calculated using the time, distance and velocity relationship. With only one receiver the position can be defined only as a circular locus. Three receivers are required to fix a position uniquely in two dimensions and four are required for a position in three dimensions.

An example of a three-receiver network calculating a two dimensional position is given in Figure 2.6.

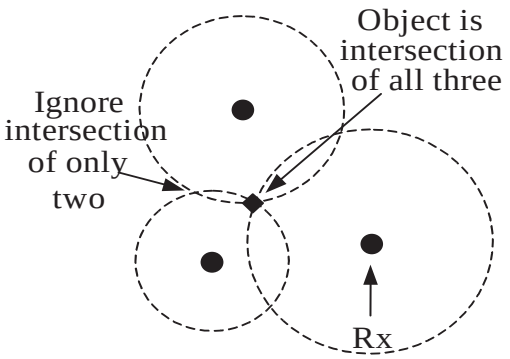

Figure 2.6: TOF positioning method.

In Figure 2.6 the solid circles are the receivers (whose locations are known), the solid diamond is the object of unknown position and the dotted circles represent the circles defined by the range calculated from time-of-flight measurements.

In a TOF-tracking system using multiple receivers, the position of the object is found directly as the intersection of a number of circular loci, centred on each receiver. Solving the equations simultaneously yields the unique intersection point common to all loci. The general form of an equation for a circle is given in equation (2.2).

$$(x - X_r)^2 + (y - Y_r)^2 = R_{tr}^2 \tag{2.2}$$

Where:

- x and y are the co-ordinates of the point of intersection that is to be found.
- X_r and Y_r are the coordinates of a particular receiver.
- R_{tr} is the radius of the circle, which represents the TOF path of the signal between object and the receiver.

This simultaneous-equation solution relies on there being a single point of intersection of all loci. It is a certainty that noise, at some level, will appear in the measurements, resulting in no single point of intersection. Such noise might have a number of sources including: poor synchronisation, limited timing precision at the receiver, or multi-path interference. The likelihood of a situation like this occurring increases with more receivers. Using a direct calculation method in this case is much more difficult.

An alternative, iterative approach is commonly used. For a given set of time-of-flight measurements and knowledge of the receiver locations it is possible to find a position that best fits the measured loci in a least-squares sense. Equation (2.3) represents the objective function of this iterative calculation[50].

$$e = \sum_{r=1}^{N} [\sqrt{(x - X_r)^2 + (y - Y_r)^2} - v(T_r - t_0)]^2 \tag{2.3}$$

Where:

- e is the sum of squares error of the current iteration.
- N is the total number of receivers.
- x and y are the current coordinates of the position to be evaluated.
- X_r and Y_r are the receiver coordinates.
- v is the velocity of the signal.
- T_r is the observed time of flight of the signal.
- t_0 is the absolute time of the start of this sample or epoch.

The term within the square root is the current estimate of the distance between the object and the receiver, while the second term represents the distance based on the time-of-flight measurement for the current sample, or epoch.

The value of t_0 may be non-zero in TOF-tracking systems where the start time of transmission is unknown. An example of this type of system would be a tag-based TOF-tracking system where the tags transmit on a pseudo-random basis. In these cases, the value of t_0 is calculated as another variable to be optimised in the iterative process.

The level of synchronisation in this type of system is dependent on the exact way in which transmissions are triggered. If the system uses an out and return method or a pseudo-random transmission, then there need only be synchronisation among the receivers. If however, the transmitters output at regular intervals then there must be synchronisation between receivers and transmitters. The accuracy of time measurements across the receiver network will strongly affect the accuracy of calculated positions.

Multi-path interference must be considered in any time-of-flight system, for it represents a potentially crippling source of error. This type of interference is seen when the signal from the transmitter reflects off a nearby object before arriving at the receiver (see Figure 2.7). When a signal is thus reflected, the path length is longer than the line-of-sight path. The two or more signals arriving at the receiver will have a phase difference, which can cause destructive interference, reducing to zero the signal strength at the receiver.

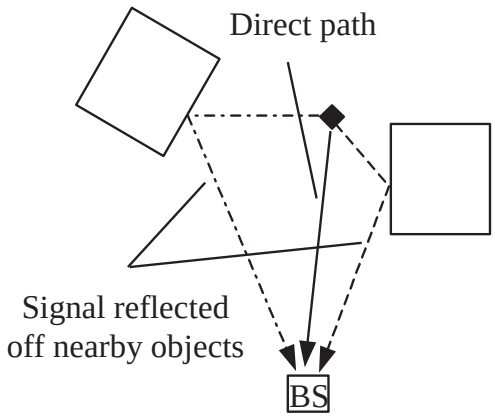

Figure 2.7: Base station experiencing multi-path fading.

A more detrimental effect on positioning systems occurs when the line-of-sight path signal is dominated by one of the indirect paths. In this situation the positioning system will calculate the time of flight based on the longer, indirect path resulting in an incorrect position estimate.

Note how in Figure 2.7 the reflected paths are longer than the direct path. If the line-of-sight path is obstructed, then the receiver will see a signal from the object that has an erroneously high time of flight.

There are several methods currently being investigated to combat the problem of multi-path interference: shaping the transmitted pulse so that the available bandwidth is used more completely as proposed by Runtz, *et al.*[51]. Grosicki, *et al.*[52], Li, *et al.*[53] and Mathias, *et al.*[54] include, as part of their trilateration algorithm, a measure of non-line-of-sight and near-far-effect in order to select only the most reliable sources of time data. By using the ultra short transmissions of Ultra Wide Band (UWB) the effects of multi-path interference can be partially mitigated since it is less likely that extremely short transmissions will overlap[55]. Some decision must still be made, however, on which signal arrived on the direct path from the object.

Any signal whose propagation speed can be modelled is suitable for use in a time-of-flight tracking system. Examples of signals suitable for use in a TOF system include:

- Sound waves[56],
- Electromagnetic waves, specifically radio and microwave frequencies, and
- Laser, a coherent subset of electromagnetic waves in the visual or near visual band. Buzinski, *et al.*[57] give an example of a laser based TOF tracking system.

The suitability of a particular signal would be governed by the requirements of the application. For example, the laser and ultrasonic are best suited to applications where the objects need not be individually identified. Radio and microwave transmissions are ideally suited to applications with multiple moving objects that must be tracked individually.

Applications of TOF tracking include:

- Surveying for road, bridge, tunnel and other civil engineering applications;
- Finding the range to target in either recreational hunting or military applications,
- Airborne surveying[58], and
- Range finding of objects to increase the functionality of both autonomous and teleoperated robots[59].

There are a number of different tracking systems that use TOF. Chadwick, *et al.*[60] discusses a vehicle location system to track one or more asynchronous and aperiodic transmitters. The system is over defined and uses GDOP measurements to select the best receivers with which to calculate positions. Manolakis, *et al.*[61] define a generic TOF system that can calculate positions in three dimensions. They describe the system as being simple and efficient and provide performance analyses under various conditions. A system that is similarly described as being efficient at calculating positions in three dimensions is presented by Farve-Bulle, *et al.*[62] for tracking robots.

Distributed wireless sensor networks have, in recent years, become popular methods for making measurements over wide areas. This project, the tracking of multiple moving objects over a contained area, is a subset of wireless sensor networks and as such, research in this area has an impact on this project. The problem of quickly and accurately calculating the positions of each node has sparked new research. Cai, *et al.*[63] and Yu, *et al.*[64] both suggest algorithms for determining the positions of nodes in a wireless sensor network using a technique called alternating combination trilateration: a form of trilateration that uses a varying number of triangles to calculate a position. Once again in the area of wireless sensor networks, Yang, *et al.*[65] have studied the limitations of using trilateration and have proposed a novel and elegant system that improves the performance of the system over the original. Another high precision, low cost tracking system is presented by Tian, *et al.*[66]

A subset of wireless sensor networks is the swarm system. Biologically inspired, this type of system uses a number of separate robots working collaboratively towards a common goal. Like wireless sensor networks, it is vital for the swarm to know the positions of each robot if they are to work effectively. Sun, *et al.*[67] use multiple, energy-efficient antennae to address this problem in the swarm.

2.1.5 Time Difference of Arrival (TDOA)

The TDOA positioning method is similar to the TOF method in that both use the signal travel time to infer a range based on a measurement of time. Both methods are referred to as multilateration. Also the same signalling techniques are applicable. The measurement

of time used in a TDOA position-tracking system is taken as the difference of an incident signal arriving at two receivers that have a spatial separation. A TDOA system only ever requires synchronisation between the receivers since it is not necessary for the time of transmission of the signal to be known.

Dong-Ho, *et al.*[68] present a comparison between TOF and TDOA systems, focusing on GDOP, position estimates and error covariance.

For a position in two dimensions, a TDOA system requires a minimum of three receivers. The calculations to find the position, however, are complex and require significant knowledge of the approximate location of the object (see Appendix A: Equation Derivations for a derivation of the three-receiver solution). There is a body of research that deals with this issue. For example, Chaffee, *et al.*[69] discuss the situation in which there are multiple solutions to a set of pseudo-range equations and methods for determination. Conversely, Gayer, *et al.*[70] constrain their TDOA system so that only one solution will be valid.

If four receivers are used then the calculations become much simpler (see Appendix A: Equation Derivations for derivation of the four-receiver solution)[71]. Gillette, *et al.*[2] developed a simple, closed-form solution to localisation using TDOA data from five microphones. Yang, *et al.*[72] describe a direct approach to derive target positions from TDOA measurements over a wide area. Chan, *et al.*[73] analyse another system whose performance approaches the Cramer-Rao lower bound.

For each time-difference value, a locus of positions can be defined which lies on a hyperbola. The total number of hyperbolae available for use in position calculations equals the combinations of ways the receivers can be split into pairs. This is summarised in equation (2.4)

$$h = {}^{r}C_{2} \tag{2.4}$$

Where:

- h is the number of hyperbolae available.
- r is the number of receivers involved in calculation.
- C is a representation of the statistical-combinations function.

As with TOF, either a direct or an iterative solution can be employed to calculate the object's position. The derivation of the direct solution for a four-receiver system can be found in Appendix A: Equation Derivations. Fontana gives an equation that can be used for the iterative solution[50]. Ho, *et al.*[74] provide an alternate, iterative solution.

A TDOA system requires receiver level synchronisation so that measurement of time difference between the arrival times of the signals at the various receivers has meaning. The transmitters need not be synchronised since the epoch start for any given transmission is irrelevant. As well as describing a direct TDOA solution, Yang, *et al.*[72] and Chaoshu[75] also extol the virtues of using satellite clocks to synchronise their widely-spread receiver network.

The most widely known example of a TODA positioning system is the Global Positioning System (GPS). In this architecture, there are a number of orbiting satellites and each user has a receiver. Each satellite transmits a signal that is encoded with the current time and the position of that particular satellite. With a signal from three or more satellites, the receiver is able to use multilateration to calculate its position. GPS uses a system that is the reverse of what has been described previously (multiple, known transmitters used to locate a single, unknown receiver) but the process is the

same. Braasch, *et al.*[76] give a review of GPS and its associated hardware.

As discussed in the previous subsection, localisation of nodes in wireless sensor networks, such as the ones that would be applicable to this project, is very desirable. Some researchers have chosen to implement TDOA in order to achieve this. Zhao, *et al.*[77] propose a practical implementation of such a system. Their major claim is that it minimises assumptions on architecture in order to realise a generic localisation method that can be applied to a wide range of networks. Cheng, *et al.*[78] have also produced research in this field. Their work uses local signals to identify and localize the nodes in the network.

There are some applications of multilateration in sporting applications. Guizzo[79] reports on the use of TDOA to track a soccer-ball on the field of play. This particular system uses receivers mounted on stadium light posts that are fibre-optically connected to a central server to get 0.1m accuracy.

Examination of the body of research suggests that the use of TDOA tracking systems is very popular as a form of secondary radar at airports. The term secondary radar is in common usage in the aviation industry and refers to the method for tracking aircraft movements that is the alternative to conventional radar. Secondary radar involves a transponder on the aircraft transmitting (known as 'squawking') an ID code as instructed by an air traffic controller. Stations on the ground compute the position of the aircraft from these transmissions. Transponders that are 'Mode-C' enabled embed altitude information in the transmission, giving the air traffic controllers situational awareness in three dimensions. The advantages of secondary radar over primary are so great that in many places (e.g. Ohakea Air Force Base in New Zealand) the primary radar site has been decommissioned.

Anyone who has flown from a busy airport knows how many movements there are and how important it is that air traffic controllers have accurate, real-time information on those movements. The work of Geyer, *et al.*[70], Chaoshu[75] and Wood[80] has been aimed at increasing situational awareness at airports.

An extension of secondary radar research is 'Mode-S' that allows controllers to track the movements of aircraft that are approaching or leaving the terminal. Leonardi, *et al.*[81] discusses the use of multilateration in this setting.

Position tracking of mobile-phone users is another application that has sparked research. Government regulation has been a driving force for this; it will soon be a requirement that an emergency call made from a mobile phone be localized to at least 150m[82]. Spirito[83] discusses a method for enhancing the existing GSM (Global System for Mobile communications) to achieve this. It uses the existing, synchronised cellular network to implement multilateration positioning. Wylie-Green, *et al.*[84, 85] have written an algorithm that can simulate the accuracy of such a system. Laitinen, *et al.*[86] highlight the difficulty of localisation in a Non-Line-of-Site (NLOS), urban environment and discusses alternate methods of positioning in such situations. Zhao[87] reports on the progress of standardising position tracking in the 3G cellular network.

2.1.6 Angle of Arrival (AOA)

This method of position tracking involves using several receivers that can calculate the angle of incidence of an incoming signal with respect to a reference and the principles of basic triangulation (see section 2.1.2). Figure 2.8 illustrates operation of a two-receiver, AOA, position-tracking system.

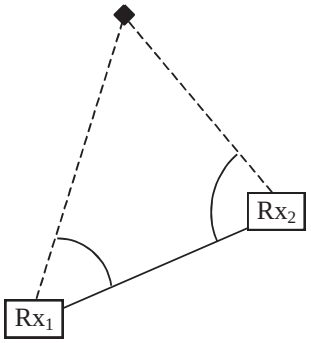

Figure 2.8: Basic operation of a two receiver AOA positioning system.

In RF applications, measuring the angle of incidence of a signal is accomplished by using some form of directional antenna array.

One example of such an antenna is the phased array antenna (PAA) in which the angle of incidence is calculated by measuring the phase difference of the signal at several receiving elements aligned in a known pattern. PAAs are generally used to control the direction of a transmission rather than measure the angle of incidence of a signal but the process is commutative.

PAAs are becoming more popular in many disciplines because, with no moving parts, the antennae have a long useful life and are very resilient to extreme conditions. As an example of an application where the resilience of this type of antenna is desired, NASA's Messenger space probe has been fitted with a PAA for its voyage around the inner solar system where it was deemed that the standard gimballed antenna would not tolerate the harsh conditions[88]. This use of PAAs highlights the robustness of the solution event though it is not an example of positioning.

Another directional antenna type is the Doppler antenna in which the Doppler shift of an incoming signal is measured with respect to an antenna array that simulates a single element moving in a circle[89]. While in principle a Doppler shift antenna can be constructed by mounting a single receiving element on a rotating base, in practice it is infeasible to operate in this manner, because of the high rotational speed required for operation.

Consider an antenna constructed in this way (shown in Figure 2.9). In order to observe a change in the frequency of the signal at the receiving element, that element must be moved fast enough to produce a Doppler shift great enough to overcome any noise or modulation present within the signal. Driving the antenna element to this speed is inefficient, the antenna would have to be able to handle being rotated at speed and the rotation would require hardware to transport the signal from the antenna to the electronics. This hardware might be some kind of rotating mechanical connection similar to the commutator in a DC electric motor or a wireless transmission system. Either would increase the complexity of the antenna system.

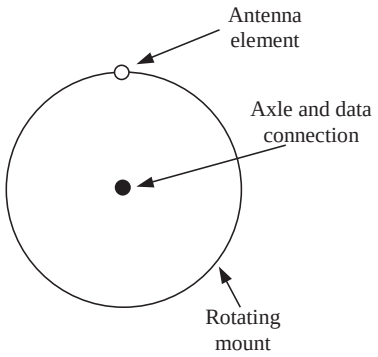

Figure 2.9: A single element rotating Doppler antenna.

In order to construct a functional Doppler antenna, it is possible and advantageous, to use an array of static receiving elements mounted around the circumference of a circle, as shown in Figure 2.10. The signals on these antennae are sampled in turn by the electronics of the system that calculates the angle of incidence. This emulates the characteristics of a single rotating element. Sampling the signals at sufficient speed and without adding noise

requires specialized, high-frequency switching components.

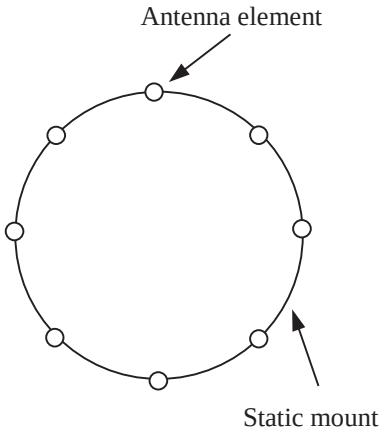

Figure 2.10: A multi-element antenna array that operates as a Doppler antenna.

By observing the change in frequency of the signal seen by the simulated rotating antenna, it is possible to determine the signal's angle of incidence.

When using a purely angle-based position-tracking system, there is less need for time synchronisation in the system. The receivers need only be time synchronised to a level where angle measurements can be identified as applying to a single object.

It is possible to have an AOA system that is self-configuring even if the positions of only two receivers are known. If two receivers are set up a known distance from each other and a number of other receivers in the system transmit an identifying signal, then the positions of these extra stations can calculated as though the receivers were objects. A centralised management unit would then be able to keep track of all of the objects and receivers in the system. An AOA system of this type would enjoy the benefits of minimal configuration time. The concept of a self-configuring AOA system was presented by Kirchner[90].

A common reference of direction is required across the entire AOA receiver network for the angle measurements to have any meaning. This implies that AOA antennas must either be static or employ some way of measuring the orientation of the station with respect to a reference.

A minimum of two receivers is required to calculate a two-dimensional position using the AOA method. More would increase the accuracy of measurement but they would not necessarily allow for calculation of a three-dimensional position. It is assumed that the receiving stations are able to calculate the AOA in only one plane. An antenna array that could measure the AOA in two planes would be required to calculate a position with a third dimension.

More receivers in the network will result in more accurate position measurements. As with other position-tracking methods, there exists an iterative solution that lends itself to over-defined situations such as this. The iterative solution minimises the sum of squares error between the measured AOA and the AOA calculated for the current position estimation. The objective to be minimised is shown below in equation (2.5).

$$e = \sum_{r=1}^{N} (\alpha_r - a_r)^2 \qquad (2.5)$$

Where:

- e is the error of the current iteration.
- r is the number of the receiver from 1 to the total number of receivers N.
- α_r is the calculated angle between the object and a given reference for the current iteration variables.
- a_r is the measured angle from the hardware.

A minimal case, two-receiver AOA system will fail when the object is in line with both

receivers, as the denominator of equation (2.1) sums to zero. To avoid this, the receivers could be configured such that no valid position will ever appear on the connecting line, or to include more receivers to compensate.

Like other positioning methods, investigations are in progress to advance knowledge in the area of AOA positioning. AOA has been applied to a number of wireless sensor networks such as the one proposed by Kirchner[90] or the one by Niclescu, *et al.*[91] that exploits the AOA capabilities of some wireless network nodes.

The effects of mulitpath interference are not mitigated by using AOA. Curry, *et al.*[92] address this issue in their filtering system applied to an AOA positioning system. The filter is designed to remove incident signals that are likely to be other than the direct path. Similarly, Wang, *et al.*[93] investigate the effect of special fading on the symbol error rate of a transmission. Both Lo, *et al.*[94] and Dai, *et al.*[95] have implemented neural networks as a means of performing AOA calculations in multipath environments.

2.1.7 Combined angle and time-based methods

A combination of the angle and time-based methods of positioning described in the previous sections could result in a compact and powerful system for position tracking.

With the ability for an angle-based system to calculate the angle of incidence and a time-based system to calculate range, a complete position-tracking system could be realized with a single receiver. However, careful design of such a system would be required to ensure that the limitations of the individual methods do not adversely affect the system.

In a single receiver realization of a combined system, any obstruction of the objects would result in a failure to calculate a position, thus reducing the reliability of the system. This limitation renders the single-receiver, combined solution unacceptable for tracking multiple, moving objects where the probability of occlusion is significant. Adding redundant receivers would make the system feasible for this type of situation.

Regardless of this, the advantages offered by hybrid systems have been explored by some researchers. Cong, *et al.* [96] use a hybrid TDOA/AOA system to locate mobile users of a code-division-multiple-access (CDMA) network. A similar hybrid system is proposed by Li, *et al.*[97] to provide support to military surveillance.

The use of multiple techniques for positioning requires methods for combining the results from each component of a given system to give the best results. Kleine-Ostmann, *et al.*[98] present a model for efficiently combining position data from both time and angle-based positioning methods for both military and civilian applications.

Yang, *et al.*[99] use a hybrid TDOA/AOA that uses information in the AOA calculation to estimate multipath effects. Their tests show that accuracy is increased when using this system.

2.1.8 Tracking Multiple Objects

Many present day applications of position calculations require that the positions of multiple objects be tracked. Examples of the applications of multiple object position tracking include:

- Finding positions of sensors in a distributed wireless sensor network, tracking animal movements and biometric data[100].
- Increasing the safety and security of the elderly or infirm by facilitating their remaining in their own home, while remotely monitoring their movements[101].

- Tracking product shipments, assets or employees in businesses[55].

- Sports analysis to help train players in both general and specific tactics in order to increase their performance[44-48].

Vermaak, *et al.*[102] have proposed a unifying framework for tracking multiple objects. This framework deals with some of the issues associated with tracking multiple targets such as the, sometimes, invalid assumption that the number of objects will be constant (or that they will be the same objects). The group uses the idea of existence (a concept also used by Musicki, *et al.*[103, 104]), data relating to the birth and death of objects, in the area of interest to the time horizon. Vermaak's group have used probabilistic methods to achieve this goal. Emphasis is placed on making a system that can be used in conjunction with existing position-tracking architectures.

All of the position-tracking methods described could be used to track the positions of multiple objects, but they would require some modification. The most important of these is that each object to be tracked must have some uniquely identifiable characteristic.

In a reflected signal system, individually identifying each object appears to be relatively easy since the transmitter is directed towards the object whose position is required. Problems arise when the objects are moving, the system must predict the movement to keep that object aligned with the beam. When there are multiple objects that converge and diverge, the system can no longer distinguish which is which.

Radio-tag based systems are ideally suited to tracking multiple objects since they transmit a unique ID as part of the position-tracking signal. This neatly circumvents issues of multiple objects requiring unique identification and existence.

2.1.9 Position Tracking Literature Conclusions

Several methods for position tracking of objects have been described and discussed. All methods have been shown to be compatible with either a direct or iterative solution. The range capability of each system will be entirely hardware specific. The systems have also been shown to be suitable for use in most distributed wireless sensor networks. A comparative summary of the characteristics of each position-tracking method is given in Table 2.1.

Table 2.1: Comparative summary of position tracking methods.

Positioning method	Receivers required for a 2D position	Level of synchronisation required	Potential sources of error	Special requirements
Video-based tracking	1	Frame synchronisation when using multiple cameras.	Image distortion, occlusion.	Lens distortion characteristics should be known.
Time of Flight	3	Receiver level time synchronisation plus some knowledge of the sample or epoch start.	Poor time synchronisation or precision, multi-path interference, configuration geometry.	
Time Difference of Arrival	4	Receiver level time synchronisation.	Poor time synchronisation or precision, multi-path interference, configuration geometry.	
Angle of Arrival	2	None	Low angular resolution, configuration geometry.	Receivers must have angle sensitive components
Combined Angle and Time-based system	1	Receiver and transmitter synchronisation for single receiver configuration.	Dependent on the configuration of the system.	

2.2 Synchronisation

Time synchronisation of electronic systems has gained in importance as the requirement for more accurate time measurements has increased. This is especially true for systems where distributed nodes must collaborate to perform a task, such as sensor-data fusion operations, coordinated actuation and power-efficient duty cycling.

As the requirement for greater accuracy has increased, so too has the difficulty. A calculation requiring the measurement of nanosecond-scale events between multiple nodes also requires those nodes to be synchronised within a nanosecond with sufficient

stability to retain synchronisation for all measurements.

2.2.1 The Need for Synchronisation

Some examples of the need for time synchronisation are given below.

In a time-division-multiple-access (TDMA) system, each device is given a time slot in which it is allowed to transmit data. Occasionally all the local clocks must be synchronised to ensure that a device does not transmit out of turn, thereby interfering with other transmitters. GSM mobile equipment is an example of such a TDMA system[105, 106].

Another example is a position-location network (PLN) in which a number of receivers, of known location, track the position of mobile nodes using a time-based positioning method (TOF or TDOA)[107]. GPS is an example of such a PLN, and its method of keeping accurate synchronisation will be discussed later in this section. Since the subject of this thesis is concerned with a PLN it is vital that there is an understanding of synchronisation methods. Of particular interest within the scope of this project is the synchronisation of distributed networks and synchronisation for the purpose of accurate event timing.

It should be noted that synchronisation is only required in applications where several nodes in a network must collaborate. In other situations it may be possible to perform operations in an asynchronous fashion, such as in frequency division multiplexed communication systems or asynchronous communication protocols.

The predominant cause of timing errors, and the need for frequency synchronisation, is jitter. This affects the oscillators (usually quartz crystals) that control the speed of the local clocks leading to random fluctuations in their output.

Many things cause jitter, including: supply and substrate noise[108], temperature fluctuations, analogue signal jitter[109], and variations in the manufacturing process[110]. Jitter causes the local clocks to advance at different rates and become unsynchronised. Zamek, *et al.* [111] propose a more accurate method for estimating jitter in oscillators. Their research can be applied to increase the accuracy and stability of all time-based PLNs.

Conversely, Doherty, *et al.*[3] and Lisoweic, *et al.*[112] suggest that Loran-C and GPS respectively could be used to synchronise less stable oscillators such as those that would be found in a low-cost, quickly-deployable PLN. The concept of using such a system for synchronisation is not new; Doherty's group published their work in 1961.

Kim, *et al.*[113] discuss Binary Offset Carrier signal synchronisation as being an important step in recovering information from a global navigation system (they give Galileo and GPS as examples). Kim's group extends earlier work on reducing error from this synchronisation which could be used to increase the position accuracy of these, and other, navigation systems.

2.2.2 Synchronisation Difficulties

One of the major difficulties involved in synchronising components of a distributed network is the time it takes for a synchronisation signal to reach the outlying components. Consider a network of three nodes on a line, as in Figure 2.11.

Further consider the situation where node 1, the master clock, sends out a synchronising reset signal to nodes 2 and 3. A signal travelling at the speed of light will arrive at node 3 when node 2 shows 1ns and the master clock shows 2ns. A global accuracy of better than 2ns is clearly not possible when synchronising in this manner. However, it is the aim of this review to show that techniques exist to in-

crease the global accuracy in systems such as this.

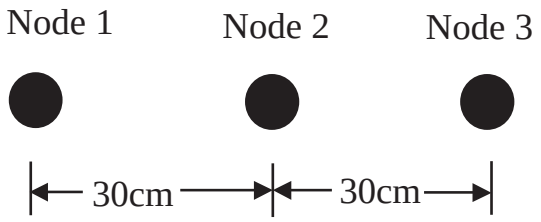

Figure 2.11: Three node, distributed network.

Gershenfeld, *et al.*[114] describe the advantages of having a bit size greater than the physical size of a network, where bit size is defined as the length of time required to transmit one data bit. Adapting this argument it can be said that, in general, if the required global synchronisation is shorter in time than the time it takes for a synchronisation signal to reach all of the nodes in the network, then synchronisation based purely on transmission of a reset signal is not effective, and it is necessary to use a more intelligent form of synchronisation.

Another major difficulty in synchronising networks to a high degree is the non-deterministic delay in instructions being executed at a node. This will affect the time it takes to implement a synchronisation operation and can result in nodes in the network being out of synchronisation with other nodes. Temperature, humidity, power supply fluctuations and manufacturing differences between the nodes all have a random effect on the speed of execution.

2.2.3 Reference Broadcast Synchronisation (RBS)

In this scheme a signal is broadcast to all nodes. The signal's time of arrival is used as a reference by all of the nodes in the network[115]. The local time of arrival of this reference at each node is retransmitted to all surrounding nodes. An average time is then computed at each node, which can be used as a local time reference for any operations that the network must perform. Since all of the nodes have the time of arrival information data from the whole network, the average value calculated is the same at each node. Figure 2.12 illustrates this.

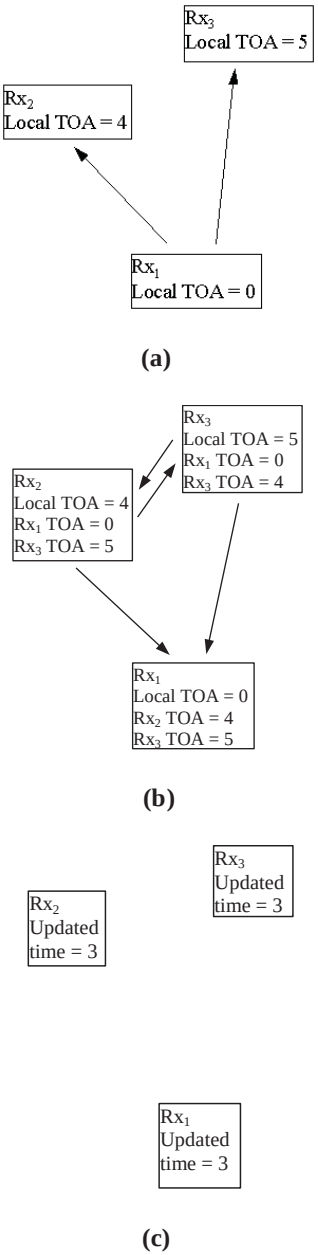

Figure 2.12: (a) Rx_1 sends a synchronization command. (b) Other stations transmit local TOA. (c) All clocks updated with the average signals.

RBS has been described here in an implementation for a single hop network. That is, a

network where each node can communicate with any other directly. The RBS method can also be used in networks where multiple retransmissions of a signal are required in order for the signal to reach all of the nodes in the network (referred to as a multi-hop network). RBS is implemented in multi-hop on a hop-by-hop basis, so long as each pair of nodes along the chain have a reference to which they can synchronise[115].

The major advantage of this RBS is that all of the nodes in the network are able to automatically come to a consensus about synchronisation. The nodes are also able to synchronise regardless of the distance between nodes. RBS can also be applied to mobile sensor networks so long as a consensus regarding synchronisation can be reached fast enough to ensure that the movement of the nodes relative to each other does not adversely affect the accuracy of synchronisation.

The limitation of RBS is that while events within the network will be synchronised they cannot be time stamped with a reference that has meaning outside the network.

Elson, *et al.*[116] describe a reference broadcast system that can be used in distributed wireless sensor networks that has an accuracy of approximately 2 microseconds.

Maróti, *et al*[117] propose a similar protocol called the Flooding Time Synchronization Protocol (FTSP), which has a reported accuracy of 1 microsecond which they claim is ideally suited to wireless sensor networks that have limited resources.

Palchaudhuri, *et al.*[118] present an adaptive synchronisation protocol that can be tuned to the specific requirement of the application and the resource constraints.

2.2.4 Measured Propagation Delay Synchronisation

This method employs a synchronisation command as described earlier. However, this method also includes some calculation to increase the accuracy of synchronisation.

Each node in the network must know its distance from the source of the synchronisation signal. When the node receives the signal it can then adjust its local time taking into account the time of flight of the signal over the known distance. The effectiveness of this method will be based on the accuracy of the measured distance.

The advantage of this method is that it is simple to implement. Once the distance to each node is known the propagation delay will be approximately constant. If the network is spread over a large area then other factors such as temperature, humidity and line of sight obstructions affect the propagation delay of the synchronisation signal. This variation may not be significant, depending on the synchronisation requirements and the distance between the nodes.

The major limitation of this method is that, unless the nodes are stationary, the distance of each node from the synchronising node must be recalculated before each synchronisation event to ensure that the correct offset is incorporated. This may be infeasible if the method of calculating the distance is done using a manual ranging method. This limitation will not affect the system if the overhead in recalculating the distance is low enough to allow the system to remain viable.

2.2.5 Atomic Clock

It has been stated that one of the major disadvantages of local oscillators is that they have a component of jitter, which causes the various clocks in a network to get out of synchronisation. One solution to improving this level of synchronisation is to use a more stable oscillator. One such oscillator is a cesium atomic

clock. These are significantly more stable than conventional quartz crystal oscillators and therefore stay in synchronisation longer.

Greater stability reduces the need for synchronisation events, thus synchronisation methods with higher overhead are acceptable.

However, currently available atomic clocks are too bulky and expensive to be used efficiently in distributed networks. Atomic clocks are priced between $50,000US and $100,000US each[119], but research is in progress to miniaturise these atomic clocks to fit on an integrated circuit. One such small atomic clock is pictured in Figure 2.13 and is the approximate size of a grain of rice[120]. This type of small-scale atomic clock is reportedly accurate to less than 0.5μs per day (compared with a quartz oscillator's 0.5s per day accuracy).

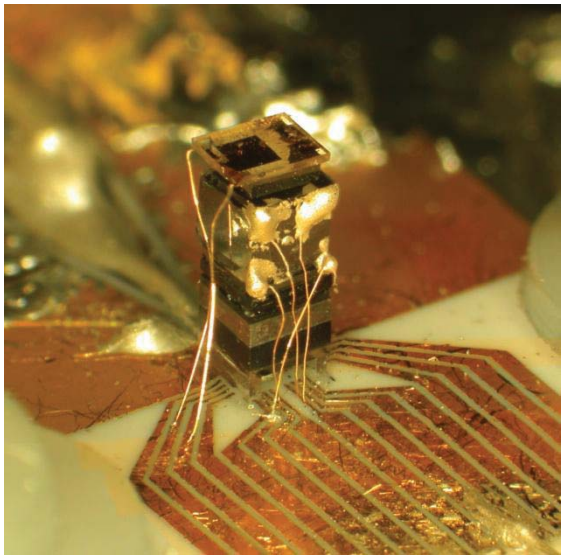

Figure 2.13: Atomic clock suitable for use at the circuit level.

Chip-scale atomic clocks are still some way from mass commercial use however, as the cost per unit is approximately $1500US[121].

2.2.6 GPS timing

There are several methods of GPS timing. One is to use the time signal present in GPS satellite transmissions to synchronise the node of a distributed network. Installing a GPS receiver in each node of the network does this; the receiver uses the redundant information in the satellite transmissions to synchronise the clock at the node with GPS time. The accuracy of such a system is reported to be on the order of 110ns for stationary, commercially-available GPS receivers[122].

The major advantage of this GPS timing is that the only cost involved is that of the receivers and integration; it is free to use any non-encrypted signals transmitted by the GPS satellites. The other advantage is that this system can be implemented anywhere in the world with the same equipment. It should be noted though, that the precision of GPS measurements is dependent on the receiver's distance from the equator. As the receiver moves away from the equator, the incident signals from the satellites appear from the same part of the sky, reducing precision (the effect of the spread of transmitters and receivers is called geometric dilution of precision and is fully explained in section 2.1.1).

The major limitations of using GPS signals to synchronise a network are that the GPS receiver will increase the size and cost of each node and each node must be able to receive a signal from the GPS satellites. Because of this, GPS time synchronisation may not be suitable for applications where the size and cost of each node must be kept to a minimum, and also, applications where the nodes are not in view of at least four satellites.

One of the major disadvantages of the GPS system itself is that corrections must be made for issues arising from relativity. The first issue is that the GPS satellite is moving with a high velocity relative to a base station, this causes the satellite's onboard clock to run approximately 7μs slower per day compared with a ground clock. The second issue is that the satellite experiences less gravity, which causes the clock to run approximately 45μs

faster each day compared with a ground clock. Therefore the total correction that must be applied to the clocks in the satellites is –38μs over the course of each day [123].

Zhao[87] discusses the advantages of including a position location function in 3G mobile equipment. With position data, there are many services that can be provided to the equipment and hence it is currently a very lucrative market.

Martin[124], Pflieger[125], Dierks[126] and Houlei[127] each describe GPS synchronized systems for taking measurements from mains power grids. These measurements are aimed at improving the efficiency and robustness of mains power systems.

D'Antona, *et al.*[56] suggest GPS timing for synchronising nodes placed around a lake. The method used here could be applied to any wireless sensor network and by extension, any PLN.

Lewandowski, *et al.*[128] propose a GPS timing system with an accuracy of 4 nanoseconds. The system uses measurements of the delay caused by the ionosphere to increase the accuracy of calculations. This work is aimed at achieving timing accuracy for events occurring at intercontinental distances.

2.2.7 Network timing protocol (NTP)

This protocol is often used in network servers to ensure that the time stamps on log entries are correct[129]. This is advantageous in the event of a security breach or power failure, as it allows accurate reconstruction of the events that led to the failure.

In NTP a server is coupled to an atomic time standard. This server is defined as the primary server or 'stratum 1'. This server is used to synchronise another level of computers, called 'stratum 2'. This can continue until stratum 16. Computers further down the hierarchy are less accurate compared with those higher up.

Figure 2.14[129] shows an example of this. Stratum 1 comprises several servers connected to atomic clocks. These servers can be used to synchronise either client computers (labelled as 'BX' and 'BY') or another layer of servers (labelled as 'B1' and 'B2'). This figure can be extended to the last stratum with synchronisation precision dropping with each stratum. Multiple strata of servers are used to allow a much greater number of client computers to synchronise (servers can only process so much traffic, limiting the number of clients that can connect and synchronise).

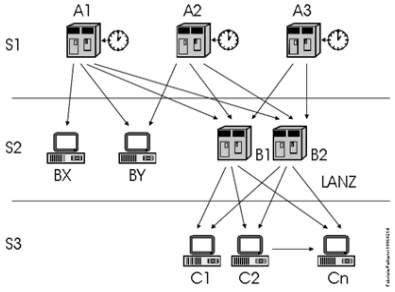

Figure 2.14: Network timing protocol hierarchy.

The actual synchronisation process involves multiple packets of data being sent back and forth between the computers. These packets are used to update the clock on the computer that is querying the NTP server. The synchronisation continues until the two clocks are in agreement.

The level of accuracy that can be attained using NTP depends heavily on the real time property of the operating system that is being used[130]. Approximations of the accuracy of clocks on machines running NTP software suggest that accuracy is of the order of milliseconds[131], however research is in progress to increase this level of accuracy, including that of Mills[132] who proposes changes to the Unix operating system kernel to realise submillisecond accuracy.

It is interesting to note that NTP does not aim to synchronise individual computers to one another but instead to synchronise all clients to the Universal Coordinated Time (UTC). Synchronisation of machines to one another is a serendipitous effect.

The advantage of NTP is that it can be used to synchronise the internal clocks of a large number of computers over a wide area. Each computer will automatically attempt to synchronise with an NTP server that is as close to the UTC as possible (defined by its stratum number).

The major limitation of NTP is the time required for the client's clock to converge with the server. NTP requires several hours to a day for the various clocks to synchronise with UTC, hence NTP is inappropriate in situations where nodes must be synchronised often (on the order of more than once in a week). However, Butner[133] proposes a system that can be used for timing nanosecond scale events that is able to converge to synchronisation in approximately one minute.

For further reading on NTP and its applications see Gurewitz[134] and Nakashima[135].

2.2.8 Ultra Wide Band Synchronisation

Ultra wide band (UWB) is a special class of radio communication system with some unique synchronisation requirements. A UWB transmitter works by sending pulses with duration on the order of hundreds of picoseconds to tens of nanoseconds. These transmissions occur at specific times. In order for the receiver to correctly receive one of these signals it must sample the input from the antenna at the specified time. Because of the short duration of the pulses, the nodes of a UWB network must have a high level of synchronisation so that the window can be kept to a minimum to reduce interference.

Successful implementations of UWB synchronisation schemes lead to communications systems that are very resilient to multi-path fading.

Yang[136] presents a blind method for time synchronisation that involves an integrate-and-dump operation over the duration of one symbol period (the time taken to fully transmit one symbol by whatever method) and exploits the multi-path diversity allowed by UWB transmissions. This method has a better performance when compared with other blind templates and is comparable to data-aided methods.

Zhang[137] describes a system called transmitted reference ultra wideband (TR-UWB) and also describes how the relatively poor bit error rate of TR-UWB can be improved while retaining the ability to resolve multi-path signals and timing attributes.

Tian[138] compares blind methods of timing acquisition with training based methods and proposes a new blind method that has an enhanced timing offset estimation that is comparable to most training based methods.

UWB is an interesting technology with respect to position tracking since the high degree of timing accuracy required for operation also allows for very precise event-time measurements. UWB is examined further in later chapters for application to our research questions.

2.2.9 Synchronisation Literature Conclusions

This section has highlighted both the need for synchronisation in distributed wireless sensor networks and some of the major difficulties in achieving it. Several methods have been described that cover a wide range of situations. The selection of technique and the accuracy that can be attained is heavily dependant on the equipment and implementation.

While in some systems it may be easier and advantageous to use asynchronous protocols, systems in which synchronisation is a requirement can be implemented by using one of these methods.

By using one or more of the methods described, synchronisation can be achieved in systems that are distributed in the metres to hundreds of metres range by using one of the transmitted reference methods. Synchronisation can be achieved in the kilometres to thouthousands of kilometres range by using either NTP or GPS timing synchronisation.

A comparative summary of these synchronisation methods is given in Table 2.2.

Table 2.2: Comparative summary of synchronisation methods.

Synchronisation method	**Precision**	**Range**	**Comments**
Clock reset signal	Dependent on relative ranges of nodes	Dependent on transmission hardware	
Reference broadcast synchronisation	Microseconds	Dependent on transmission hardware	Can be adapted for multi-hop operation
Measured propagation delay	Dependent on atmospheric characterisation	Dependent on transmission hardware	
Atomic clock	Nanoseconds		Stays stable longer than quartz-based oscillators so requires fewer synchronisation events
GPS timing	Hundreds of nanoseconds	Global	Requires GPS receiver in every device
NTP	Nanoseconds	Thousands of kilometres	Can require days to converge to synchronisation
UWB	Nanoseconds	Dependent on transmission hardware	

3 Methodology

In which the logical progression of the research is explained – testing procedures are described – results handling.

It is in the nature of this research that a number of disciplines must be drawn on in order to reach the final solution that answers the research questions.

This chapter will explore and explain the progression of the research including the logic behind decisions made and paths chosen. This chapter will follow the structure of the rest of the thesis so that the sections that follow will correspond to the following chapters. As befits an investigation of this nature, this chapter contains some results of subsequent chapters as these affected decisions for further investigations. The reader is asked to suspend analysis of these results until later.

The exact method and materials for each of the tests is explained fully in the relevant chapters and hence will not be discussed here.

3.1 Background

It was unclear at the start of the research what systems were available and specifically, what precision and accuracy could be achieved in a given position-tracking system. It is for this reason that it was decided to cover a wide range of positioning and synchronisation techniques in the background; the optimum could be selected from these.

3.2 Specifications

The definition of the specifications of the final system was crucial to this research. It was necessary to know quantifiable measures of system performance in terms of visibility, generating results in a timely manner and also physical constraints such as equipment volume, durability and the potential to harm players or interfere with the game.

The specifications defined the direction of the research. For example, a position-tracking system that requires three receivers in order to calculate a position is superior to a system that requires four receivers to calculate the same position.

The specifications are a major part of answering the research question 1: *'how best can a minimalist tracking system be implemented?'* since the specifications define the 'desired result' that needs to be achieved (see chapter 1).

The specifications are, in places, contradictory and require balancing in the final system. The shortfalls of these balanced specifications are the focus of research questions 2 and 3. Question 2: *'does an irregular configuration of receivers give superior results?'* is posed on the grounds that performance gains may be realised while using irregular system configurations, especially in conditions where there is a high probability of an object being obscured from the majority of receivers in the position-tracking network. Question 3: *'can a software model be developed as an expert system for position-tracking systems?'* addresses the issues with the maximised-performance and ease-of-use specifications. A properly developed software package will remove operator-training burden and the need for in-depth un-

derstanding of system theory as well as maximising the usability of the output data.

3.3 Tracking system selection

The initial idea in this research was to produce a tracking system from scratch using the most suitable equipment and techniques that could best fit all of the specifications (especially accuracy and equipment volume).

As research continued, it became clear that systems existed that could be manufactured to higher standards than an in-house system and with greater accuracy and long term cost.

It was also noted that there would not be enough time to scratch build a positioning system and perform the other research that would be dependent on that system working.

At the time this research was conducted there were few commercially-viable, position-tracking products being advertised hence, not a large number of products could be assessed.

Each system was checked for conformity with the specifications. The system with best conformity was chosen to move forward with research.

3.4 System characterisation

Characterising the system allows prediction of the behaviour of that system under different circumstances. The information from the characterisation is critical to defining a minimalist system.

The Multispectral Solutions Sapphire system was selected for the research (an active-RFID-tag based TDOA system) and the information on the behaviour of each of the components was required.

The tags, the receivers and full system operation were the major areas that needed characterisation and so, most of the tests concern these components.

The tag-transmission spectrum was measured and originally these results would have been used in conjunction with receiver gain data (obtained from the manufacturer) to estimate and model the visibility of tags to the tracking system. This test would have included an examination of how the signal is attenuated by a human body and how much, if any, signal diffracts around a body. Unfortunately, it was not possible to measure the tag-transmission spectrum to any significant range with the equipment available and on location at the testing site. Instead, this experiment was simplified to test how a human-body simulator (a bottle of water in this case) interacted with the transmission in a lab setting. A human-body simulator was used in place of an actual human to achieve consistency in the results (different body shapes may affect the signal differently) and ensure that the experiment is repeatable.

The tags were also mounted at different locations on the body of a player to find where the best signal reception would occur. It was expected that the human body would have a significant effect on the tag transmissions (evidence from the spectrum tests confirmed this) and so several locations on extremities were trialled. Trials on the torso were included for comparison and completeness. It was hypothesised that tag placement would be a critical factor in realising an effective minimalist system.

Given the failure to fully measure the tag spectrum at range it was decided to measure the ability of the system to receive signals with specific receiver/tag combinations. In other words, the range of each receiver type for a given tag type was measured. It was decided that this would be an acceptable test, as it would give quantifiable results that answered the pertinent question of range using equipment that would be in use in the final system.

Results are given for one tag type only, however, it is noted that a scale factor can be applied to convert ranges between tag types (e.g. micro tags have half the range of asset tags for the same receiver type).

Finally, initial trials were conducted with minimal receiver numbers to see if the system would be capable of tracking players in a rugby game with sufficient visibility. A three-receiver configuration was tried since this is the minimum number required to calculate a 2D position according to the manufacturer. A four-receiver configuration was trialled to see what effect that had on visibility. This was the maximum number of receivers available at the time and so no configurations with greater numbers of receivers were tried.

The decision was made in all of these experiments to keep outside influences out of the test results wherever possible. This means that only one variable at a time was changed (tag type, receiver type, tag placement location). Most importantly, interactions with the human body were limited by keeping the body away from the transmission or reception elements as much as possible (e.g. tags were held away from the body in the receiver lobe experiment). It is realised that by testing like this, the results will not match the situation in which the system will be used; however, it allowed factors affecting system performance to be gauged individually and an accurate cause-and-effect chain to be constructed. It was the intent that that the results of the characterisation would be used in models that added in the factors expected in a game situation (human-body-signal interaction most notably).

It was found after consultation with the system manufacturer that a characterisation with this detail did not exist; therefore this chapter contains significant new knowledge.

3.5 Player visibility modelling

The system characterisation suggested that there would be times that the tracking system would not be able to output a valid position for a player. This would be especially true for a minimalist system. It was recognised that one could take the brute-force approach to tracking where more receivers are added to increase the robustness of the system; however, this would go against the first research question of finding a suitable minimalist system.

It was decided that the first step in maximising visibility is to predict where and why visibility would be lost.

The player-visibility model was developed using data from the system characterisation and pre-tracked data from the Industrial Partner's manual system. The model was also to be the first part of the software expert system that the third research question requires.

It was important that real game data were used in the development of this model as this allowed it to be fine tuned for real game situations. This further allowed prediction of how well the system would perform in a game situation. This last part would become important for the optimisations.

The factors most affecting signal receptions were considered (as identified in the system characterisation). These were range, angle and obstruction. Each factor was implemented in the model based on information gained from the system characterisation.

The model assumed that if three receivers had visibility with a tag that position data would be output even if more receivers were in the network. Trials in the system characterisation showed this to be an accurate assumption.

As the name implies the visibility model calculates only player visibility and does not look at measurement accuracy (which will change, based on the GDOP of the receiver configuration with signal visibility). No accuracy measurements were included here, since the tracking system would output a position if the tag were visible to enough receivers, regardless of accuracy (hence the model represents the system).

It was recognised that the model could be used in a more general setting if there were the added functionality to check visibility in environments with static obstructions. For this reason, time was given to extending the model to deal with static obstructions.

The model was verified with a number of manually generated example situations to ensure that each block was working as intended. Finally the model was applied to the full game dataset in order to make inferences on how a tracking system will perform in a real-world situation, especially with regard to tracking with a minimalist system.

3.6 Receiver-placement optimisation

A key factor in answering the first research question was the fact that different receiver configurations would result in different system performance. Finding the configuration that yielded the highest performance was the aim of this chapter. Performance is defined here as a combination of accuracy and visibility where higher performance is at higher accuracy and/or visibility.

It was originally intended that the optimisation process would be performed for the rugby pitch on which the game would be played prior to configuring the system on site. Thus the optimisations would be used in the planning stage of position tracking. The optimisation would only need to be performed once per pitch so long as no changes were made to that pitch (changing field dimensions for example). It was further hypothesised that knowledge of how each team played could be used in the optimisation process (by way of the occupancy signature for each team). This was found to be inappropriate for this type of tracking.

It was seen from the results of the visibility model that changing the receiver positions would change the resulting visibility and it was also known from the literature review that the receiver positions would affect accuracy based on GDOP (i.e. receivers spread evenly around the object would produce the most accurate results and vice versa). In this way the optimisation process helped answer the second research question.

The problem of optimising the system configuration in this case is made difficult by the degrees of freedom (each receiver represents a new degree of freedom). This results in a total number of receiver-placement permutations that is too high to optimise by brute force. Furthermore, when optimising this problem there are a number of local minima (configurations that are optimal within given boundaries but are not the overall optimum configuration) which simple optimisation techniques have trouble handling.

Genetic algorithm (GA) optimisations were chosen as a tool to find the receiver configuration that gave the highest performance. A search of the literature and trials with a small number of other optimisation types revealed the suitability of GA optimisations for this type of problem.

It was hypothesised that the optimisation process could be used for any game situation, finding the optimum configuration of receivers for individual pitches (there is variation in the size of rugby pitches allowed in the game rules) and thus the optimisations would form the second part of the software expert system (research question 3).

A number of different objective functions were trialled in order to find one that produced useable results in acceptable times. A number of these were based on the visibility model and the real-game data; however, it was found that the optimisation process could not handle sufficient data in the required time to produce a stable result. Other objective functions used a combination of maximised coverage and minimised GDOP. Wherever possible, data from the system characterisation experiments was used (especially in calculating ranges of receivers and acceptance angle).

Analysis of a truncated solution space showed that the system was insensitive to receiver placement[2] and that most well-spaced receiver configurations would produce similar results in the rugby situation.

The useable rugby optimisations were compared using the player visibility model. Each solution's receiver configuration was passed as input to the model along with the same 3000-fame sample from the real-game dataset. The solution with the highest visibility is the fittest.

Given that it was found that the solution space for the rugby problem was flat and building on the static-obstruction-visibility model, the optimisation process was applied to a number of static-obstruction problems. It was hypothesised that these situations would represent a more complex problem to optimise. Objective functions were still a combination of maximised coverage and minimised GDOP. A similar examination of the solution space was conducted. The outputs of these optimisations were assessed on both coverage and GDOP and inferences are made.

[2] Insensitivity to receiver placement is not the same as insensitivity to receiver configuration. The system will output inaccurate data if the receivers are located at positions different from the configuration data. Note that 'placement' refers to the receiver's physical location where 'configuration' refers to the stored coordinates of the receivers in the hub processor (i.e. where the receiver 'actually' is and where the hub processor 'thinks' it is).

These optimisations were helpful in showing that irregular receiver configurations yield better tracking results. These optimisations gave conclusive results for the second research question.

3.7 Position Reconstruction

It was established in the previous two chapters that some level of invisibility was inevitable in the tracking system. This was especially true when dealing with rugby, since there were some formations in which perfect tracking is almost impossible (rucks for example).

It was therefore decided to implement a reconstruction algorithm to recover as many of these missing data as possible by using the data present. This algorithm was demanded by the third research question but helped to realise a minimalist system from the first research question. In other words, if positions would be lost, then instead of adding more receivers to improve coverage, an algorithm could infer position from other data.

The algorithm was developed in a functional block architecture that allowed different recovery techniques to be added and evaluated as necessary. Each technique was applied in order of descending effectiveness (i.e. the most effective techniques that returned the most data were applied first).

The blocks were split between two broad categories: those that could be applied to tracking systems in general and those that could only be applied to a rugby game.

The algorithm was applied to selected frames from the data collected from field trials and pre-tracked game data and the effectiveness of each block was evaluated. The input frames were assessed for visibility using the model described in chapter 7 and were known from visual inspection to contain a given situation (e.g. ruck or scrum). The algorithm was passed 10 frames that contain at least one in-

stance of the test situation. In order for the algorithm to be successful it must correctly identify the frame or frames in which the situation occurs and correct for it with sufficient accuracy. The evaluation was conducted with ten instances for each situation. Statistical analysis was performed on the results.

3.8 Simulated tracking-system performance

Over the course of this research the relationship with the Industrial Partner deteriorated. They were under the impression that the purchase of the test system was the end of the research and that they would be in possession of a fully working tracking system that would suit their needs; this was false.

Also, the insurance policy that was taken out on the equipment required that a staff member from the Industrial Partner be present whenever testing in the field was undertaken. This severely limited access to the equipment and held back research.

Finally, when the Industrial Partner did not see performance as high as they expected from initial trials they sent a delegation to Multispectral Solutions to discuss the system. There they were told about the way the system is used for other tracking applications (such as personnel in an office building) and they became convinced that the tracking system could not be applied to tracking rugby with suitable performance (despite evidence from the system characterisation experiments that showed that it could).

Because of all of this, access to the system for a full game trial was denied. Instead it was decided to pool the results from all of the experiments and apply them to the game-data suite. This involved configuring the visibility model to use the calculated optimum receiver positions and then applying the model to the entire set of pre-tracked game data. This resulted in a dataset where some of the player positions were invisible. The reconstruction algorithm could then be applied to attempt to recover those missing positions (based solely on nearby visible positions). The reconstructed positions could then be compared with the known positions from the game data-suite to check reconstruction accuracy. By doing this, the results from the system characterisation, visibility modelling, receiver placement optimisation and position reconstruction are used to simulate how the system would work in practice.

Results were analysed based on volume of recovered data and the accuracy of those data. Evaluations of the simulated system were made based on the specifications and the final definition of a minimalist system is made with respect to the system specifications.

While it is realised that a test of the actual tracking system in a game situation is an important verification of the results, the simulated system provides a suitable indication of how the real system would perform. The use of real data from pre-tracked rugby games and the system characterisation (erring on the conservative side) further raises confidence in these results.

The simulated tracking-system performance chapter ties together each of the three research questions in a single system. By doing this, the relationships between the research questions can be gauged and balance between specifications and the research questions can be found.

The results are analysed in several ways. The visibility is measured, specifically how it changes with different numbers of receivers (6, 8 and 10-receiver configurations are used, positions of which are taken from the optimisation data). The reconstruction algorithm is assessed based on how many positions it recovers (and hence how the visibility changes) and how accurately it can recover positions. Inference will also be made on the how effectively the reconstruction algo-

rithm can identify and correct missing positions in application-specific situations.

The analysis performed here will be used to justify the final minimalist solution based on the requirements of the specifications.

3.9 Repetition

A number of different experiments were performed in this research. Each experiment has sufficient repetition to ensure consistent and repeatable results.

In the case of physical experiments, repetitions are performed to ensure that outside or unconsidered influences do not affect results and to account for inaccuracies in measuring results.

In the case of experiments involving only software, repetitions were either performed with different input data (e.g. receiver placement optimisations used a different subset of the game-data suite) or it was accepted that repetition using the same data would not produce different results (e.g. the simulated system performance test would produce the same output for the same receiver configuration when applied to the game-data suite).

The number of repeats is stated and justified in the method for each experiment. Software simulations with no random elements are not repeated (however they are verified using several test cases) since the results would not change.

4 Specifications

Qualitative and quantitative requirements are discussed – issues and constraints are presented – initial specifications are reached.

Successful design projects share the common characteristic that the requirements are captured early in a series of specifications that are comprehensive, clear and testable. To this end, consultation was undertaken with the Industrial Partner on the requirements of a system suitable for real-time tracking of player positions in sports matches. These specifications will be used to gauge the requirements of a minimalist system.

From the beginning, it was decided that the player-positioning system need only track the positions of the players and perform any filtering or processing so as to improve the accuracy and availability of this information. The Industrial Partner, if they wished to do so, could then apply additional information and perform statistical modelling on the data.

The specifications that follow are the result of the consultation process and helped to define the direction of the research. They are presented here in descending order of importance.

4.1 No physical interference

The equipment in the system must not interfere significantly with the play of the game. In an ideal system, neither the players nor the spectators would be able to detect the presence of the equipment.

The characteristics of any tracking equipment that is attached to the players can be specified and tested in the following manner:

- It must be incapable of injuring the player during a collision or fall. Impact tests will highlight any possibility of injury. The equipment must not mark or bruise the player in any way in order to be acceptable. Padding may be added to further decrease the injury probability.

- Tracking equipment must be placed where it cannot be knocked off during a game, as there will be no opportunity for the operator to reattach devices once play has commenced. Testing of the system under game conditions will indicate fulfilment of this specification. The loss of no more than two devices per half game is considered acceptable.

The characteristics of tracking equipment that is stationary around the field of play can be specified in the following manner:

- It should cause minimal interference with the game; hence any receivers should be placed outside the tracked area and also out of the way of ancillary game staff (linesmen, camera operators, etc). An exclusion boundary can be used to successfully implement this specification. The ancillary staff are usually within 5m of the sidelines and thus no receivers may be placed closer than that.

- The visual footprint of any receivers must be small enough that the spectators' view of the game is not affected. In the ideal case the receivers would be small, be outside the

spectators' line of sight. Successful implementation of this specification can be gauged by the angle subtended by the equipment on the spectator's vision. This angle should be less than 5° to fulfil this specification.

4.2 Durability

Any required tracking devices must be able to handle the impacts and shocks that a player will encounter during the course of the game. Devices must operate correctly for the whole game, regardless of what happens to them. The devices will be placed on the players before they take the field for the first time. Once this happens there will be no opportunity to replace or adjust the devices. Therefore, the devices must be constructed to resist shock and be placed in an area where shock is minimised.

Equipment placed around the field of play must also continue to operate correctly following any collision. Examples of such collisions include ancillary game staff accidentally walking into the equipment or a ball striking the equipment when it is kicked out of the field of interest.

In the short term this specification can be tested by putting the system in a game-like situation and examining both the physical equipment and the data for evidence of damage or loss. If after these tests the equipment and the data are intact then this specification has been met.

In the long term this specification can only be gauged after a large number of game situations. One season would be sufficient to see how durable the system is.

4.3 Portability

This system is not intended to be a permanent installation mounted in a stadium; instead, it is to be portable and travel with a team to different playing venues. As such, the entire system must be able to be stored and transported as part of the team's luggage and should require minimal set-up and calibration time.

A corollary of this specification is that the volume of equipment that must be transported, stored and set up is kept to a minimum.

An ancillary interpretation of the portability requirement is that, with minimal modifications, the system should be applicable to live tracking in different sports. While the primary application of this system will be in the sport of rugby, the system should also be applicable to sports like soccer, netball and basketball.

Success in meeting this specification will be determined by the volume and mass of equipment required to operate the system and a subjective assessment of the complexity of the configuration process.

An entire tracking system (core system, power supplies and ancillary equipment) that can be fitted into a single transport case (such as the Pelican 1780T transport case) would be ideal. Such cases have an internal volume of approximately $0.02m^2$. Any system requiring more than three such cases (total volume of $0.06m^2$) for transport would be unacceptable. Ideally the total mass of the equipment should be less than 30kg. A system will be judged portable if it can be configured in less than 1 hour.

Examination of the modifications required to operate the tracking system in a sport other than rugby will give an indication of the system's portability between applications.

4.4 Maximised operation time

The system must be able to operate for an entire game without attention to the receivers or transmitters, e.g. battery charging.

Fulfilment of this specification must take into account the various additional events that can occur beyond the regular duration of play (80 minutes for a game of rugby). These events in-

clude: the singing of national anthems, injury time, stoppages of play, tiebreaker over times and halftime.

Examination of the power consumption of any battery-powered equipment in the system will give an indication of the expected endurance of the system.

Tests of run time under game conditions will be the final evaluation on this specification. The system will be capable of tracking continuously for at least 3 hours on one battery charge in the ideal case. Continuous tracking for between 2 and 3 hours will be acceptable but not ideal. Less than 2 hours of continuous tracking capability will be unacceptable.

4.5 Minimal cost

The system should be low enough in cost such that a professional team or franchise can afford to own its own tracking system. Competitiveness in the market place will require that the cost of the system be on a similar level to any other form of tracking system currently available.

The cost of extending the operational time of any tracking devices will have significant impact on the system cost specification, as these will be the most numerous components in the system. The cost of replacing or recharging either the devices or their batteries should also be minimal.

Volume production of the tracking system will impact directly on the total cost. Therefore the ability to use commercial-off-the-shelf (COTS) components wherever possible will help to minimise the overall cost of the final design.

Consultation with the Industrial Partner yielded an upper cost limit of $100,000NZ. Any system that costs more than this will be unacceptable.

4.6 Measurement accuracy

The system must produce an output of the positions of each object to be tracked at approximately one-second intervals with accuracy better than 0.5m in any direction on the field. These positions need only be returned in two dimensions, but the ability to return positions in three with elevation, would be advantageous for further analysis. The Industrial Partner indicated that positions in three dimensions would be desirable but is not required. Furthermore, it is considered that data that are missing are inaccurate and thus the system must have less than 20% missing data in the output dataset.

Simulations can give an indication of the errors in positioning that can be expected from different position-tracking systems, but this specification will be evaluated by taking position measurements at known positions on the field. Simulations will also be used to optimise the configuration of the tracking system to maximise its performance. Furthermore, analysis of raw data from a particular tracking system can be used to increase the accuracy of measurements and to recover data that are missing.

4.7 Ease of use

A person with little or no experience in the workings of the system and with minimal training should be able to operate the system correctly during a game. The system must be simple enough that one person can set up, activate and have the system operate reliably and correctly by following a simple list of instructions.

Evaluation of a tracking system meeting this specification will involve an assessment of how well an operator with minimal training and a list of instructions can set up and configure the system. An operator with a list of instructions should be able to configure the system for use in less than 1 hour.

4.8 Timely availability of data

The time taken for the system to produce data suitable for statistical analysis must be significantly shorter than that of the Industrial Partner's current, manual, video tracking system. An initial specification is that data should be available for statistical analysis within 1 hour of the end of the game. However, if the delay in the availability of data were shorter, this would be advantageous.

The system should also require minimal human intervention. A particular example of this is that, if an object is lost to the tracking system for a length of time, then human intervention is not required when the object reappears.

The elapsed time between the end of a game and the data being available in a format suitable for statistical analysis will be the test for this specification.

4.9 Issues and constraints

Beyond these specifications there are a number of constraints imposed on the system by the physical dimensions of the playing field.

A rugby field is generally 120mx70m, including the area behind the goal posts (Figure 4.1). There are situations where these dimensions can be altered, but for the purposes of this research it will be assumed that these dimensions are fixed. The area outside this field of play is available for the placement of equipment; however, this placement may be further constrained by the characteristics of the system and obstructions around the field of play.

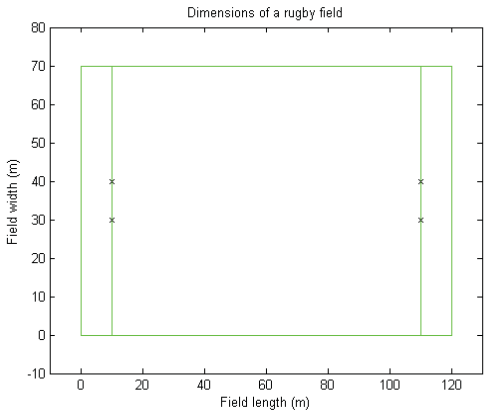

Figure 4.1: Boundaries of a Rugby pitch. Crosses denote goal posts.

It is likely there will be conflict between the various specifications of this system that will require compromise. Foremost of these is cost, since in almost every case the cost increases when any particular system characteristic is improved.

Other specifications that may be in conflict include:

- Accuracy of measurement and portability: Increasing accuracy will most likely require more receivers; hence the portability of the system will be reduced.
- Non-interference and maximised operation time: The batteries required to attain the required endurance may be large, thus increasing the visual footprint of the tracking equipment.
- Accuracy of measurement and durability: Preventing any tracking devices from being damaged may necessitate their being placed in a location that is not conducive to accurate measurements.

The conflicts of characteristics and the resulting compromises will be explained in the text as they arise.

5 Development of a prototype tracking system

The road from in-house development to COTS – Nordic – Paric – Multispectral Solutions – review and refinement of requirements – ultra wideband communications – compromise of requirements.

At the beginning of the research it was intended that the tracking system would be developed entirely in-house. As the literature review progressed and the specifications of the system were refined, it became clear that an in-house system would not be suitable, both from the perspective of its performance and marketability. The literature review (chapter 2) suggested that significant work had been done in the area of position tracking and it was reasonable to assume that a commercial-off-the-shelf (COTS) system existed that could be modified to suit the requirements of the research. A COTS system was also required to meet the minimal-cost requirement of the system (see section 4.5).

In chapter 2 it was shown that the best system would be time-based with unique ID's transmitted from tags mounted to the players and thus the research focused on locating such a system and identifying the modifications that would be required to construct a position-tracking system suitable for the game of rugby. Three companies were identified as having potential solutions: Nordic Semiconductors[139], Paric Measurement[140] and Multispectral Solutions[141].

5.1 Nordic Semiconductors

This company, located in Norway, produces a number of consumer RF transceiver systems with integrated micro controllers. A development kit of the 2.4GHz transceiver was obtained for evaluation purposes. A similar system was requested from the South Pacific agent for Ubisense (a British company that produce a similar line of products); however, they did not respond to inquiries.

A significant advantage offered by the Nordic development kit was the flexibility of design. The operation of each of the units was totally under the control of the developers. Also, with the addition of the appropriate antennae, one set of hardware could be used for both the player tags and the receivers.

The transceivers measured approximately 50x25x10mm as an exposed circuit board with 2.4GHz short dipole antenna, which would have made them suitable for use as player tags with an appropriately padded case.

Given that the Nordic units were transceivers, several architectures were available when implementing a position-tracking system. One that was investigated was a poll-and-response architecture, in which a master unit sends out a request for a particular player's position. The tag on that player responds immediately and several receivers around the field calculate a position based on the time characteristics of this response. This method has inherent inaccuracies due to non-deterministic latency in the signals. Also, poll-and-response would not scale well; increasing numbers of players to track would degrade the system's ability to produce positions of all players in a suitably timely manner. For these reasons, the tracking system was not implemented in this fashion.

A second architecture involves the transceivers on the players periodically transmitting data. The contents of the transmission would uniquely identify a particular player and a position could be calculated (again, using the time characteristics of the signal). Transmission on a pseudo-random basis would help reduce the occurrence of signal collisions that would impair the performance of the system.

In either architecture, the biggest issue would be calculating accurately the time of arrival of transmitted signals at the receivers, as this would greatly affect the accuracy of any position measurement. Upon being received, a transmission must be decoded and validated, and then a time stamp applied (accurate time stamping would require a high level of time synchronisation across the receiver network). These actions are time critical. They must be completed as quickly as possible to avoid the non-deterministic delays from program execution.

The contents of the transmitted data packets can be fully controlled using the Nordic hardware. Some form of identification of a particular player would be required; however, with appropriate transducers, it would be possible to transmit some biometric data from the player (heart rate for example). These data could further increase the information available to teams.

The flexibility of these transceivers also led to their unsuitability. In order to produce a positioning system as defined by the system specifications, the architecture would have to be programmed from the ground up. This would have taken most of the research period. Furthermore, the literature review suggested that complete positioning systems already existed and so development with these transceivers would have reproduced existing knowledge rather than extending it.

Finally, a system developed based on this technology would not be easily deployable by the Industrial Partner in the marketplace without significant product development.

Nordic, and all other programmable tags, is an example of a COTS-Hardware solution in which the economies of scale from mass producing the miniaturised hardware are gained (thus reducing the cost and increasing quality) but the software could be tuned to sporting applications in-house. In practice however, while the hardware is capable of meeting the requirements for functionality and cost, the software would need to be written from scratch and would merely replicate the work of others.

Once again, the literature review indicated that a COTS-Hardware solution would detract from research and thus the decision was made to abandon research down this route in favour of a COTS position-tracking system.

5.2 Paric Measurement

This company is located in Auckland, New Zealand, and advertises a time-based positioning system for industrial indoor usage called the P^3[140]. On the face of it, this was exactly the kind of system that the Industrial Partner required.

A potentially significant advantage of this system was local production. Because Paric was based in New Zealand, it was a simple matter to have discussions with the company regarding the suitability of their technology for this application.

The operating principle of the system well matched the system requirements. A number of receivers are placed around the area of interest and receive signals from tags within the area. Position calculations are made from the time characteristics of the signals.

This system proved unsuitable for several reasons. While the system could produce positions accurate to within 0.3m of the true position of the object, the range of the system was typically

only 50m; this was too low to provide adequate coverage of a rugby field. Secondly, the transmitter tags measured 80x50x15mm which is slightly too large to be comfortably placed on players. Finally, Paric suggested that for their system to operate to the accuracy that was required then at least 12 receivers would be needed. At the time this seemed an excessive number that the Industrial Partner would not accept (the volume of equipment was too high, see section 4.3).

5.3 Multispectral Solutions Inc

An examination of the manufacturers of UWB positioning systems in the early months of 2006 showed that Multispectral Solutions Inc[141], was one of the few companies producing a complete tracking system using UWB technology. They appeared to have suitable hardware that matched the project requirements.

In the same way that the Industrial Partner is a leader in the field of analysis, Multispectral appeared to be a leader in the field of commercial, high-accuracy, time-based positioning. The company had been in business since 1988 and had developed products for several UWB applications. These applications included high data rate communications, UWB obstacle avoidance radar, asset tracking, inventory control and intelligent transportation systems.

This project falls under the Multispectral Solutions category of precision asset-location systems for which they had two solutions of interest: the Sapphire Dart and the Sapphire Dimension. Both were based on the same concept of a time-based tracking system where the tracked objects are tagged with UWB transmitters which send signals to receivers placed around the field of interest. The range of the system was quoted to be in excess of 200m in ideal conditions. A two-dimensional position is output if the transmitted signal is visible to three or more receivers, supporting the requirement for portability (section 4.3).

The major components of the Sapphire Dart system are shown in Figure 5.1[141]. Included in this picture are the receiver units (top), the hub processor (blue box, middle) and the tags (foreground, left to right: badge, micro, asset).

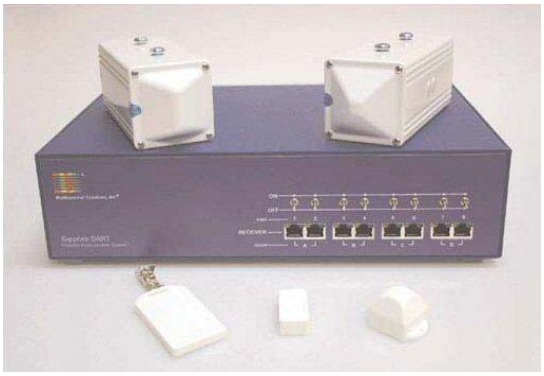

Figure 5.1: Sapphire Dart precision asset location system.

The position calculation side of the system conforms to the theories of time-based tracking explained in chapter 2. At least three receivers (but usually more) are placed around the area of interest. The locations of these receivers are provided to the hub processor as part of the configuration. The configuration page of the web interface is shown in Figure 5.2

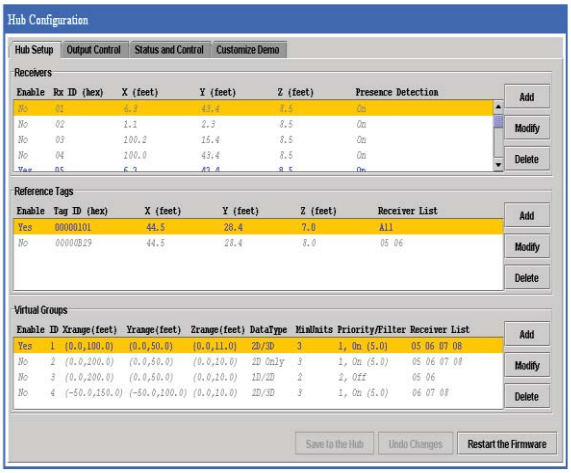

Figure 5.2: Configuration page of the Multispectral tracking system

Three major windows can be seen here. The topmost window is where the positions and

ID's of the receivers are configured, the middle window is where data about the reference tag are entered and the final window is where virtual groups are configured. The configuration of the receivers and the reference tag are self-explanatory (the use of reference tags will be explained shortly).

Virtual groups define an operational subset of all of the receivers. Positions are calculated using data only from a particular virtual group. In this way it is possible to prevent the system from outputting a position that has been calculated from signals arriving at receivers that cannot possibly have received valid data. An example of this would be TOF data from one tag appearing at two receivers when a large metal wall separates those receivers; the data at the receivers cannot have come from a valid line-of-sight path and can therefore be discarded as invalid data.

Each of the receivers is connected to the hub processor via CAT 5E shielded network cable. The receivers may be connected to the hub in a star configuration or daisy chained to each other, or some combination. Figure 5.3 is taken from the Dart manual and illustrates the connection possibilities.

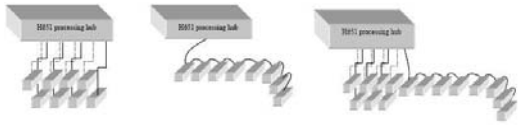

Figure 5.3: Receivers may be connected in a star, daisy chain or combination configuration.

The method of synchronising the receivers in this system is twofold. The hub is able to synchronise the receivers with signals sent through the network cables. Also included in the system is the requirement for there to be a 'reference tag' present. This tag is simply a transmitter tag placed at a known location that is in range of all of the receivers. There is provision for using multiple reference tags if the system is spread over a wide area where using only one reference tag is not feasible. The receivers are able to use TOF data from the reference tag combined with knowledge of its position to correct for small clock drifts.

The Sapphire Dart was used in the initial trials (see chapter 6) for the system because it was the only system that Multispectral had for sale at the time. The Dart system was deemed unacceptable for the final implementation of the tracking system because of the requirement for the receivers to be connected via wires to the hub processor. These wires presented an unacceptable level of interference to the game (see section 4.1). An enquiry to Multispectral about a wireless system that had the same principle of operation identified that a suitable system called the Sapphire Dimension was to be released early in 2007.

The major difference between the Dart and the Dimension was the ability for the Dimension to work using standard network components such as network switches and routers. From the perspective of the project this meant that the receivers could communicate with the hub processor over a wireless connection. If batteries were used to power the receivers and wireless-bridge equipment, then the receiver units could be physically independent of the hub processor.

The reported accuracy for both of the systems is 0.3m with no averaging and better than 0.3m with averaging (compare this with the required accuracy stated in section 4.6). The output of the system is designed to allow for a significant amount of post processing. This was within the bounds of the requirements of the project.

After negotiating with the directors of Multispectral, we were able to lease the Dart system for initial testing of the equipment to determine the suitability of the technology for tracking rugby players. Later the Dimension system was purchased.

5.3.1 Tag hardware requirements

Several different tags were available, each with different case dimensions. The difference between the tags was in the type of antenna and hence the radiation pattern of the transmissions. Selection of the correct tag for the application was therefore essential. The evaluation of the hardware is described in chapter 6.

Four tag types were available. Table 5.1 gives the physical characteristics of each (these data relate to the physical interference requirement in section 4.1). It should be noted that each tag has a default transmission repetition rate of 1Hz and that at this rate the batteries in the tag are said to last for 4.3 years. With this battery life, a single group of tags can be used over many games, thus the operating time is well in excess of the specifications and the cost can be minimised since it is not necessary to purchase new tags every game (see sections 4.4 and 4.5). The battery used is a lithium button cell, similar to what one would find in a wristwatch. This battery is easily replaceable.

Table 5.1: Physical characteristics of Multispectral tracking tags.

Tag name	Dimensions (mm)	Weight (g)
Badge	42.42 x 67.82 x 7.11	16.6
Asset	24.46 x 24.46 x 25.4	12.8
Micro	12.7 x 38.1 x 19.05	8.4
Player	44.7 x 32.26 x 6.35	10.3

The tags transmit on a pseudo-random basis. Recall that this is preferable to poll-and-response architecture, as it removes the need for both transmit and receive capability in the tags and thus it reduces size, weight, complexity and cost of the tags. The tags are small enough to comply with the non-interference specification. They also come with hard plastic cases that, with padding, will prevent either the tags or the players from harm.

The measurement accuracy specification requires that, for each object, a position be calculated at least once every second, suggesting that the transmission rate of 1Hz will be sufficient to meet this specification. Greater rates would be useful to increase the average measurement accuracy[3] (see section 4.6). It should be noted however, that increasing the transmission rate would increase the probability of signal collisions when the tags transmit. UWB transmitters naturally avoid signal collisions by transmitting in extremely short pulses (on the order of tens of picoseconds), resulting in a low probability that two or more transmitters will be active simultaneously. However, there is a limit; increasing the transmission rate will increase the data traffic in the positioning network, eventually causing a catastrophic number of collisions so that the system will not function.

The amount of traffic the Multispectral systems can handle is also dependent on the number of receivers in the network. The receivers communicate with the hub processor using Ethernet protocols and thus, each receiver will increase the traffic in the network and again, there is a volume of traffic beyond which the reliability of the network will fail. Multispectral state that their systems can handle 5000 tag transmissions per second with four receivers[142]. No figures are given on how the system's capacity changes with different numbers of receivers.

As an example of a calculation of the highest tag rate that is useable, if it is assumed that the relationship is linear between receiver numbers and total tag transmission rate (product of individual tag transmission rate and number of tags) and if it is assumed that 50 objects need to be tracked (2 teams of 15 on-field players plus 7 substitutes, the referee and provision for 5 other objects to be tracked) then with 8 receivers the theoretical maximum individual tag transmis-

[3] The transmission rate could be changed at the factory on request. 5, 25 and 50Hz configurations were offered. There was no ability for the user to alter the transmission rate.

sion rate would be 50Hz. So long as the tag rate is less than or equal to this the system performance should not degrade due to too much traffic.

A test was performed with this system to gauge the nature of this failure. 6 receivers were setup and 75 50Hz transmitters (calculated to be the maximum total tag transmission rate that the system can handle) were activated one by one. It was noted that as the number of active tags approached 75 that some of the earlier tags began disappearing from the display. When some of the tags were deactivated the missing tags reappeared. This suggests that the system had reached the maximum network traffic it could handle; this might have been due to either packet collisions or some form of buffer overflow internal to the server. This test also showed that the system will not crash with too much traffic but the performance will drop (i.e. some tags will become invisible).

The tags conform to the ease-of-use specification as they become fully operational once switched on and continue to operate until either switched off or the battery goes flat (see section 4.7).

5.3.2 Receiver hardware requirements

The Sapphire receivers are 165.2x63.5x63.5mm and weigh around 425g each, which compares favourably with both the requirement that equipment does not impinge on the view of the spectators (see section 4.1) and the requirement that both the mass and volume of the equipment be minimised (see section 4.3). Three different antenna types are available on the receivers; these are known as omni-directional, mid-gain, and high-gain. Each antenna type trades range for acceptance angle (See chapter 6 for the analysis of the range and acceptance angles of each).

Loss of line-of-sight visibility prevents a position from being calculated. Similarly, multipath interference can reduce the effectiveness of the system. Avoiding both without increasing the number of receivers will be required to make the system feasible. One way in which to achieve this is to mount the receivers high enough that tag visibility is maximised but not so high that the portability of the system is impaired. It has been suggested that mounting the receivers at a height between 2 and 3m would be an acceptable trade off between overall height and portability.

Several stand types were considered for receiver mounting. These types included camera tripods, microphone stands and speaker stands. In-house constructed stands were considered briefly but were discounted on the grounds that a suitable stand would exist as a COTS product. All types other than speaker stands were discounted based on their stability when loaded and the height they would be able to lift a receiver.

A speaker stand would be suitable for the purpose of mounting the receivers. Speaker stands are usually used to mount large speakers at concerts. The stands are designed to raise heavy speakers to heights above 2m, yet be stable enough to do this safely. An example of such a stand is pictured in Figure 5.4.

Figure 5.4: Speaker stand.

The particular stand pictured here is a Proel FREE300 Speaker Stand. It has a maximum load capacity of 70kg, weighs 3.3kg, can be ex-

tended up to 2.18m and the tripod base can be extended out to 1.1m[143]. The load capacity of this stand is far higher than the expected weight of the receiver units (receivers and power supply) however this is considered advantageous in this situation as it will lead to increased stability in high winds. This increase leads to greater reliability in the system (high wind will not remove a functional receiver from the system) and protection of investment (the receivers are less likely to be damaged).

Another requirement of the receiver hardware from the perspective of the project is that once the system is set up, it should be resilient enough to continue to operate for the duration of the data-recording process (see section 4.2). While it has been noted that the receivers will not be placed in the path of any of the players or the game support staff, there may still be times when the receivers are knocked. Reasons for a knock could include collisions from people moving around the periphery of the stadium, a strike from a kicked ball or high winds in the case of games played in inclement weather. In all situations the receivers must remain stable and active. It is unlikely that the operators of the tracking system will be able to correct a receiver should it be downed for any reason. The use of speaker stands, like the ones described above will be useful here also. Because these stands are normally used to support large and heavy speakers, they have been designed to be very stable at the base. In most cases it would be sufficient to simply extend the tripod legs to their maximum and place them on the ground. However, for extra security, the legs could be pegged to the ground so that in the event that the stand is struck, the receiver will be held in place.

5.3.3 Ultra wide band

A recurring theme from the discussions of the system is that in almost all cases, the effects of multipath interference and visibility issues (maximised measurement accuracy) will have serious detrimental effects on the performance of the system. Also, in a time-based system, increasing the time resolution will increase position-measurement accuracy. The UWB transmission scheme employed by the Sapphire system attempts to circumvent these issues. UWB is a relatively new technology and so a brief description is given here.

Pulses of extremely short duration – on the order of tens of picoseconds to nanoseconds – are used to generate UWB transmissions. These pulses are driven directly into the antenna without modulation of a carrier signal. For this reason, UWB is sometimes referred to as carrier-less transmission. Fourier transformation and analysis shows that any signal that is of short duration in the time domain will have a wide bandwidth in the frequency domain, and vice versa. This means that narrow pulses have a wide bandwidth, the shorter the duration of the pulse, the wider the frequency bandwidth of the signal. Figure 5.5 gives an example of a UWB transmission in both the time and frequency-domain.

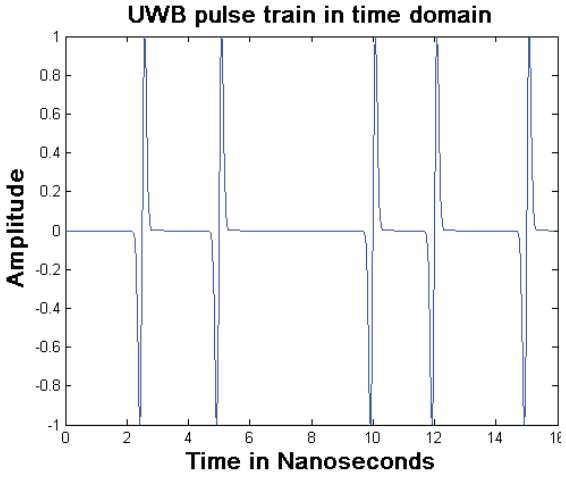

(a)

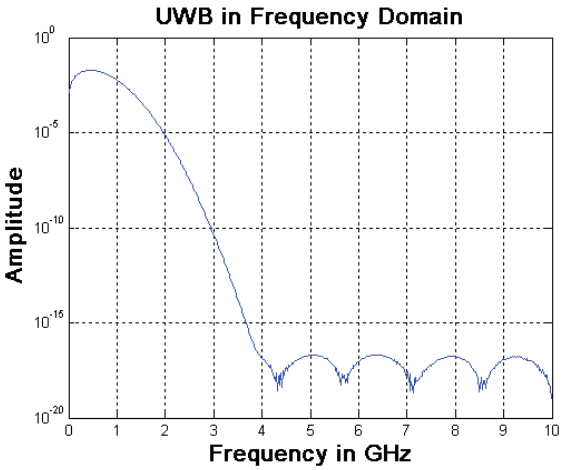

(b)

Figure 5.5: A UWB signal in both (a) time domain and (b) frequency domain[144].

UWB systems avoid interference with other transmission systems by remaining below the noise floor of conventional, narrow-band systems.

Because of the short duration of the pulses in the time domain, both the transmitter and the receiver must have very accurate timing components to facilitate communications. This facility for high accuracy timing can be used by a time-based positioning system to increase the resolution of measurements, hence increasing the positional accuracy of the system. Short duration in time also gives an advantage in terms of multipath interference. In UWB, the signal duration is short enough that very few of the multipath components will affect the reception of the signal. The ability to resolve multipath signals can even be used to increase the strength of the signal[145], however, only the direct-path signal is of use in positioning; multipath components will introduce error into the position calculations.

There are currently several competing UWB standards, but most apply to short-range, high-data-rate communications. One standard that covers the use of UWB for positioning is IEEE standard 802.15.4a[146]. This standard calls for a lower data rate and a higher range to allow the tracking of objects over several hundreds of meters. These tracking systems must not interfere with any other transmissions that use the same spectrum.

Currently, no one standard has been adopted worldwide. As of January 2009, The United States of America is the only country that allows the use of UWB under a general-purpose licence. Much like the use of IEEE 802.11 equipment, this general-purpose licence means that anyone may use UWB equipment in America so long as that equipment conforms to the regulations set down by the Federal Communications Commission (FCC), America's spectrum management body. The European Union has elected to follow the American standard for UWB[147]. It is likely, given that these two major economic powers have adopted the same standard, that New Zealand will follow, but this has not happened yet.

For the purposes of testing and using UWB equipment in New Zealand it is currently necessary to obtain a licence. Radio Spectrum Management (RSM), the regulatory authority on spectrum usage in New Zealand, issued a licence for this project in August 2006. To their knowledge, this project is the first and only to employ UWB for the purpose of position tracking in New Zealand. The licence limited operation of the system to the grounds of the

Institute of Rugby at the Palmerston North campus of Massey University in New Zealand. Measurements were taken in the area by RSM employees to ensure that the system would not interfere with microwave transmitters in the vicinity.

6 System Characterisation

In which the nature of the Multispectral system is inferred – spectrum measurements are taken – propensity for occlusion is observed – basic trials are conducted.

Once the best candidate for the trial system was identified and obtained, it was tested to gauge its effectiveness when used to track the positions of rugby players (and by extension, sports players in general). It was expected that the limitations imposed by the project constraints would cause less than optimal performance. For this reason, and to understand the workings of the system better for future planning and the generation of predictive models, several experiments were devised to test the performance of each component of the system.

The trial system was not waterproof and a 230V power supply was being used. To avoid damage to the system and injury to the user, no tests were performed outdoors in the rain.

All tests were repeated to ensure reliability. The following sections present the typical and, where appropriate, averaged results. The specific number of repeats is given in each section since it varied between test types.

In each test the procedure for testing will be explained. Also any extra equipment used will be described. For all appropriate tests the order of testing was randomised to mitigate the effects of any slow varying influences on the system.

6.1 Field trial equipment

The Sapphire Dart system, leased from Multispectral Solutions contained the following.

- A Sapphire Dart hub processor
- Four mid-gain receivers
- One omni-directional receiver
- One high gain receiver
- Four, 50m lengths of CAT 5E shielded network cable
- 11 asset tags with 1Hz transmission repetition rates
- 6 micro tags, 3 with 1Hz transmission repetition rates and 3 with 50Hz rates.

A petrol-powered, 230V_{RMS} generator provided power for the system. Interface with the tracking system and recoding of data was done using a laptop computer and the built-in Java web interface that runs on the hub processor.

The hub processor outputs the data either in real time in a TCP packet, or it can write the data to a text log file on the interface computer for later examination. The data are in the form of a string of ASCII, comma-separated values. Each line gives data about one object. Each line is formatted as follows.

<Data Header>, <tag #>, <X>, <Y>, <Z>, <battery>, <time-stamp>, <unit><LF>

The data header is a single character that defines the rest of the packet. In the standard operation of the system it defines the level of

positioning for a string, it can also signify a diagnostic packet used to help troubleshoot a system that is not working. The possibilities and meanings are given in Table 6.1.

Table 6.1: Table of data header possibilities.

Symbol	Definition
R	Three-dimensional calculation, valid for x, y and z. At least four receivers are required.
T	Two-dimensional calculation, valid for x and y only. A minimum of three receivers are required
O	Two-dimensional estimation. Occurs when the signal is visible to only two receivers. The object is shown to be on the line connecting the two receivers.
P	Presence data. Only one receiver has detected the tag transmission. No valid position, but the tag is shown to be in the network.
D	Optional diagnostic data output. This is useful when installing the system and for continuous monitoring of the system performance.

The system can be commanded to filter any of these packets depending on the user's requirements. For example, in the tracking of rugby players P-type data are unimportant for analysis and so can be filtered out.

The tag number field in the data string refers to the ID of the tag in question. This is an 8-character hexadecimal string. Each tag has its unique ID printed on it so that the user is easily able to reconcile the data with the object.

<X>, <Y>, and <Z> are the coordinates of the object in question. As has been shown in the previous table, the values of these fields are valid only for some calculations. The data header character defines which are valid in which situation. The coordinates are with respect to the user-defined origin. This origin is implicitly stated when the receiver and reference tag positions are calibrated.

The battery field gives a measure of the level of charge of the battery in the tag. It is shown as an integer between 0 and 15, where 15 indicates a fully charged battery.

The timestamp field gives the time that the particular calculation was completed. This value in no way represents the individual TOF values recorded by the receivers. The time is in UTC format. The default setting is for the times to be given to one-second precision but it is possible to command the hub to output the values with millisecond precision.

The unit field contains data on the group of receivers that calculated the position. In complex tracking situations such as in a factory it is advantageous to set up what are known as Virtual Groups (see section 5.3).

<LF> is the line feed character that signifies the end of the packet.

One field that is left off this data packet is the <DQI> field. Here DQI stands for Data Quality Indicator. This is an optional output that gives an indication of the reliability of the measurement. No units are given in the manual for this field, the only information given is that the higher the number, the more reliable the position. It is possible to filter low DQI data packets from the output stream.

An example of a typical data stream from the Multispectral system is shown in Figure 6.1. Here the system has been set to output only 2D calculated data (T packets). This means that anything of lower quality is ignored and anything of higher quality is expressed only as T data.

```
T,0000A7B2,60.34,12.79,1.52,8,1200007672.285,21
T,0000A7B2,60.56,12.10,1.52,8,1200007672.285,21
T,0000A7B2,60.60,12.15,1.52,8,1200007672.286,21
T,0000A7B2,60.34,12.04,1.52,8,1200007672.350,21
```

Figure 6.1: Typical output of the Sapphire Dart tracking system.

6.2 Tag spectrum measurements

This test is intended to check the spectrum of the signal transmitted by the tags. This information allows objective measurements to be made on the transmission characteristics and attenuation of the tags when obstructions are present. All of this makes it possible to gauge the suitability of the tags for use in tracking people.

The equipment used in this series of experiments was as follows:

- A Rohde and Schwarz FSV 9kHz-30GHz Signal and spectrum analyser.
- A quarter wavelength antenna and clamp stand
- One of each type of Multispectral UWB transmitting tags.

The antenna was specially constructed for these tests. The parameters were suggested by Dr Slava Kitaeff, physicist and radio astronomy expert at the Auckland University of Technology, New Zealand.

The tests were conducted in a radio 'clean room' that has been shielded against external radio interference and has only minimal electronic equipment inside. The experimental setup is shown in Figure 6.2. Note that the transmitter and the receiver are separated by 1m.

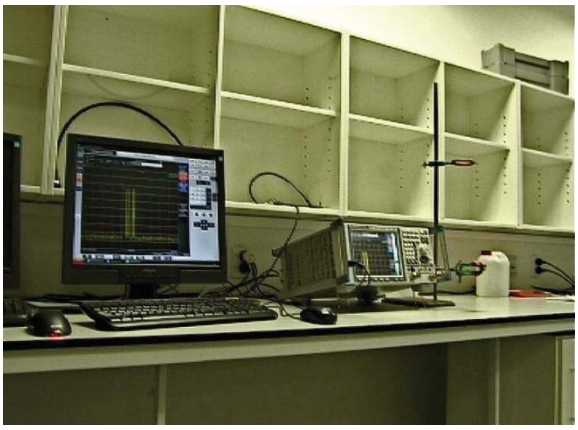

Figure 6.2: Tag-spectrum measurements, experimental setup.

Based on the recommendations of the technical staff of Multispectral, the analyser was configured using the following settings.

Table 6.2: Spectrum analyser configuration for tag PSD measurement.

Parameter	Value
Centre frequency	6.35GHz
Span (x-scale)	2GHz
Amplitude (y-scale) • Reference • Scale	 -40dBm 5dBm/Div
Resolution BW	3MHz
Video BW	3MHz
Sweep time	20s
Trace mode	Max hold

The spectrum of the system was given in one of Multispectral's technical papers[50]. Figure 6.3 shows this spectrum. The measurement of the spectrum in our particular operating environment was essential for obtaining meaningful results in later tests. Note here that the plot centres on approximately 6.1GHz and has a 4GHz

span and the spectrum of the transmitters is obvious.

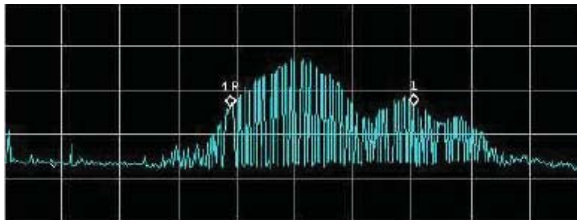

Figure 6.3: Multispectral's measurements of the PSD of the transmitter tags.

These tests will examine the transmitted spectra of the tags in free space and with an obstruction that simulates a section of the human body (a bottle of water in this case)[4]. The obstruction will be interposed halfway between the receiver and the tag, close to the tag and finally, halfway between the receiver and tag with the edge of the container on the direct path between the tag and the spectrum analyser antenna. These experiments will check for attenuation due to a complete obstruction by another body, attenuation due to a body wearing a tag and diffraction around a body respectively.

Zasowski, *et al.*[148] suggest that any obstruction of this type will result in significant attenuation. This is the hypothesis for this experiment.

10 traces are taken for each of the transmitter-receiver-obstruction configurations to ensure that there is no long term variation in the results. It is expected however that there will be no significant variation given that each configuration is static and the transmitters are designed to output the same signal for each transmission.

6.3 Receiver lobe patterns

The results of the spectrum tests were not sufficient to be able to fully map the interaction between the tags and the receivers in order to fully map the envelopes of both units. For the purposes of the project however, this is not a serious problem. Here all that is required is that the range of the system is known for various combinations of transmitter and receiver orientation. In this fashion, the shape of the receive lobe can be gauged and used for planning purposes. The actual range of the system will depend on the combination of receiving antenna and the transmission antenna.

Multispectral was questioned about the shape of the receive lobes and provided information on the electric and magnetic field shapes of the mid-gain and omni-directional antennae. They made no data available on the high-gain antenna. Unfortunately, it was not clear how these were related to the range of the system, which was what was required.

The ranges of these receiving antennas were specified by Multispectral to be 200m, 130m and 65m for the high-gain, mid-gain and omni antenna types respectively. We decided to measure the range of each of the receiver types to confirm Multispectral's data. Knowing the receiving footprint of each type of antenna allowed for planning the optimum configuration of receivers to yield the best results. The equipment for this test was as follows.

- The Sapphire Dart hub processor and appropriate cabling.
- One of each type of receiver unit (omni-directional, mid-gain and high-gain).
- A tripod stand on which to mount the receiver unit.
- Laptop computer for interface with the hub processor.
- One of each type of transmitter tags (Mini and asset tags), each with a 50Hz transmission repetition rate.

[4] A bottle of water was chosen over a human body to ensure that the test was completely repeatable with respect to obstruction shape and position. This test is independent of body type or shape.

- Distance measurement tools.

To conduct the experiment, the single receiver to be tested was mounted on the tripod and connected to the hub processor. Only one receiver was used in this experiment so configuration of the system required only that the receiver ID be made active; the position of the receiver and a reference tag were not required. The bore sight of the receiver is defined as being directly in front of the receiver unit and this has also been defined as the 0° point of the measurements. The receiver was mounted 1m from the ground; excessive wind conditions on the day made it dangerous to mount the receiver any higher[5]. This low mounting did cause some issues that are fully explained in the results section (section 6.6.1)

The hub processor was configured so that, where appropriate, the output was a P-type packet indicating that a transmission from the test tag had been received. When the tag was fully in range of the receiver, there was a stream of packets at a frequency of approximately 50Hz (the repetition rate of the test tag). When the tag was moved beyond the range of the receiver, the stream of packets ceased. It was hypothesised that there would be a transitional period where the tag was only partially visible to the receiver. Within this region, which will be referred to as the penumbra of the receiver, the signal coverage will be unreliable and hence the accuracy and number of output positions will be compromised.

The test tag was held at the orientation that gave the strongest signal, found through observation. A single tag orientation was used to ensure that the changes in the receiving range of the particular antenna were being measured. In addition to the fixed orientation, the tag was held away from the body so that the body of the carrier did not impinge on the transmission[6]. The tag was moved away from the receiver until the signal was no longer detected by the tracking system. The distance from the receiver to the point of failure was then defined as the maximum range of that receiving antenna at that orientation. Distances up to 50m were measured using a tape-measure (longest tape available); distances above were measured using a hand held GPS that used the receiver location as a reference.

For the purposes of this experiment it was assumed from examination of the Multispectral data that the receiving ranges of the system are symmetrical about the bore sight. Measurements of range were made on orientations to the right of bore sight in 10° increments. A small number of measurements of range were taken to the left of bore sight in order to confirm the symmetry assumption.

This test was performed in triplicate. After three tests it was found that the standard deviation of measurements was acceptably low (standard deviation of ranges less than 3m).

6.4 Tag placement

The placement of the tags on the objects is critical for optimum operation of the tracking system[149, 150]. The main criteria provided by Multispectral for placement requires that the tag be as high up on the object as possible to avoid ground bounce and placed such that the transmissions from the tags are not shielded.

The actual placement of the tag will be dependent on the tracking application. In the player-tracking application, it is a requirement that the tag be placed where it is not going to injure the player in the event of a collision. This must be

[5] In a game situation with these conditions the stands would either be weighted or pegged down. At the time of testing it was assumed that the stands' wide base would be sufficient to stop a tip over. The decision to reduce the height of the receivers was made after one of the stands tipped during set up.

[6] It is recognized that, in a game situation, the tags will be close to the body and thus the transmission will be affected however, this test aimed to measure reception lobes independent of other factors. Body shielding is handled in a later test.

traded off against the requirements of height and non-shielding.

Testing was performed with the tag in the following locations: shoulder blade, right elbow, back of the leg and the rear of the shoe. Padding was used to prevent damage to either the player or the tag. Each placement was evaluated based on the proportion loss of visibility when the player moved along a set course. The receiver configuration was held constant for all tests and in all cases a player type tag was used (see Table 5.1).

The track traversed by the player in all of the tests is shown in Figure 6.4. The player starts at *(60,0)* and moves towards the centre of the field; from there he traverses the circuit in an anti-clockwise direction. Because the exact path of the figure of eight was difficult to hold constant it has been marked on the figure as a hexagram *(110,60)*. This path need not be constant since this particular manoeuvre was intended to show the effects of player orientation near the corner of the field. Because this section of the player's path is not constant, the figure-of-eight section of the track is removed from accuracy measurements.

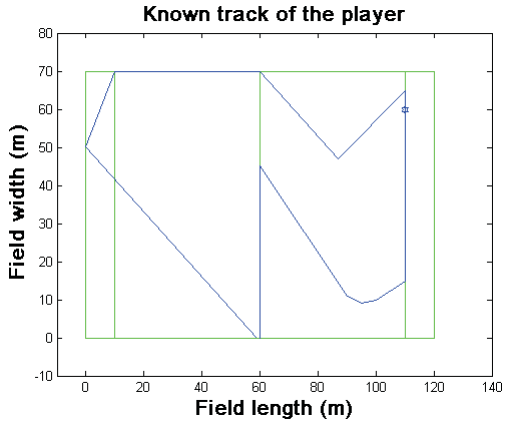

Figure 6.4: Known track of the player for the tag placement tests.

The tag placement tests were performed using the Sapphire Dimension system. Six receivers were used in this experiment. The receivers were evenly distributed around the area of interest as shown in Figure 6.5. Six, mid-gain receivers are used here to ensure full coverage of the area of interest. This number and configuration was arrived at after examining the receiver lobe patterns. It was desirable that for this test that there be no areas of partial coverage in order that tag invisibilities be caused by tag placement and not receiver configurations.

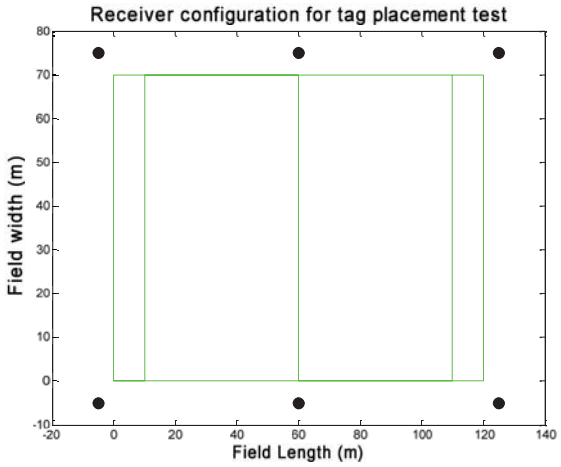

Figure 6.5: Receiver configuration for tag placement test.

6.5 Sapphire Dart trials

A series of tests was performed using the leased system to determine its effectiveness in a player-tracking situation. Of particular interest in this series was to determine the minimum number of receivers required to allow the system to operate satisfactorily.

In terms of the specifications for the real world application, these tests allowed the ease of operation, accuracy, and timely recovery of data constraints to be gauged. Some further information was gained that allowed the operation of the system to be modelled so that the best configuration of the system can be found.

The equipment for this series of tests was as follows:

- The Sapphire Dart hub processor.

- A laptop computer for data monitoring and collection.
- Four mid gain receivers, including the shielded CAT 5E Ethernet cables that came with the system.
- Four receiver stands that can be extended to 2.5m from the ground.
- A mixture of asset and micro tags.
- A 230V AC power supply. In this case, a petrol powered generator.

One of the rugby fields at Massey's Institute of Rugby was selected for these tests. The markings on the pitch were used as known locations on the field, thus making analysis of the accuracy of the data possible. The placement of the receivers was planned in advance to save in setup time. It is expected that this kind of planning will take place when the full system is implemented.

The length of the network cables (~50m) in this trial system limits the possible separation of the receivers in these tests and the placement of the receivers must be altered accordingly. Longer cables would allow the maximum separation between the receivers to be increased; however, there is a limit to the length of cable that can be used. According to the manual the cable length must not exceed 91m (300ft).

The setup time of the system with a single operator was 45 minutes. This included unpacking all of the equipment, measuring the exact locations for receiver placement (with approximately 0.5m precision in these trials), erecting the receivers on their stands, running the cables back to the hub processor and configuring the system for operation. With sufficient planning, an experienced pair or trio of operators should have the system operational inside 30 minutes. The system itself is easy to set up with minimal experience. The instructions manual is very clear on the setup procedure and the labels on the equipment are easy to decipher.

Recall that when the receiver and reference-tag positions are configured the input coordinates imply the origin of measurement. For all Sapphire trials, the origin is at the intersection of the sideline and the goal line. Note that this origin has remained constant across the vast majority of this research.

Each trial is conducted only once since the trials will include a large number of tag positions measurements.

6.5.1 Three receiver system

Three receivers represent the absolute minimum number required to calculate a two-dimensional position for an object. The receivers were placed as in Figure 6.6. Each receiver was oriented with the bore sight towards the centre of the field. The locations for the receivers were measured using a 50m tape measure with the lower-left-most, dead-ball line as the origin (as the field is aligned in Figure 6.6). The receiver locations shown in Figure 6.6 are with respect to the front of the receivers where the antenna is housed. The centre pole of the receiver stand was used to correctly locate the receivers with an offset included in measurement to account for the difference between the axis of the stand and the location of the antenna (approximately 180mm). The receivers were placed with better than 0.3m accuracy.

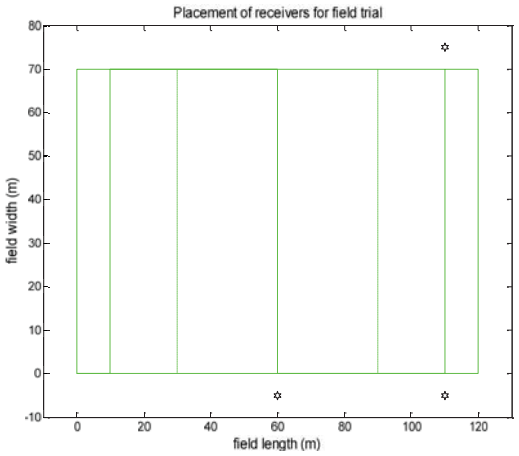

Figure 6.6: Configuration of three receivers for field trials.

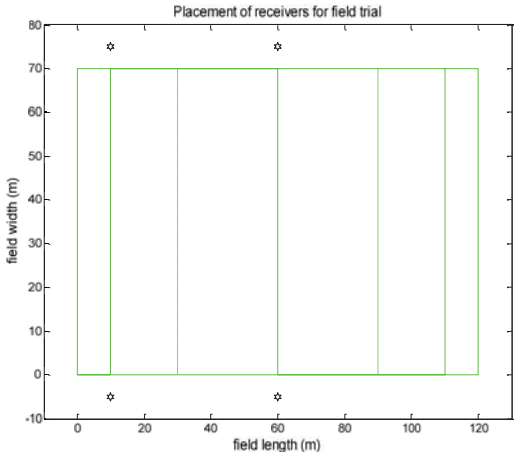

Figure 6.7: Configuration of four receivers for field trials.

In order to test for accuracy, a person carries a tag along the major field lines; the tag is held in the right hand and away from the body of the person such that the antenna of the tag is not obstructed[7]. At each intersection the person halts and rotates on the spot in order that the output data show a concentration of points, which should be easily identifiable in the output.

6.5.2 Four receiver system

Four receivers represented the maximum number that could be used with the trial system because only four Ethernet cables of appropriate specifications were available. Once again the maximum straight-line separation between receivers could not exceed 50m due to the cable length[8]. The receivers were arranged as in Figure 6.7.

This configuration was chosen to simulate tracking of half the field in a game situation to check the effectiveness of the system, where effectiveness is defined as a high visibility and accuracy. Another aim of this test was to evaluate the level of care required when setting up the receivers of the system in order to get sufficient accuracy of position measurements. The more care that is required in placing the receivers the longer the setup process will take. For this test the receivers were placed to within one metre of the location recorded in the system configuration.

6.6 System characterisation results

The tests described yielded the following results.

6.6.1 Tag spectrum measurements

The initial test for this experiment aimed to find the spectrum and power output by the tag. The transmitter tag was placed 360mm from the receiving antenna so that there was approximately 10 wavelengths separation ($2.99 \times 10^8 / 6.35 \times 10^9$) in order to escape the near field effects. Figure 6.8 shows a time-domain trace of the signal received from the tag. Here the trace is triggered by the first pulse from the

[7] Again, this test aimed to show the performance of this system independent of any form of interference.

[8] In this particular test, the hub processor is located under the left goal post and the receivers are daisy-chained to either side. The cables stretched over the play area and thus this would not be a suitable configuration in a game situation.

tag and it is clearly visible above the ambient noise.

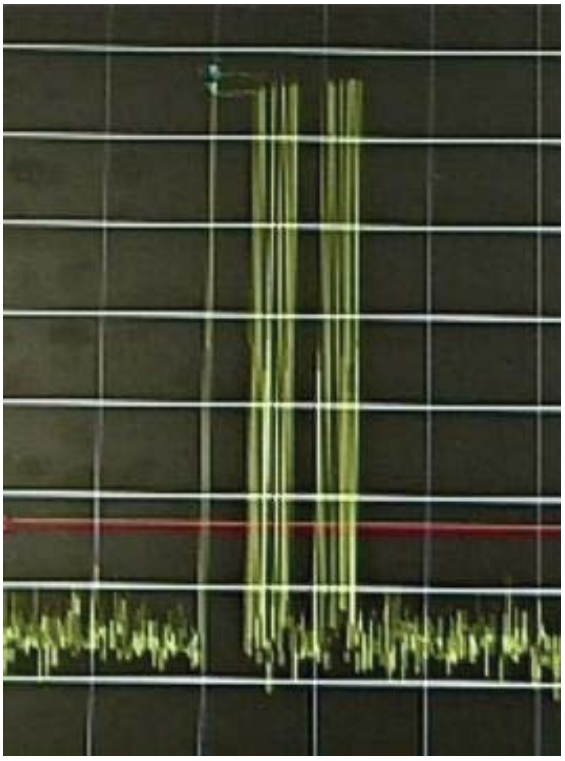

Figure 6.8: Free-space, time-domain signal from the badge tag.

A trace in the time domain is a measurement of signal rather than spectrum. Dr Kitaeff suggested time domain after measurements in the frequency domain failed to produce useable results (Figure 6.9 shows the result of the same test in the frequency domain). Some review of the tag operation is useful here. The microcontroller on the tag encodes information by modulating a 1MHz square wave (modulation occurs via a logical and operation). The antenna is excited every time there is a leading edge. The tag therefore outputs a UWB pulse lasting approximately 10ns at a rate of up to 1MHz where the position of the pulses in time confer the information (pulse-position modulation). The spectrum analyser can only measure one frequency at a time and the pseudo random nature of transmission makes it unlikely that the spectrum analyser will be measuring the correct frequency at the correct moment to obtain a complete spectrum. Also the tags spread their energy over a 2GHz bandwidth which means that the energy of at any frequency from the tag will be difficult to distinguish from noise. Repeated measurements could only come up with traces similar to Figure 6.9. Time-domain measurements can be used to measure attenuation of the signal more effectively.

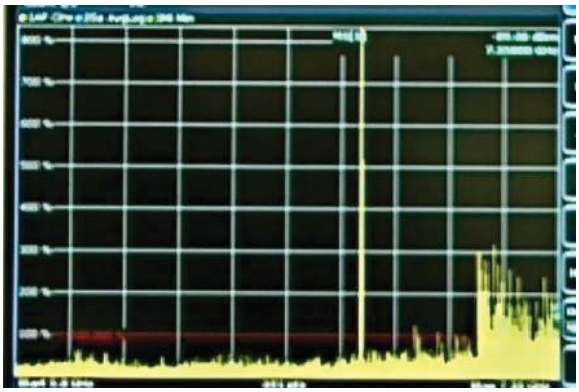

Figure 6.9: Attempt to capture frequency-domain trace of UWB tags.

The next test involved interposing a bottle of water (with measurements of 200(L) x 100(W) x 300(H)mm) half way between the transmitter and receiver to see how a body attenuates the signal. Figure 6.10 shows the result of this test.

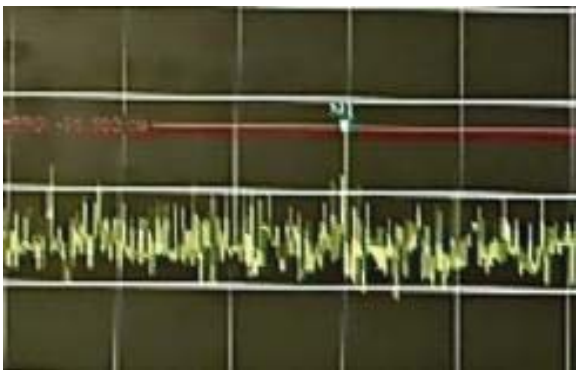

Figure 6.10: Water obstruction halfway between transmitter and receiver.

Here there is almost complete attenuation of the signal. It appears that no diffraction has occurred around the obstruction. If this tag were attached to a person who was to be tracked, the signal would be occluded to this receiver.

The next test involved moving the water container close to the tag to simulate a tag being attached to an object. The resulting signal trace is shown in Figure 6.11.

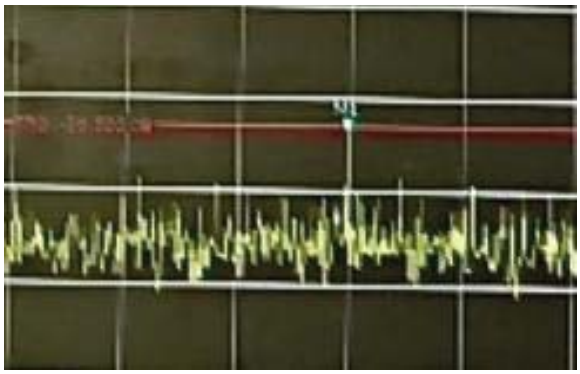

Figure 6.11: Water obstruction close to the transmitter.

The position of the obstruction along the direct path between receiver and transmitter appears to have no effect on signal attenuation. If there is an obstruction between receiver and transmitter then the signal will be fully attenuated.

Finally, the last test places the water bottle half way between the transmitter and receiver along the direct path but steps it off to the side such that only the edge of the bottle is on the direct path. The resulting signal trace is shown in Figure 6.12.

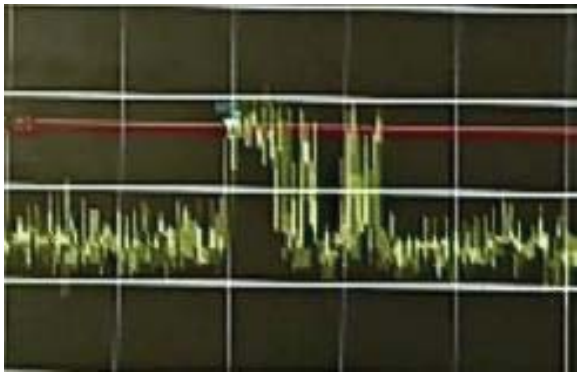

Figure 6.12: Edge of water obstruction on direct path.

While the signal is visible in this plot, it has been significantly attenuated. The shape, however, has been maintained which suggests that all transmitted frequencies from the tag are being attenuated by a similar amount. Were they being attenuated differently the shape of the signal in the time domain would have been deformed.

Given that the transmitters have an approximate wavelength of 0.05m and that the obstruction is halfway between the receiver and transmitter (separated by 1m) then the perpendicular distance where the reflected signal off the obstruction is half a wavelength longer than the direct path is 0.11m. Given that the obstruction is 0.2m wide this suggests that the attenuation is caused by the water obstruction encroaching on the Fresnel zone of the transmission and the reflection off the edge of the obstruction is causing destructive interference at the receiver.

In order to confirm destructive interference, the water obstruction was moved further off the direct path. Similar results were obtained (shown in Figure 6.13). As can be seen, the signal is attenuated by a similar amount even though the water container is not obstructing the direct path.

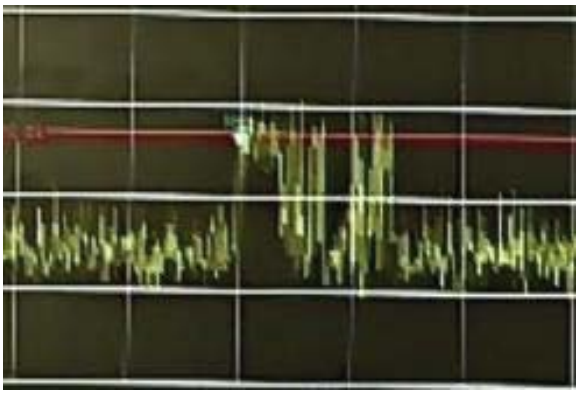

Figure 6.13: Water obstruction close to, but not on, the direct path.

These tests were performed on each tag type in turn, with similar results. Note that it was not possible to test how the signal would be attenuated when the obstruction is further off the direct path because of time and space constraints with the lab. Inferences can be made

from the tests performed that the further away from the direct path an obstruction is, the less its attenuating effects.

In terms of this research, these results suggest that the hypothesis was correct: Obstructions will cause total attenuation of the signal when they are on the direct path. This also confirms the work of Zasowski, *et al.*[148]. Therefore, obstructions must be considered when the performance of tracking system is measured and when planning for the system configuration.

6.6.2 Receiver lobe patterns

The receiving lobe shapes are shown in Figure 6.14 as blue lines. They are shown in polar form, as this is most suited to visualisation and planning purposes in this situation. Where the manufacturer's stated lobe shape data were available, it is also shown in red. Note that the values from the manufacturer have been converted from gain measurements and normalised to the maximum measured range and thus the manufacturer's data are indicative of shape only and not range. The measured values are shown with plots for ±1 standard deviation (green lines) to indicate the spread of the measurements. Note that is some cases the spread is so low that the green lines overwrite the blue.

Measurements of range at angles to the left of bore sight (angles 0° to 180° in Figure 6.14) confirmed that the maximum range is symmetrical about the bore sight of each receiver. Furthermore, measurements of range using the micro tag confirm the hypothesis that the shape of the receive footprint remains the same regardless of the tag. The values of range are greater or lesser depending on the tag. This difference in range can be attributed to the different transmit antennas used in the tags.

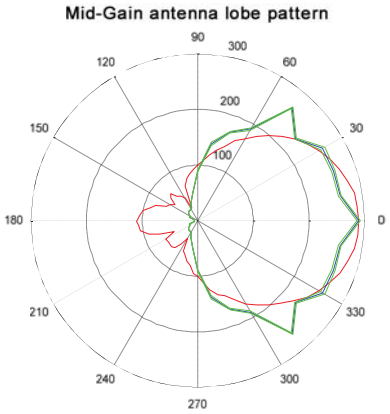

(a)

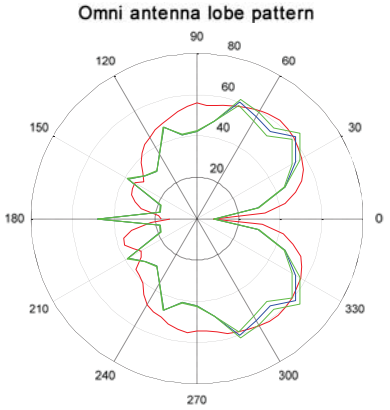

(b)

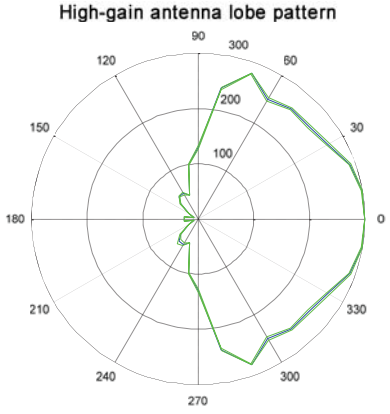

(c)

Figure 6.14: Polar plots of measured range of each receiver type in conjunction with a mini tag. (a) Mid Gain, (b) Omni, and (c) High Gain.

It should be noted that in Figure 6.14(c) the signal at 0° was still strong at 300m. Unfortunately, this was the boundary of the testing area. Moving the tag beyond this boundary would have introduced other obstructions into the measurements. The most notable obstruction is a large hedge around the perimeter of Massey University's Rugby Institute. The standard deviation of measurements drops to zero at 300m because of this boundary (measurements from all three tests had range saturating at 300m).

During the measurements, it was noted that there were regions where no signal packets were output by the hub, initially suggesting that the tag had gone out of range. However, when the tag was moved further from the receiver, the signal became strong again. Consultation with Multispectral revealed that this loss of signal was due to destructive interference from the transmitted signal being reflected from the ground (ground bounce). Ground bounce is a form of multi-path fading and given that UWB technology is resilient against multi path fading this suggests that the path difference between the LOS path and the indirect path was very small. The suggested solution to fix this issue was to increase the height of both the tags and the receivers where possible. A test with the receiver at 2.5m from the ground confirmed this solution.

One feature that is particularly worthy of note here is the range of the omni antenna at bore sight. It was expected that the range of all receiver types would be greatest at bore sight yet both our results and the Multispectral measurements show that the range of the omni receiver at bore sight is almost zero. Multispectral was queried about this but no reply was forthcoming. Given that the measured results approximately match the manufacturer's measurements it is plausible to assume that this zone is a result of the antenna type used; though it is unclear why this would be favourable in the antenna design.

Also worthy of note in these plots is the range and receive arc of the receiver fitted with the high-gain antenna. Here it was expected that an increase in range due to antenna type would be at the cost of a decrease in the receiving arc compared with a lower gain antenna type. What has been observed suggests that the range of the high-gain antenna is greater than the mid-gain antenna in all directions.

The spread of measurements in all cases is low suggesting that the tests are reliable. In all cases the standard deviation of the measured ranges is between 1-6% of the measured range (calculated as one standard deviation as a percentage of measured range). The standard deviation of measurements below 50m was significantly lower than those above 50m; this was expected and is due to a GPS unit being used to measure greater distances and a tape measure for lower distances (the GPS had a precision of 2m with full satellite coverage).

Examination of these footprints suggests that either the high-gain or the mid-gain receiver would be best suited to a tracking system for open field sports. Both have their best range to the front of the antenna and minimal range at the rear. This lends itself to a tracking system where the receivers must be placed around the periphery of the field of interest and signal reception from outside the field is not desirable. In the ideal system, we would like to have each receiver able to receive signals from any point on the field (achievable in sports such as rugby, soccer, hockey, etc). This suggests that in the worst case of receiving a signal from a tag at one extreme corner of the field from the diagonally opposite corner would require 150m of range from the receiver-tag combination. Additionally, the range of the receivers at ±90° from bore sight should be great enough that the field close to the receiver location will have coverage. Both the mid-gain and high-gain antenna types fit these criteria and indeed both types will almost completely cover the whole field of interest. Limiting the coverage area to the field of interest will help to prevent the system from

receiving unwanted signals from outside the field of interest. For this reason, the mid-gain receivers are the best choice for the final system.

6.6.3 Tag placement

Results from this series of tests have been ordered anatomically from the player's head to their feet. In all cases, the intended player track has been included in the results to indicate the approximate track of the player. Recall that six receivers are equally spaced along the length of the field on either side (see Figure 6.5).

Figure 6.15 shows the results of the trial with the tag on the shoulder blade.

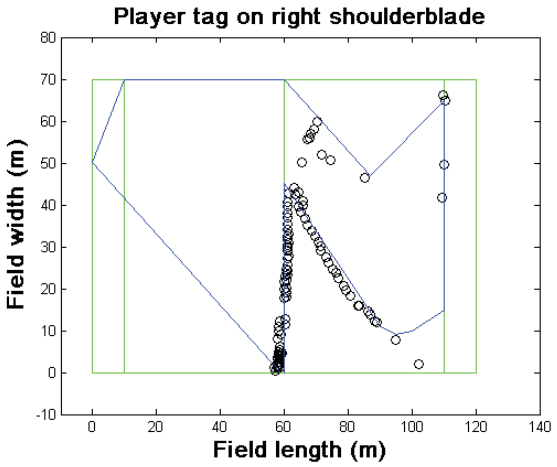

Figure 6.15: Tag placement test results with the tag on the right shoulder blade.

Overall, this placement can be seen to be unacceptable, as the tag is invisible over a large proportion of the field. Over the course of this test, the tag was invisible 45% of the time. This is a high invisibility rate given that the receiver lobe patterns in Figure 6.14 suggested that the tag should be within range of at least 4 receivers for most of the track. This suggests that the human body is having a detrimental effect on tag transmissions (corroborated by the results in section 6.6.1).

Figure 6.16 shows the results of the trial where the tag was placed on the player's elbow.

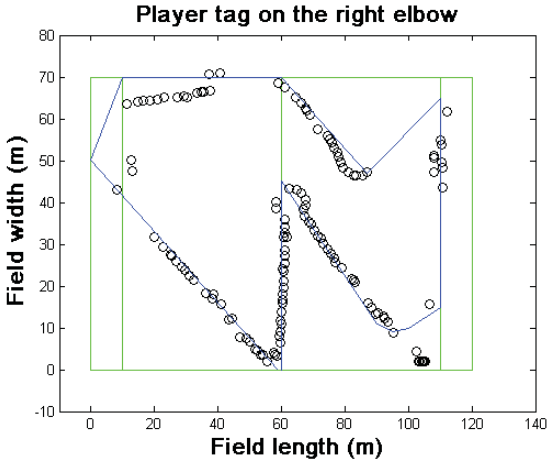

Figure 6.16: Tag placement test results with tag on the player's elbow.

The tracked positions here are much more spread out and appear over the whole field. However, there are still large gaps in the track where the signal was lost. Over the course of the test, visibility with the tag was lost 41% of the time, making this placement only slightly superior to the shoulder blade in this respect.

An added failing to this placement apparent in Figure 6.16 is that the ability for the system to track the player's centre (key to getting a smooth track) is compromised when the tag is placed on the elbow since its position relative to the player's centre will change significantly throughout the game. Furthermore, given the degrees of freedom of the elbow, it is highly likely that there will be periods where the orientation of the tag is entirely non-optimal, thus increasing the rate of invisibility.

Notice in Figure 6.16 that the player has performed wide arcs at the waypoints, but for most of the straight-line movements near the centre of the field, the recorded track of the player corresponds to the intended track. Near the edges of the field however, the recorded track differs significantly; more than could be attributed to variations in player path.

Figure 6.17 shows the results of the test with the tag on the back of the player's left ankle.

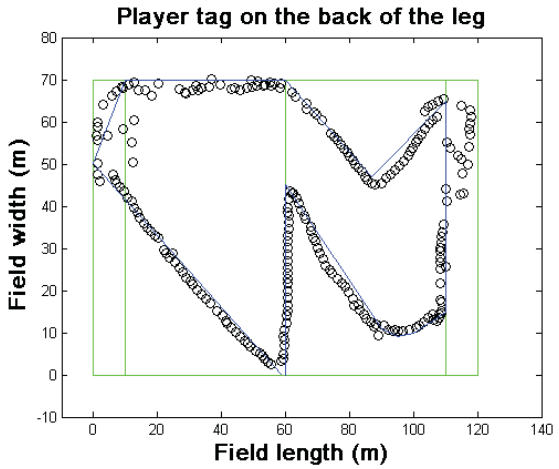

Figure 6.17: Tag placement test results with the tag on the back of the leg.

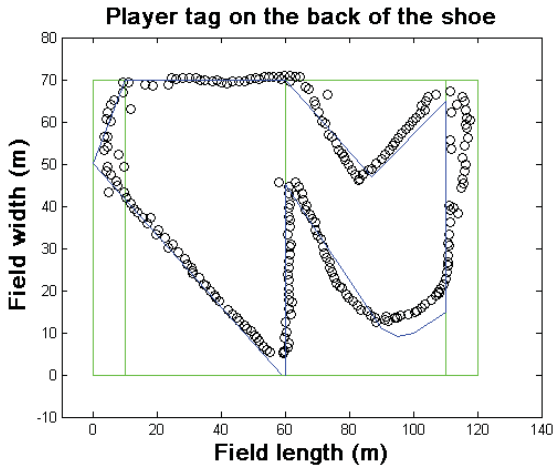

Figure 6.18: Tag placement test results with the tag on the back of the shoe.

On the basis of signal visibility, this tag placement is the best so far. In all areas of the field there are a significant number of visible positions. Over the course of the test, visibility with the tag was lost less than 1% of the time. Furthermore, in most areas of the field, the accuracy appears high; this is especially true near the centre of the field. Note that near the edges of the field there are some points that deviate considerably from the known track of the player. This inaccuracy is most likely caused by the orientation of the tag being away from most of the receivers. These deviations would be identified and removed by the reconstruction algorithm (see chapters 9 and 10). The deviations from the intended track on the right side of Figure 6.17 are mostly caused by the figure of eight performed by the player although some positioning error still exists here as evidenced from the discontinuity when the player begins to move towards the centre of the field again.

Figure 6.18 shows the results of the test with the tag on the back of the player's shoe.

In this test, the signal from the tag was invisible only 1.6% of the time. This is slightly lower performance than the tag on the back of the leg but the accuracy is greater, especially near the top of the field. The reconstruction algorithm would be able to repair or recover this track effectively. Note that while the recorded track does not follow the intended track the general shape conforms. This indicates that the player followed the approximate shape of the intended track but did not move to all waypoints precisely; for the majority of this test, there are no significant deviations from the player's track. Once again, deviations are occurring near the edge of the field where the tag will be oriented towards the minimum number of receivers.

From these tests it is apparent that placing the tag on either the player's legs or shoes produced the best performance. This is the opposite to the instructions from Multispectral which said to place it as high as possible to counter the effects of ground bounce. Multispectral's advice did not take into account the shielding effect of the human body. In all of the upper body tests performed, the player's torso would have significantly interfered with the transmission of the signal. Body obstruction is therefore a more significant impediment to tracking than ground bounce. This highlights the gap in the knowl-

edge and confirms the similar findings of See, *et al.*[149]

Low down on the leg provides the minimum of body mass to interfere with the signal. The performance with the tags placed on the shoe and on the back of the leg is sufficiently similar that it can be left to the operators to decide which is more appropriate, especially in terms of player comfort and tag durability.

6.6.4 Sapphire Dart system trials

The results of these tests are shown in Figure 6.19. The field lines have been included for perspective. Recall that the first test was conducted with three receivers (the minimum required to calculate a two-dimensional position) placed to cover the lower right of the field (see Figure 6.6). The second test was conducted with four receivers covering the left side of the field (see Figure 6.7.

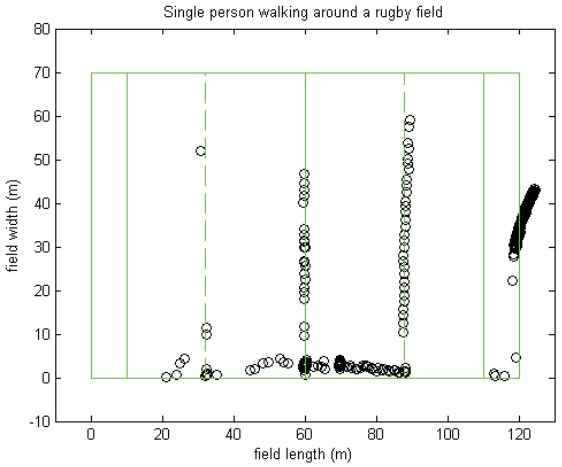

Figure 6.19: Result of a single person walking the field lines in a three-receiver system.

Note here that a significant portion of the upper left side of the plot is devoid of points. This shows that the there was not sufficient coverage by the receivers to calculate a position at all times. Though it is not clear from the plot, the missing positions correlate with the person walking in the negative y direction. Thus a probable cause for invisibility is that the body of the person was interposed on the direct path between the tag and the receiver, effectively shielding one receiver from the signal and precluding the calculation of a position. This is a serious detriment to the system as it was configured here. It suggests that the three-receiver system (mathematically the minimalist system for tracking in a two-dimensional TDOA situation) is not sufficient to track even a single human especially when approaching the maximum range of some of the receivers. The minimalist system capable of tracking a game of rugby will need more receivers.

It can be clearly seen that where the object approaches the receivers, the accuracy degrades. At all times during the tests the tag was moved along the field lines whereas in Figure 6.19 the positions can be seen to significantly deviate from the field lines. Recall that the geometry of the receivers with relation to the tags will affect the accuracy of position measurements.

The high level of invisibility in the upper left portion of the field in Figure 6.19 is also a concern. This is a clear violation of the requirement that there be at least one position measurement for each player every second.

This minimalist system test confirms that it is possible to calculate a position of an object in two dimensions and that the system will function in this minimalist fashion. However, the lowered accuracy and visibility makes this configuration unsuitable.

The track from the four-receiver trial is shown in Figure 6.20. At all times the person carrying the tag moved at walking pace. The field lines have been added to this plot as appropriate.

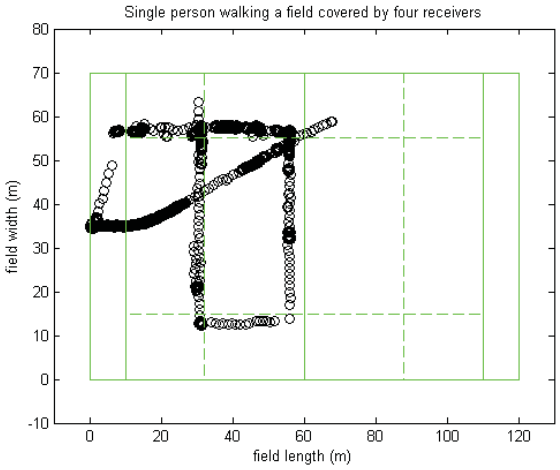

Figure 6.20: Results of a single person walking field lines in a four-receiver system.

In this plot it can clearly be seen that more positions are being recorded compared with the three-receiver test, suggesting that the hypothesis that three receivers were simply too few to get a two-dimensional position for every tag transmission was correct. With four receivers there is some redundancy, a valid position can be calculated when the tag is invisible to one receiver. However, there are still some areas where the positions are sparse (the lower right, for example). The receivers near the corners of the box form these areas. These corners will provide the least coverage, especially if the body of the person is interposed between the tag and the other three receivers. Away from the corners, there is sufficient visibility for the tags to satisfy the requirement of one position every second.

The accuracy of the positions in this test was less that expected. Note in Figure 6.20 that while the track of the person moves parallel to the field lines it does not directly overlap the field lines that the player was walking. The track is between 0 and 4m from the nominal positions of the field lines. While the positions do jitter a small distance about the trend of movement, the major component of the inaccuracy in the positions is a constant offset of approximately 2m in both dimensions. It is the offset that is most concerning here. With such a large offset the accuracy of the system is significantly less than required for the real-world application of sports tracking. The major contributor to this inaccuracy is the lack of accuracy in the placement of the receivers of the system, although it is unclear how each individual placement inaccuracy contributes to the position inaccuracy as a whole. The placement accuracy was taken into account in further testing of the system and receivers were placed with 0.5m or better accuracy.

Note that the diagonal track across the centre of the field was the player moving from the centre-top of the field to the hub processor. It was included here because of the high output-position density. Furthermore, the track is an accurate representation of how the player was moving based on visual inspection of the player and field lines during the test.

6.7 System characterisation conclusions

In this chapter, the selected player tracking system was characterised within the scope of player tracking. The findings have demonstrated a gap in knowledge in the application systems to track humans (specifically optimum tag placement). The findings have also produced useful data for the next phase of the research, the modelling and optimisation of the system.

The transmitted signal from the tags was analysed and it was found that any incursion into the direct path by a water obstruction (that simulated the human body) would result in total attenuation of the signal. Also incursion into the Fresnel zone of the transmitter and receiver, approximately 10cm from the direct path, would result in destructive interference that would partly attenuate the signal. From this it has been determined that, when modelling or planning a tracking system, obstructions must be considered.

The receiver antenna lobe patterns were mapped and it was found that both the mid-gain and high-gain antennae give performance suitable for use in player tracking applications such as sports played on a small field. The results of these tests will be used in the receiver placement optimisation algorithms to generate look-up tables to allow more accurate modelling of the tracking system. With these data it will not be necessary to assume that the receivers have a hemispherical lobe pattern and thus no conservative estimate of range will be required.

It was found that that best placement for the tag is low on the leg or on the shoe, where interference from the player's torso is limited. The decision on exactly where to place the tag on the lower leg is left to the operator, as the performance is similar, regardless of location on the leg.

It has been shown in the system trials of the Sapphire Dart system that a minimalist system, consisting of only three receivers, is not sufficient to track even one person around the field and would be totally incapable of tracking two full teams of rugby players.

Using four receivers to cover half the field gave much better results when tracking the movements of one person around the field. There were sufficient data from the four-receiver test to satisfy the requirement of the real world system of there being one position recorded every second for each player. Note that four receivers covering half the field imply that six receivers would be required to cover the whole field (four receivers to cover each half of the field but the two at the centre will be shared, totalling six receivers).

The three and four-receiver tests showed that both the placement of the receivers and the accuracy of that placement are critical to the performance of the system. Inaccurate placement of the receivers will introduce an offset into the results and placement of the receivers too close to the field of play will cause position inaccuracy whenever the players approach those receivers. Great care must be taken to ensure that the receivers are placed accurately, placement to the nearest metre should be sufficient to produce positions accurate to 0.3m, the specification for player tracking. In addition placing the receivers between 5 and 10m back from the edge of the field of play should reduce the position inaccuracy when players approach the edge of the field of interest.

7 Player Visibility Modelling

The method for calculating player visibility – the need for such a model – available data – invisibility causes – handling of static obstruction – validation and testing of the models.

In any type of tracking system, a lack of visibility to any of the receivers will have a detrimental effect on calculated positions, from a slight loss of accuracy to a total loss of position. The number and positioning of receivers will determine how robust the system is to visibility losses.

A visibility loss represents increased complexity in planning a position-tracking system configuration. Even an open field becomes a complex and dynamic environment when the number of objects to be tracked increases. Given the limitations from the research specifications, it is not acceptable to counter the problem of visibility loss by increasing the number of receivers. A more elegant solution is necessary to maximise the efficiency of the system. Furthermore, the development of tools that can be used to increase efficiency fills a major gap in knowledge that has, up to now, been dominated by practice of arranging receivers in regular patterns and increasing their numbers to increase visibility and efficiency. This statement is supported by the work of Gillette, *et al.*[2] who show that a linear solution exists for TDOA position calculations.

A model can be used to predict the performance of a specific receiver configuration in terms of visibility. The model can also be used as an objective function in an optimisation aimed at maximising visibility and hence, maximising the performance of the tracking system.

The visibility model covered in this chapter was implemented in Matlab 6.5 and installed on a standard PC running Windows XP.

7.1 Input data for modelling

All models require data for verification (checking that the algorithm works), validation (checking against real-world examples) as well as to run an optimisation. These data can be generated artificially or be taken from a real test case. Either way, the data should be representative of the case.

Player-position data from several games of rugby that had been tracked previously[9] were selected for input to this model as being the most representative. These data were in the form of a comma separated value (CSV) text file.

Position information within the CSV file was represented as a two-dimensional position for every player on the field (two teams of 15 players) at intervals of approximately 2 seconds (henceforth referred to as a frame).

The sample data provided by the Industrial Partner consisted of 10 previously tracked rugby games with a total frame count in excess of 30,000.

[9] Tracking of the player position data was performed using the Industrial Partner's manual video tracking system, described in section 1.4.

Neither the orientation of the player nor their stance was recorded within the input data. Both are important to the visibility model. The stance of the player was assumed always to be upright. This assumption will affect the accuracy of some visibility calculations.

The data record only the active parts of the game; the positions of the players are not recorded during stoppages of play. This is done because the data contained between plays are of no use for the Industrial Partner's analysis of a game. Without the data between plays there are significant discontinuities in player tracks.

As a consequence of these limitations, some pre-processing is needed to render the data suitable for input to the visibility model.

Orientation estimation is the most significant of the pre-processing operations. In the absence of any other data, it was assumed that the orientation of the player was the same as his direction of motion between the current and previous frame. There are situations where this assumption will not hold, such as a player back-pedalling, however, this occurs infrequently and was therefore ignored.

A requirement for a minimum of 0.5m movement in any direction was imposed to avoid noisy data rapidly and randomly changing the measured orientations of stationary players.

Similarly, orientation was not calculated using data that spanned discontinuities (such as at the beginning of a play). To calculate orientation at the start of a play, the current and following frames were used, rather than the current and previous frames. The implicit assumption in this was that the player's orientation did not change in the first two frames of each play.

The decision tree for the calculation of orientation is shown in Figure 7.1. This tree was applied to every player in every frame.

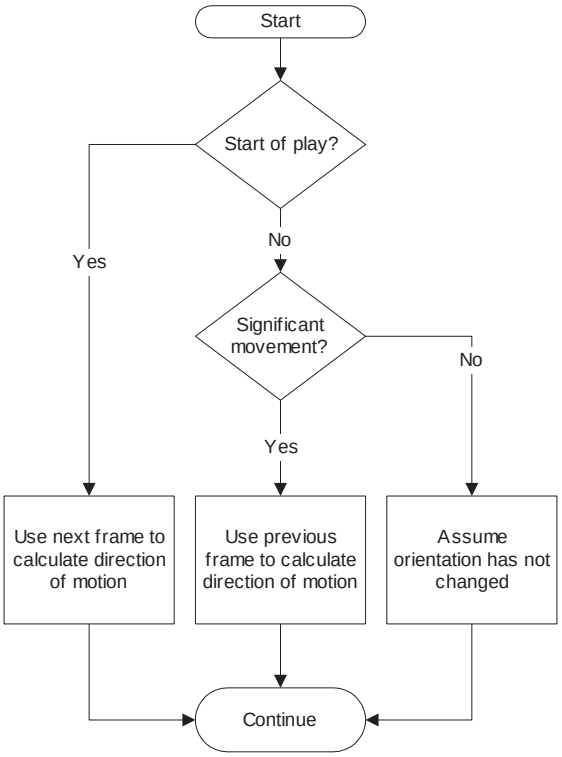

Figure 7.1: Decision tree for player orientation calculation.

Once the orientation data have been compiled, there is no longer the need to analyse the data in a sequential fashion and the visibility model can examine the data in any frame order. This allows frames to be selected at random from the data suite.

Finally, a scaling factor is introduced to convert the positions from pixels to metres. Pixel numbers are used in the manual video coding of the game to avoid the need to use floating-point values to represent positions. Since the resolution of placement in the Industrial Partner's coding system is one pixel, this was the obvious choice. The decision was made to convert to metres for the visibility model because the data are more meaningful to the observer in this form. The scaling factor for the sample data is 10 pixels to 1 m.

According to the Industrial Partner, the sample data were representative of player positions for

the game of rugby. To analyse this assertion the players' occupancy of the field was examined.

In an occupancy plot, player positions are quantised to a given resolution. The resolution is changed depending on the analysis that is being performed. For example, to determine which team is dominant in a game, the resolution could be set so that there is one area for each half of the field. A team with significant dominance would force all of the players to one side of the field for a significant portion of one half of the game (the teams switch ends at half time).

It is postulated that the occupancy plot of any particular game, with sufficient resolution, will be unique to that game. The term *occupancy signature* is defined here to refer to the field occupancy for a particular set of data. The occupancy plot for the data suite is shown in Figure 7.2. The colour bar to the right of the plot maps the colours to frequency of occupation. Here the graduations on the colour bar are the number of players (hence the dark red spot on the plot was occupied by 1500 players over the course of the games given the indicated resolution).

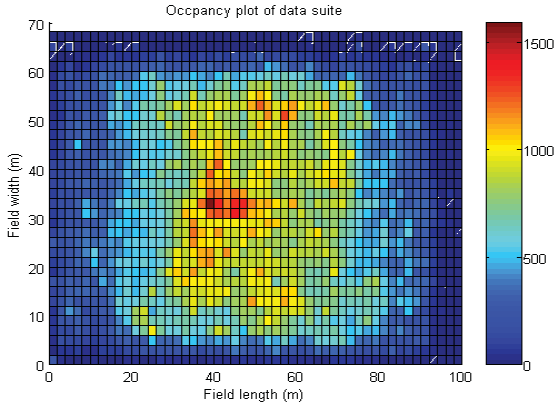

Figure 7.2: Occupancy plot of the data suite.

This plot shows the majority of players occupying an area of the field near the centre, typical of a game where neither team has dominance. The individual occupancy signatures for each of the 10 games in the data pool can be viewed in Appendix B: Plots.

7.2 Calculating visibility

Situations that cause a loss of visibility were inferred from hypothesis and results of the system characterisation. Calculations of visibility were performed between each tag and receiver in every frame. The player was invisible to the tracking system if the number of receivers visible to a tag was less than three.

The model returns an array indicating which player positions cannot be calculated due to invisibility. The returned array can be mapped directly to the input data.

To avoid redundant calculations, the model has a hierarchical structure. The particular receiver must sequentially satisfy the various criteria for visibility, beginning with that criterion producing the highest discrimination and requiring the least computation. Should a particular receiver fail to satisfy a criterion, the tag is deemed invisible and execution immediately moves to the next receiver. If the number of visible receivers for a tag drops below three, the position of that tag cannot be calculated.

The criteria for visibility (in order of precedence) are: range, transmission angle and player obstruction.

7.2.1 Range

A tag is invisible if the straight-line distance between the tag and the receiver is greater than the range of that receiver.

For simplicity of coding, it was assumed that the reception envelope of all receivers was a semicircle that had the bore sight of the receiving antenna pointing towards the centre of the playing field. Thus the value for range given to the visibility model was a radius of this semicircle. For the model described here, that radius was 100 m. This value was chosen as a conservative estimate of the range of the receiver from the receiver-lobe-pattern tests (the range at ±90° from bore sight of a mid gain receiver).

The calculation of range is expressed in equation (7.1) below. The result of this equation is a Boolean representing visibility (i.e. 0 denotes invisible and 1 denotes visible).

$$V_{range}(f,p,r) = \begin{Bmatrix} 0, if\left(\sqrt{\left(Px_{(f,p)} - X_r\right)^2 + \left(Py_{(f,p)} - Y_r\right)^2} > R\right) \\ 1, if\left(\sqrt{\left(Px_{(f,p)} - X_r\right)^2 + \left(Py_{(f,p)} - Y_r\right)^2} \leq R\right) \end{Bmatrix} \quad (7.1)$$

Where:

- $V_{range}(f,p,r)$ is the Boolean value for visibility by range for a given player p and receiver r pair within frame f
- $Px_{(f,p)}$ and $Py_{(f,p)}$ are the coordinates of the player. That player is addressed within the input data array by the frame f and the player number p
- X_r and Y_r are the coordinates of the receiver r
- R is the maximum reception range of the receivers.

7.2.2 Transmission angle

The angle criterion compares the transmission arc of the tag with the position of the receivers. If they are aligned then it is considered that this line-of-sight condition is met. Figure 7.3 shows a player whose tag is within the angle of one receiver but out of angle with the other.

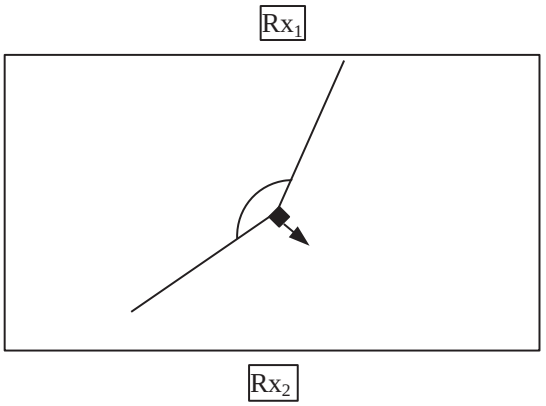

Figure 7.3: Example of Out of Angle criterion.

Here the player, who is facing the direction of the arrow, is oriented such that the tag on his back will be visible to receiver 1 (Rx_1), but not visible to receiver 2 (Rx_2). The transmission angle criterion was met if the transmission arc of the tag included the receiver position, given that the range criterion had already been satisfied and the receivers had a semicircular reception pattern.

The transmission arc of the tag was variable within the model to allow for different transmission arcs that result from either different tag antennae or placement of the tag on the player.

It was assumed that the tag was on the back of the player's leg (best location as suggested by the tag placement trials); therefore the transmission bore sight was in the opposite direction to that player's orientation, however, this is also variable within the model.

The original hypothesis was that the transmission arc around a player would be the sum of the arcs of individual tags worn by that player. For example, if the player had one tag on his front and one on his back, and each transmitted with an arc of 180°, then the total transmission arc would be 360°. Unfortunately, further analysis illuminated two critical flaws in this hypothesis.

The receivers will see each tag independently. Using the above example, if the tag on the back

is visible to only one receiver and the tag on the front is visible to only two receivers then a position cannot be calculated since, independently, the tags do not have visibility with three or more receivers.

The other flaw is that since the tags transmit on a pseudo random basis, it is improbable that they will consistently transmit together, so it will be impossible to combine and make use of the TOF measurements from both tags to get a single position.

Despite these flaws, multiple tags do increase the probability that one of the tags is visible to three or more receivers. Further processing can be applied to reconcile the positions of multiple tags on a single player after tracking has been completed.

The calculation of transmission-arc visibility is shown in equation (7.2) below. The result is a Boolean indicating invisibility.

$$V_{angle}(f,p,r) = \begin{Bmatrix} 0, if\left(\theta_{i(f,p)} > \theta_r\right) \\ 1, if\left(\theta_{i(f,p)} \leq \theta_r\right) \end{Bmatrix} \tag{7.2}$$

Where:

- $V_{angle}(f,p,r)$ is the Boolean value for visibility by transmission-arc for a given player p in frame f coupled with receiver r
- $\theta_{i(f,p)}$ is the incident angle of the centre of the transmission from the player to the receiver. This can easily be generated by applying a bias to the player orientation data (the bias will be determined by the placement of tags, in this case 180°)
- θ_r is the acceptance angle of the receiver adjusted to account for the receiver orientation within the global coordinate system (assumed here to be towards the centre of the field).

7.2.3 Player obstruction

This invisibility mode occurs when another player passes through the line-of-sight path between transmitter and receiver.

This section of the model was concerned with excluding players who were not candidates for causing obstruction; if any players were left at the end, then an obstruction occurred and the tag was invisible.

The first part of the exclusion process was to define a bounding box with the tag and receiver at opposite corners. Any other players within this box are obstruction candidates; all other players were removed from consideration. If the tag were in line with the receiver in either the x or the y dimension, then either the width or the height of the bounding box, respectively, would be zero, which would mean that only players exactly on the connecting line would be considered as candidates for obstruction. This situation was unacceptable, as there may be other players near this box who should be candidates that will be missed. To alleviate this issue, a lower limit of 0.3m was imposed on both the height and the width of the bounding box. Figure 7.4 illustrates a standard bounding box and one that has saturated to its lower limit. The solid dot is the player of interest, Rx_1 is the receiver of interest and the circles are candidate players.

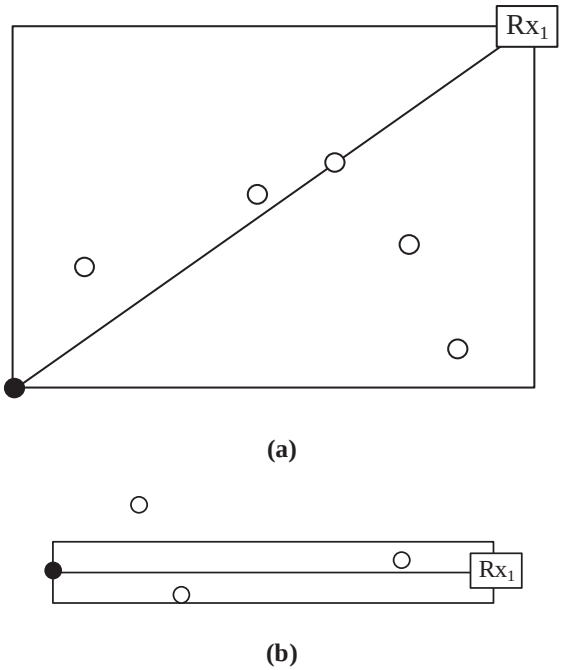

Figure 7.4: (a) A standard bounding box. (b) A bounding box that has been limited in size to account for the player of interest being in line with the receiver.

The one player outside the box in Figure 7.4(b) can be ignored because he is outside the bounding box and therefore is not going to cause obstruction.

A search of the current frame of player positions will yield those players whose current positions lie within the defined bounding box. A subset of candidate player is created from the dataset; all other players are discarded after this calculation. Equation (7.3) defines the criteria for a player being considered a candidate for obstruction.

For equation (7.3):

- $\omega(f,p,r,o)$ is a Boolean that indicates if player o has a possibility of obstructing line of sight (LOS) between player p and receiver r in frame f. If the value is true then that player is considered to still be a candidate for obstruction.

Note that in equation (7.3) the range that player o must be in is dependent on the position of player p with respect to receiver r. The position with respect to one another defines the upper and lower bounds of the bounding box. Note also that the limitations on the size of the bounding box apply.

The shortest distance to the straight line between the receiver and tag was calculated for each candidate player. This is a line normal to the straight line between tag and receiver (as illustrated in Figure 7.5).

$$\omega(f,p,r,o) = \begin{Bmatrix} 0, if\left(Px_{(f,p)} \le Px_{(f,o)} \le X_r || Py_{(f,p)} \le Py_{(f,o)} \le Y_r\right) \&\& \left(Px_{(f,p)} < X_r\right) \&\& \left(Py_{(f,p)} < Y_r\right) \\ 1, if\left(Px_{(f,p)} \le Px_{(f,o)} \le X_r \&\& Py_{(f,p)} \le Py_{(f,o)} \le Y_r\right) \&\& \left(Px_{(f,p)} < X_r\right) \&\& \left(Py_{(f,p)} < Y_r\right) \\ 0, if\left(\overline{X_r \le Px_{(f,o)} \le Px_{(f,p)}} || \overline{Py_{(f,p)} \le Py_{(f,o)} \le Y_r}\right) \&\& \left(Px_{(f,p)} > X_r\right) \&\& \left(Py_{(f,p)} < Y_r\right) \\ 1, if\left(X_r \le Px_{(f,o)} \le Px_{(f,p)} \&\& Py_{(f,p)} \le Py_{(f,o)} \le Y_r\right) \&\& \left(Px_{(f,p)} > X_r\right) \&\& \left(Py_{(f,p)} < Y_r\right) \\ 0, if\left(\overline{Px_{(f,p)} \le Px_{(f,o)} \le X_r} || \overline{Y_r \le Py_{(f,o)} \le Py_{(f,p)}}\right) \&\& \left(Px_{(f,p)} < X_r\right) \&\& \left(Py_{(f,p)} > Y_r\right) \\ 1, if\left(Px_{(f,p)} \le Px_{(f,o)} \le X_r \&\& Y_r \le Py_{(f,o)} \le Py_{(f,p)}\right) \&\& \left(Px_{(f,p)} < X_r\right) \&\& \left(Py_{(f,p)} > Y_r\right) \\ 0, if\left(\overline{X_r \le Px_{(f,o)} \le Px_{(f,p)}} || \overline{Y_r \le Py_{(f,o)} \le Py_{(f,p)}}\right) \&\& \left(Px_{(f,p)} > X_r\right) \&\& \left(Py_{(f,p)} > Y_r\right) \\ 1, if\left(X_r \le Px_{(f,o)} \le Px_{(f,p)} \&\& Y_r \le Py_{(f,o)} \le Py_{(f,p)}\right) \&\& \left(Px_{(f,p)} > X_r\right) \&\& \left(Py_{(f,p)} > Y_r\right) \end{Bmatrix} \quad (7.3)$$

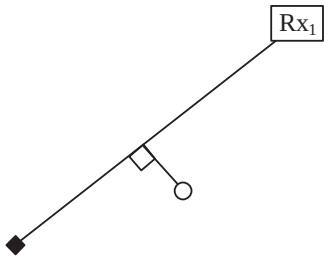

Figure 7.5: Finding the shortest distance between possible obstruction player and the connecting line.

This shortest distance is the same as the distance between intersection and potential obstructing player of the line that is normal to the direct line between current player and receiver. The equations to find this distance are given below. Note that η is used in these equations as a substitution to simplify the equations.

$$d_{path}(x_l, y_l, \theta_l, x_p, y_p) = \sqrt{(\eta - x_p)^2 + (\tan\theta_l(\eta - x_l) + y_l - y_p)^2} \quad (7.4)$$

$$\eta = \frac{x_p - \tan\theta_l\,(y_l - y_p - \tan\theta_l)}{1 + (\tan\theta_l)^2} \quad (7.5)$$

Figure 7.6 shows each of the variables to the above equations.

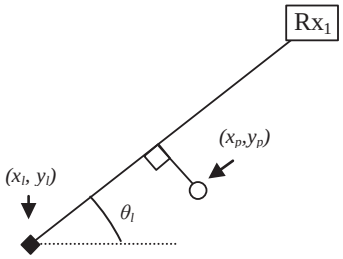

Figure 7.6: Annotated diagram to indicate variables in equations (7.4) and (7.5)

It was assumed that all players are cylinders with a radius of 0.3 m, and that their centre is the position listed for them in the CSV data file, leading to the condition that the shortest distance between possible obstructing player and the connecting line, described above, must be less than 0.3m for that player to avoid being excluded. It was assumed that no significant signal diffraction occurs, and incursion upon this path by another player will result in obstruction.

Even though only the plane of the field of play (henceforth referred to as the horizontal plane) is of interest for positioning, the height of the receivers compared to the tags must be taken into account when performing player obstruction analysis. The receivers are all on stands that raise them 2.5 m above the ground, the tags are assumed to be on the back of the legs of the players approximately 0.3 m from the ground and each player is 1.8m tall.

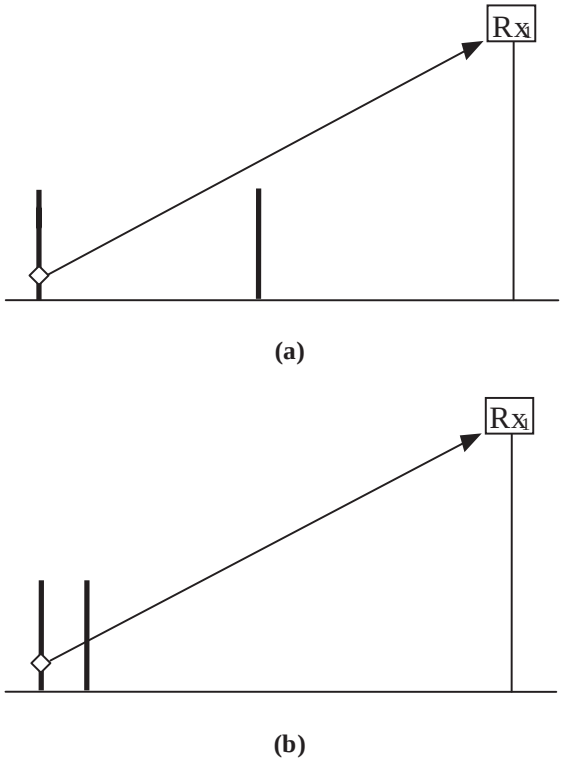

Figure 7.7: (a) No obstruction in the z plane. (b) Obstruction in the z plane.

Depending on the position of the obstructing player with respect to the tag and receiver in the horizontal plane, there may not be obstruction because of the geometry in the vertical plane. Figure 7.7 shows the case of a potential obstructing player in the vertical plane where the tag is obstructed and not obstructed respectively. In both cases the potential obstructing

player was a candidate for obstruction in the horizontal plane.

There is an obstruction in the vertical plane when the value of the equation defining the LOS path in the vertical plane is less than the height of the potentially-obstructing player but greater than the height from the ground to mid thigh at the distance of that player to the player of interest. It was shown in the system characterisation that the signal from the tag could pass the legs of a player when the tag is located on the back of the leg.

In the scope of the player obstruction, this means that a simple exclusion zone can be applied. If there are any candidates for obstruction within range of the player then they will cause obstruction. The specific range will depend on the distance between the player of interest and the receiver. The equations below show how to calculate obstruction in the vertical plane based on this. Once again, a true result denotes visibility while a false denotes invisibility.

$$V_{vert}(f,p,r,o) = \begin{Bmatrix} 0, if\, P_l < \frac{Z_r - T_z}{d_r} d_o + T_z < P_z \\ 1, if\, P_l < \frac{Z_r - T_z}{d_r} d_o + T_z < P_z \end{Bmatrix} \quad (7.6)$$

$$d_r = \sqrt{\left(Px_{(f,p)} - X_r\right)^2 + \left(Py_{(f,p)} - Y_r\right)^2} \quad (7.7)$$

$$d_o = \sqrt{\left(Px_{(f,p)} - Px_{(f,o)}\right)^2 + \left(Py_{(f,p)} - Py_{(f,o)}\right)^2} \quad (7.8)$$

Where:

- $V_{vert}(f,p,r,o)$ is a Boolean indicating if obstruction has occurred in the vertical plane. Only values of p that have not been excluded by previous test are included.
- Z_r is the height of the receiver.
- T_z is the height of the tag on the player
- d_r is the distance between the player p and receiver r
- d_o is the distance between the player p and player o
- P_z is the height of all of the players.

Thus, visibility with respect to player obstruction can be summarised in the following equation. Note that this equation is only applied to those players that produce a true value in equation (7.3)

$$V_{obs} = \begin{Bmatrix} 0, if\left(d_{path} \le P_{rad}\right)\&\&(V_{vert} = 1) \\ 1, if\left(d_{path} > P_{rad}\right)||(V_{vert} = 0) \end{Bmatrix} \quad (7.9)$$

Where:

- V_{obs} is a Boolean indicating a loss of visibility due to player obstruction
- P_{rad} is the radius of a player.

As discussed, all players in the input data have an upright stance. That assumption is the most detrimental in this section of the model. Consider a scrum formation: If the tags were on the backs of the players, then when the players bend down for a scrum, their tags will be facing up and will be unobstructed by the players around them. However, because of their upright stance, the model will most likely show that the players at the perimeter of the scrum are obstructing the players in the centre. There was very little that could be done about this with the data available. At the time of programming, it was best to suggest that this model is conservative and that it will over estimate the level of player obstruction. Further development of the model may allow an estimate of stance to be made based on the relative locations of players and a known behaviour model for rugby.

Selective placement of the receivers and tags is the best way to avoid obstruction from other players.

7.2.4 Combined visibility model

It is now possible to define a meta-equation for the visibility model as a whole. Equation (7.10) brings together all of the checks for visibility discussed in previous subsections. The model will output the number of receivers that are visible to a given player at any time.

$$V_{comb}(f,p,r) = \begin{Bmatrix} 0, if \overline{V_{range}(f,p,r)} || \overline{V_{angle}(f,p,r)} || \overline{V_{obs}(f,p,r)} \\ 1, if V_{range}(f,p,r) \&\& V_{angle}(f,p,r) \&\& V_{obs}(f,p,r) \end{Bmatrix} \quad (7.10)$$

Where:

- $V_{comb}(f,p,r)$ is the overall visibility of a receiver r to player p in frame f. This value can be used either as a Boolean or an incremental value.

Thus the total visibility for any given data set can be calculated from equation (7.11). Note that players visible to less than three receivers are shown to have no visibility since no useful position can be calculated with less than three receivers

$$V_{total} = \sum_{f=1}^{f_{max}} \sum_{p=1}^{p_{max}} \begin{Bmatrix} 0, if \sum_{r=1}^{r_{max}} V(f,p,r) < 3 \\ \sum_{r=1}^{r_{max}} V(f,p,r), if \sum_{r=1}^{r_{max}} V(f,p,r) \geq 3 \end{Bmatrix} \quad (7.11)$$

The mathematical mode of failure is recorded for those players whose positions cannot be calculated. This allows analysis of the causes of failures which will lead to a better understanding of the environment and the potential to develop a more efficient tracking system.

The decision tree shown in Figure 7.8 explains how the visibility-model algorithm executes. There are a number of modifications made over the pure mathematical form of finding visibility (e.g. checking for visibility for a given player stops when the number of visible receivers is less than three). These were made to reduce execution time for the algorithm.

The green boxes in the diagram are used to describe loops in the execution (in a similar fashion to the Labview development environment). The condition of the loop is given at the top centre of each box. In addition, the dots at the intersection of an arrow and the box denote how the flow of execution is handled. Green dots denote the execution entering or exiting the loop in the normal way (i.e. exit on a green loop will induce another iteration if applicable), red dots indicate an immediate break from the indicated loop. There is one red dot in Figure 7.8 that indicates that the receiver loop is broken if the number of visible receivers drops below three.

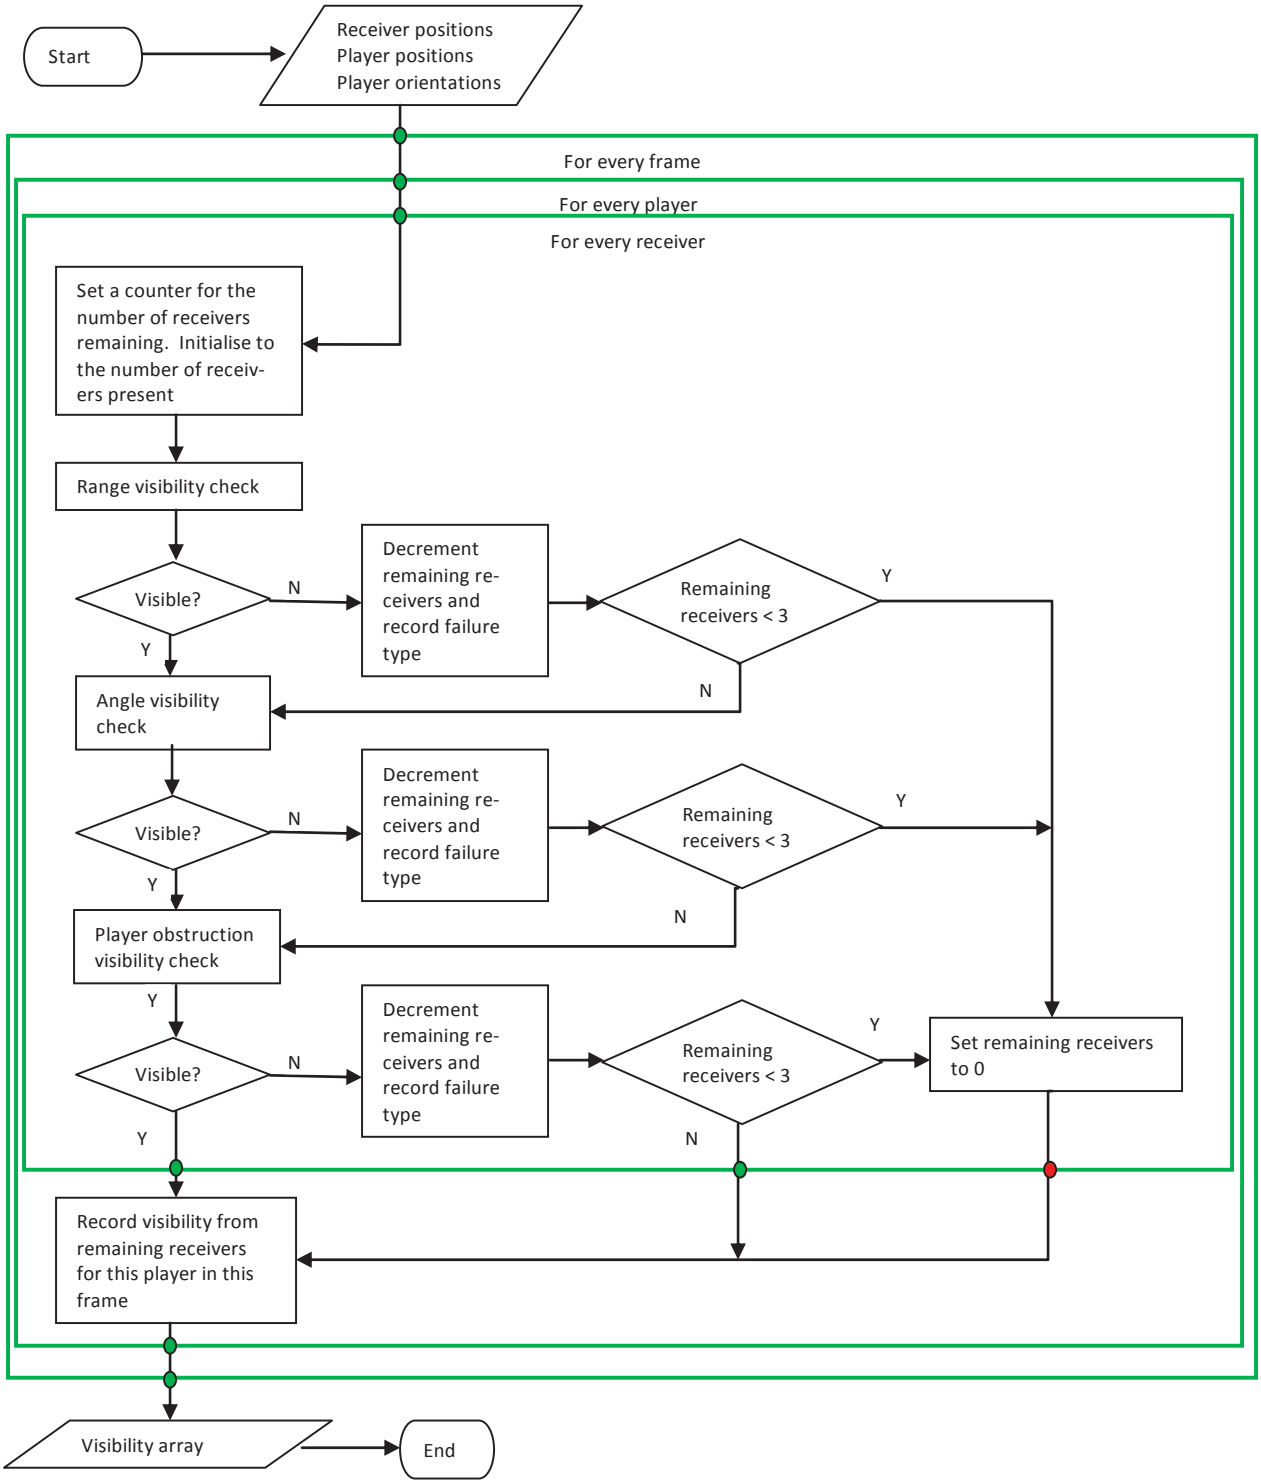

Figure 7.8: Decision tree for visibility model execution.

7.2.5 Static obstruction modelling

In its current form, the visibility model will be suitable for use in any sport played on a regular shaped and open field. Other applications would be in the planning of the configuration of a tracking system in a warehouse, industrial environment or a sport with more complex environment (such as orienteering). In these applications it is necessary to consider objects that are not tracked. These obstructions are most likely to be static and may include walls, shelves or other buildings. In order to handle these obstructions, the model is provided with a map of the area marking the static obstructions (referred to as the 'terrain map'). A ray tracing method can then be used to determine if there is visibility. The map contains information that can be related to the likelihood of losing visibility. This information might be height of the obstruction or a symbol representing the material from which it is made (and thus its permeability), however this encoding can be modified to suit the situation.

In the case of a sport like orienteering, the topography of the landscape will be important too. Courses, more often than not, contain significant inclines and the model must be extended to take account of land formations as another form of obstruction. A land formation will either change the height of a receiver or it will introduce a land obstacle to LOS. The overall height of the receiver is calculated by adding the receiver stand height (Z_r in the open field model) to the terrain height to get the absolute height of the receiver. Land obstructions are handled similar to other static obstructions; if the ray tracing intersects the terrain then there is obstruction.

Another difference in orienteering situations is that the occupancy of the area of interest will certainly not be uniform; a competitor will follow the path of least resistance between waypoints. The input data would therefore need to reflect this behaviour. This concept can also be applied to warehouse situations; there will be avenues and walkways that have increased traffic with regard to other areas. The occupancy of the terrain can be encoded into the terrain map and the model will only look to these areas when performing calculations. Redundant calculations can be avoided in this fashion. Also, in an optimisation situation, areas where there will be little or nothing to track are less significant than areas of high traffic.

In both the warehouse and sport situation the terrain map is a bitmap image. The locations of pixels map the area of interest at a given resolution. A grey scale image would allow for one variable to be encoded at each position. This variable could be the height of the terrain or an obstruction. A colour image would allow a considerably greater number of variables to be encoded. Such variables could include:

- The height of an obstruction or terrain,
- Its composition, and
- Allowable locations for receivers to be placed.

The bitmap format is easy for a user to generate and analyse, as the image would directly correlate with the situation and no special training in visualisation would be required. A bitmap of this type could be obtained from either a scan of a floor plan or generated from scratch in an image editor such as MS Paint or Photoshop.

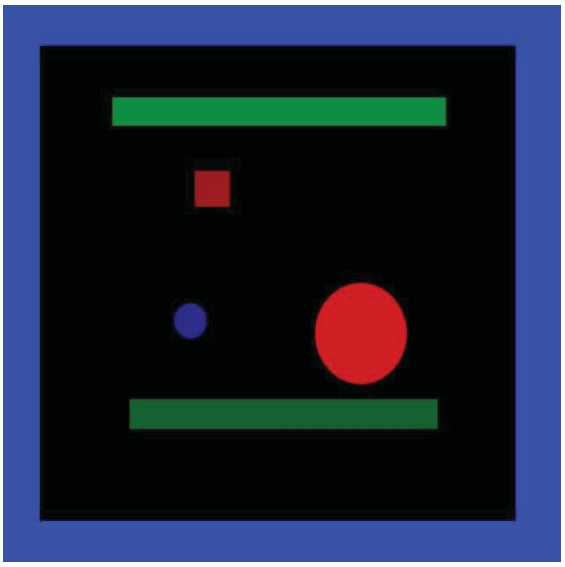

Figure 7.9: Example of a bitmap image used for visibility analysis of complex environments.

Figure 7.9 is an example of such a bitmap. An RGB scheme was used to encode data. Black represents free space, occupied by nothing of interest; blue represents the allowable locations for the receivers; red represents static obstructions that are made of a material that will absorb completely the signals from the tags, and green represents static obstructions that will only partially attenuate the signals from the tags. The saturation of the colour represents the height at which the object is located. In the case of blue colours it represents the height at which the receiver may be mounted and for red and green it represents the height of the obstruction.

In this map it can be seen that most of the space in the area of interest contains nothing of interest (black). Receivers may be placed high up in the band surrounding the area and one cylinder within the area but at a lower height (blue). There are two types of obstruction present in the map, two long rectangular regions constructed of a material that is semi permeable by RF transmissions (green) and two small regions that will completely block the signal (red). Note also that these obstructions can be seen to be at different heights based on the saturation of the coloured regions.

To illustrate how easy it is to generate a terrain map consider the following situation. An orienteering competition is held at Massey University's Turitea campus in Palmerston North, New Zealand. The competitors will move through the centre of the campus and will need to be tracked. Assume the following: that only buildings will cause obstruction to the tracking system; that buildings are complete obstructions to the system. It is also known that the terrain is flat over the area of interest. Using any piece of image processing software (e.g. Photoshop), a map of the area of interest is imported (this map is available on the Massey University web site). The area of interest of this map is shown in Figure 7.10.

Figure 7.10: Raw map of the area of interest.

Next, by adding another layer the user can essentially trace over the desired features with the appropriate hues and saturations to encode the data. Then the raw image is removed leaving a useable terrain map. In this case, only the buildings are of interest and they are to be encoded as completely opaque to RF signals (red at the highest saturation). Thus Figure 7.11 is the terrain map. This map is saved as a 24-bit RGB bitmap. Bit maps are used here because compression may blur the edges of the features (for example, JPEG would blur edges because it is a lossy compression technique). Additionally, the image in this case is only 156kB and thus does not require compression.

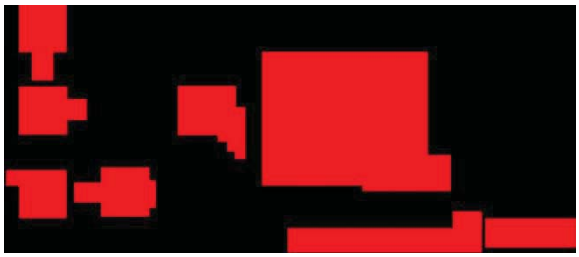

Figure 7.11: The resulting terrain map.

Note that the large building in the centre-right of the image has been coded as one solid building. This is because that complex contains an enclosed courtyard which would be inaccessible to the competitors. It is this type of extra information that is required when making the terrain map.

This terrain map can be used to calculate the performance of a tracking system based on a number of criteria. The most useful are coverage (either across the whole area or of the known track of the objects if those data are available) and the GDOP of measurements. The same equations from the open-field model can be used to make these calculations with only minor modifications.

In order to calculate receiver coverage it is first necessary to know if the point of interest is in range of the receiver (regardless of obstructions). Secondly it is necessary to know if the LOS path is obstructed. For the terrain map shown in Figure 7.11 an obstruction is defined as red colours. Finding the colour of a pixel is a matter of looking at the first byte of the colour encoding of the terrain map. The equations to find if a receiver has LOS to a given point are shown below.

$$V_{range}(x,y,r) = \begin{Bmatrix} 0, if\sqrt{(x-X_r)^2+(y-Y_r)^2} > R \\ 1, if\sqrt{(x-X_r)^2+(y-Y_r)^2} \leq R \end{Bmatrix} \quad (7.12)$$

Where:

- V_{range} is the visibility of a given location (x,y) to a receiver r considering only range.
- All other symbols have the same meaning as in equation (7.1).

This equation is similar to equation (7.1) except that instead of examining a given player in a given frame it looks at a given point within the area of interest.

$$V_{obs}(x,y,r) = \begin{Bmatrix} 0, if\left[\mathrm{T}_{(\chi,\psi,1)} \geq \mathrm{A}\right]_s^e \\ 1, if\left[\mathrm{T}_{(\chi,\psi,1)} < A\right]_s^e \end{Bmatrix} \quad (7.13)$$

Where:

- V_{obs} is the visibility of a given location (x,y) to a receiver r considering only static obstructions.
- $T_{(\chi,\psi,1)}$ is the red value of the pixel at point (χ,ψ) in the terrain map. The *1* in the index indicates that the red byte is of interest.
- A is the altitude threshold for an obstruction. This value is a scaling of real-world altitudes onto the resolution of the bitmap and can be changed to suit the situation.
- s and e are the start and end point of the ray respectively. The origin of the ray s is the current location of interest (x,y) and the end, e, is location of the current receiver (X_r,Y_r). The ray is defined in polar coordinates and must be converted to the Cartesian coordinates in order to index the terrain map.
- Note that the bracketed range notation is used in this equation to indicate that each step along the ray is to be considered separately. All pixels along the ray

must be less than the threshold for there to be visibility.

Thus coverage at a given point can be defined as follows.

$$V_{comb}(x,y,r) = \begin{Bmatrix} 0, if \overline{V_{range}} || \overline{V_{obs}} \\ 1, if V_{range} \&\& V_{obs} \end{Bmatrix} \quad (7.14)$$

Where:

- V_{comb} is the combined visibility when considering range and static obstruction.

A summation of combined visibility at a given point for all receivers will yield the coverage at that point by the tracking system and thus will show where positions of objects can be calculated.

If the entire area of the terrain map is to be analysed then the model will consecutively calculate the visibility at each point. However, another colour in the terrain map can be used to concentrate the model on a subset of the map (an example of this is given in Figure 7.12; the blue pixels show the area of interest). The use of 'track data' within the terrain map can be used to reduce computation time and to focus optimisations to the required areas.

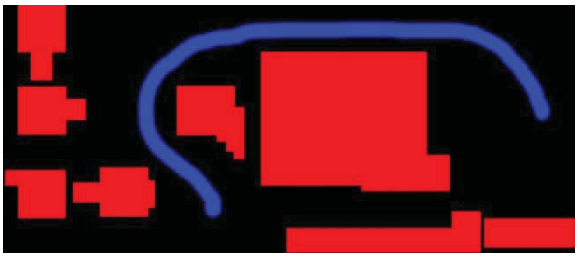

Figure 7.12: Terrain map with track data included.

The algorithmic operation of this model can be summarised in the decision tree shown in Figure 7.13. Note that the section ringed in blue is optional depending on the requirements of the analysis. If it is not required then assume that the 'yes' branch is always taken at the 'Within area of interest?' question. Similarly, once the coverage is known then GDOP can be calculated if required. GDOP is calculated using the equations shown in appendix: A, section 13.1. It is calculated using only those receivers visible to a point. Thus if there is unfavourable geometry because one of the receivers is occluded then the GDOP will be high.

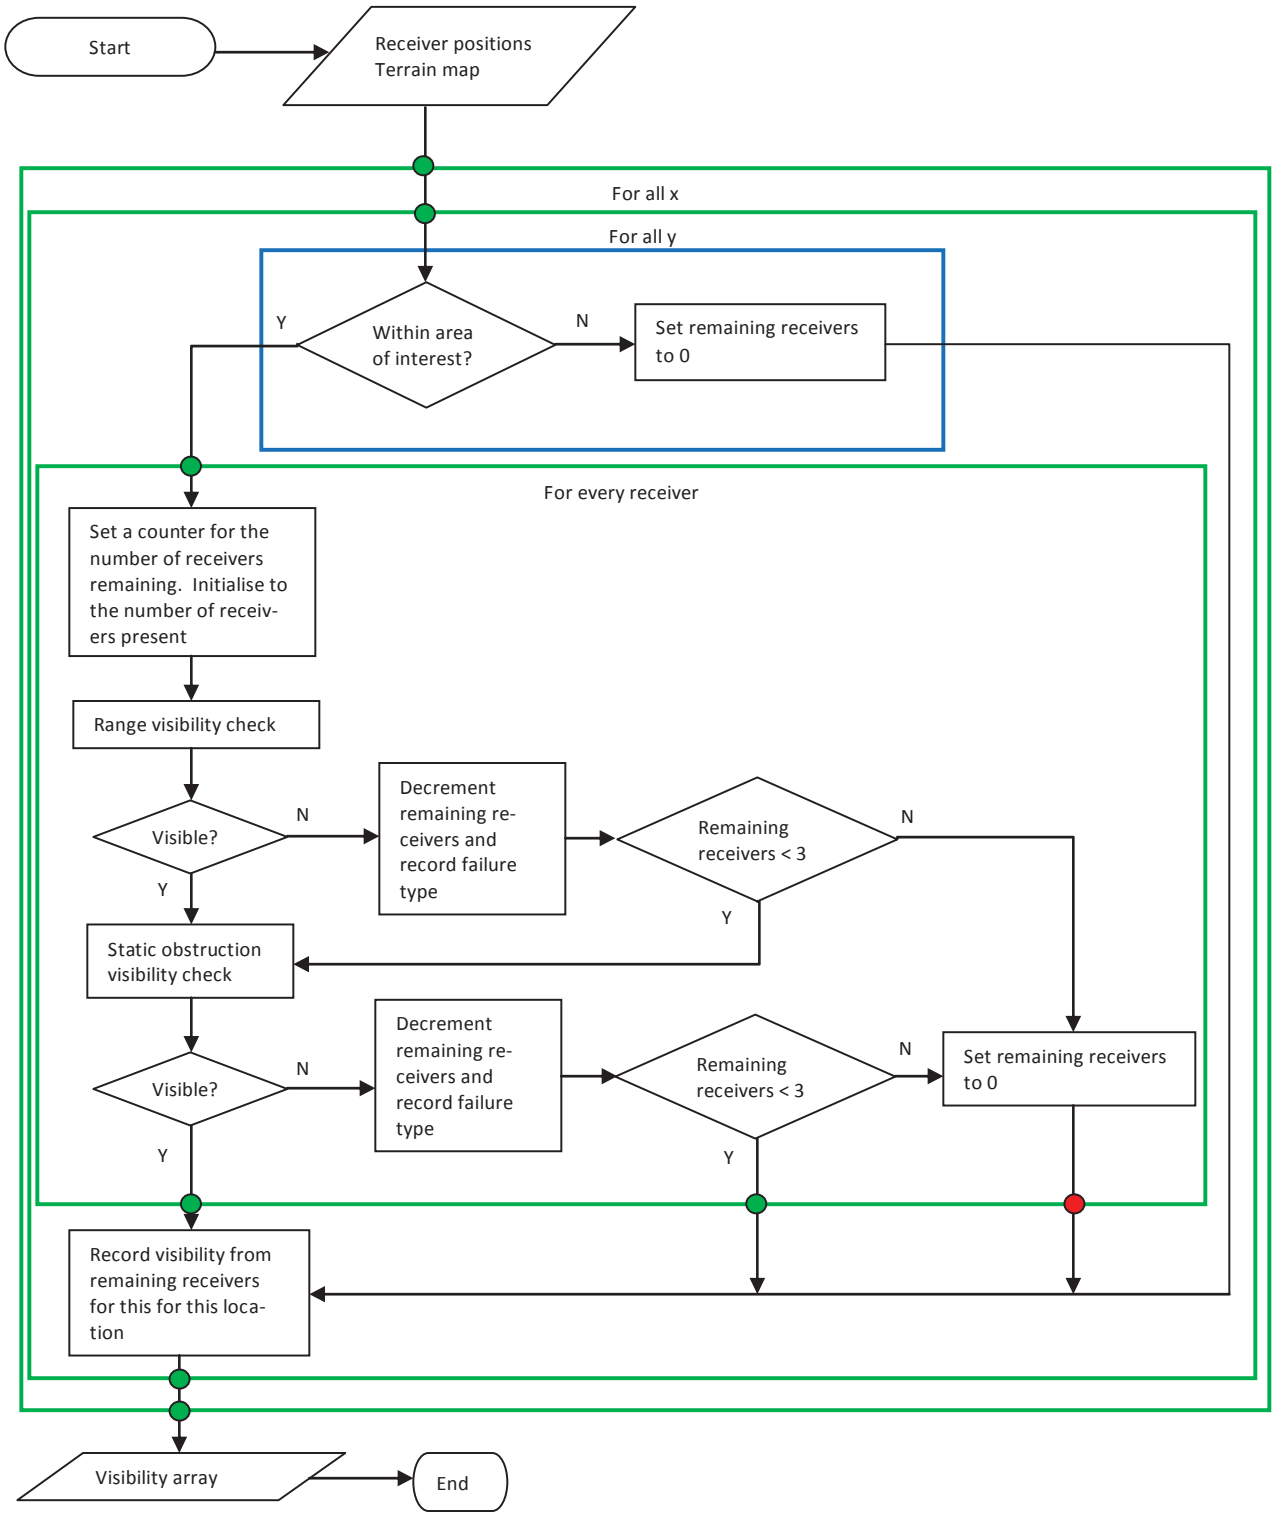

Figure 7.13: Decision tree for static obstruction model.

7.3 Visibility model results

Once the visibility model had been verified, it was run using the pre-tracked input data. The output of the analysis was an array that mapped to the player positions for each frame.

The model was used to analyse configurations of 6, 8 and 10 receivers. On average the model took 10 minutes to analyse a rugby game containing 3046 frames. In the input data, this is the equivalent of 80 minutes of game plus a few minutes overtime. For the three different receiver configurations, the percentage loss of visibility was 36%, 26% and 18% (6, 8 and 10 receivers respectively). This indicates that increasing the number of receivers will increase the visibility of tags to the tracking system. An animation of the game data after application of the player visibility model is included on the attached DVD (see Appendix C: DVD index for the path). The animation shows where each player becomes invisible and the type of invisibility encountered. The frame number is provided for reference in the legend.

Figure 7.14 shows a typical frame of data with two full teams in play. Here each receiver is 2.5 m off the ground and has a range of 100 m. The tags are worn on the back of the players' legs and are assumed to have a transmission arc of 180°. The players themselves are assumed to be 1.8 m tall and are wearing the tags at a height of 0.3 m from the ground.

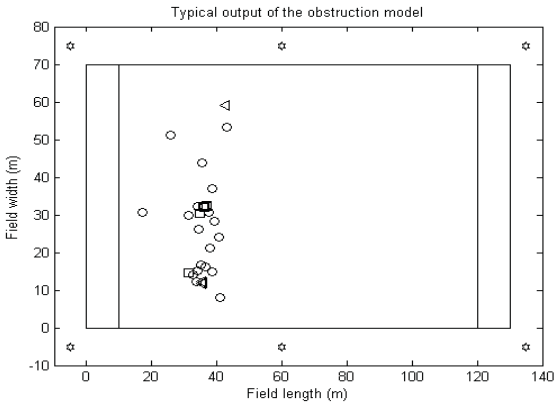

Figure 7.14: Typical output of the analysis model.

A common symbology is used here to represent the different visibility types, summarised below (note that unless stated, the symbols are the same for both teams):

- Circles denote valid positions,
- Crosses are tags invisible due to range issues (none present in this figure),
- Triangles are tags invisible due to angle issues,
- Squares represent player obstructions, and
- Hexagrams denote the placement of the receivers.

To aid in the distinction of players from different teams, the triangle symbol points in the direction of attack of that player's team. However, they do not represent the orientation of the player.

Figure 7.14 shows that players in the open are visible, while players in tight formations are invisible. This frame is typical of the entire analysis.

Closer examination of this frame reveals that there are two major clusters of players. As expected, some players in these formations are invisible. Figure 7.15 narrows the field of view within this frame to the cluster around *(35, 30)*.

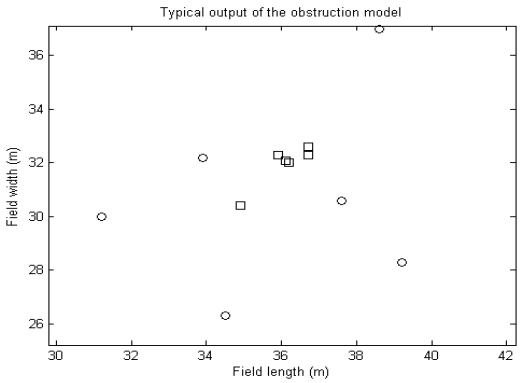

Figure 7.15: Close up of a clustered formation.

Here, it can be seen that, where players are tightly clustered, they suffer from player obstruction invisibility. The same can be said of the player near the position *(35, 30)* despite the fact that there are no players close by.

Figure 7.16 gives a comparison of the frequency of the different visibility failure modes in the analysed data.

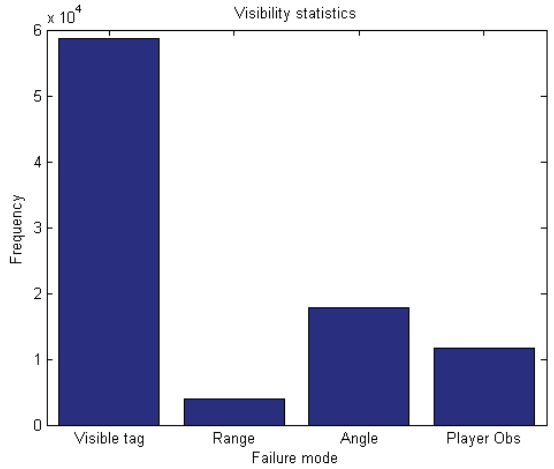

Figure 7.16: Bar plot showing the visibility statistics of an analysed game of rugby.

Here it can be seen that the orientation of the players with relation to the receivers is the cause of the greatest number of visibility failures, followed by player obstruction and then range. This differs from the hypothesised ranking of invisibility causes, suggesting that increasing the transmission arc of the tags will yield the greatest increase in tag visibility.

Another useful measurement to examine is the length of time tags are invisible at any one time. Figure 7.17 shows a histogram of lengths of contiguous visibility failures.

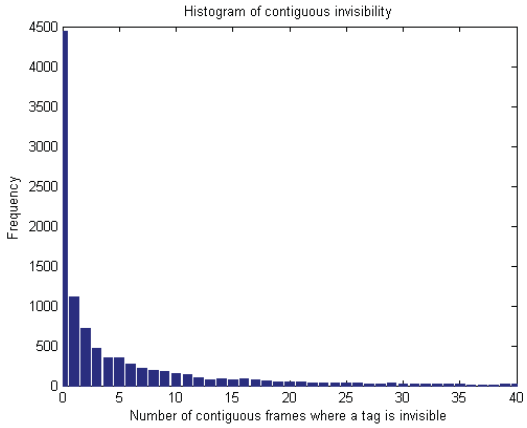

Figure 7.17: Histogram of contiguous invisibility.

There is an approximate inverse exponential trend in invisibility intervals in Figure 7.17, suggesting that in most cases, a loss of visibility is transitory. The longer contiguous invisibility events occur when a tight formation (especially a ruck) has a prolonged duration; it is possible for a player to remain inside a formation for approximately 40 seconds (based on the above figure).

7.3.1 Static obstruction model results

This model was verified by using a number of terrain maps of varying complexity. The first is shown in Figure 7.18. This terrain map has a single, round obstruction in the middle that cannot be penetrated by RF signals.

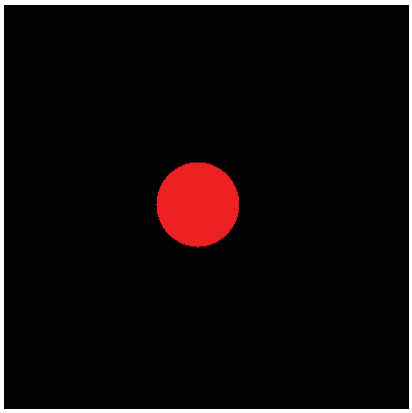

Figure 7.18: Singe-dot obstruction terrain map.

Figure 7.19 shows how the static obstruction affects the coverage of the area when using 6 receivers (shown as black dots). Figure 7.19 (a) shows the coverage when the receivers are clustered in one quarter of the terrain map and (b) shows the effect of spreading the receivers evenly around the obstruction (the configuration that will give the best GDOP in an open field situation). Note that coverage has a lower limit of three receivers since coverage by less than three receivers will not yield a useful position. The colours on these plots represent the number of receivers covering an area where a position will be calculated. Dark blue denotes an area where no position can be calculated.

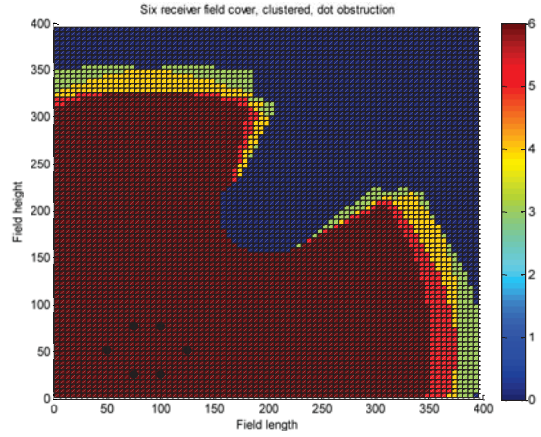

(a)

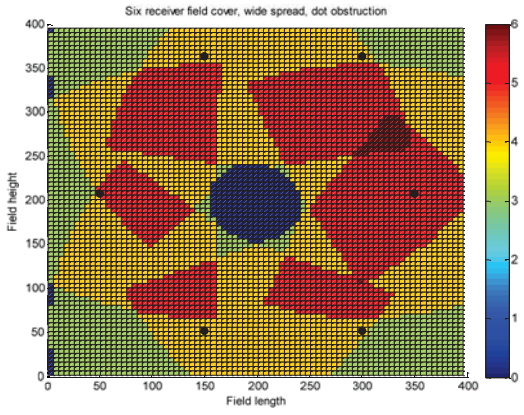

(b)

Figure 7.19: 6-receiver coverage of the single dot area. (a) clustered and (b) spread.

Here it can clearly be seen that the obstruction is adversely affecting the coverage of the area of interest. Even when the receivers are spread evenly around the obstruction there are areas of lower coverage. This suggests that even a minor obstruction will hamper the performance of a tracking system and thus careful planning is required to get optimum performance from a minimalist tracking system.

In Figure 7.20 the GDOP of the above receiver configurations are shown (the colours now represent the GDOP value of a given area). Note that where a position can not be calculated (due to low coverage) a penalty value of GDOP has been applied (i.e. low accuracy).

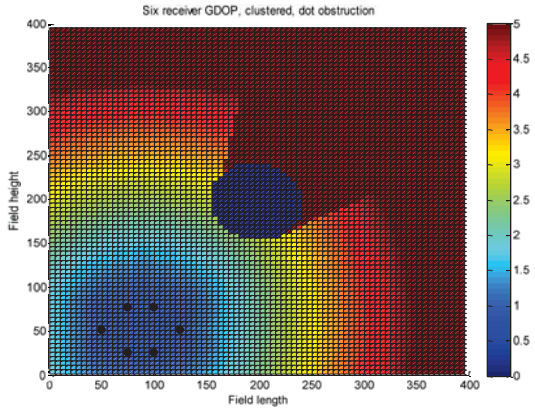

(a)

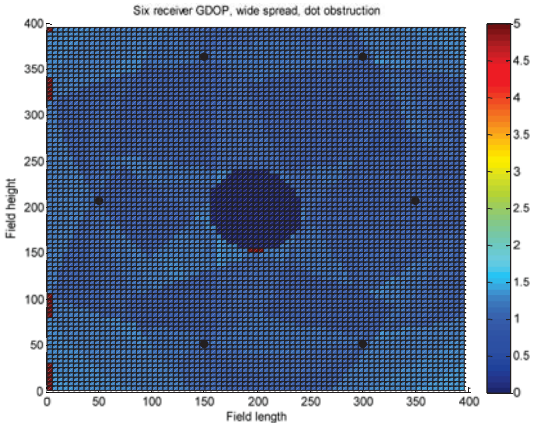

(b)

Figure 7.20: 6-receiver GDOP of the single dot area. (a) clustered and (b) spread.

It can be seen from these plots that the lowest GDOP occurs where there is the most cover within the formation of the receivers. Figure 7.20 (a) especially shows that moving away from the receiver formation will reduce GDOP.

The next terrain map is shown in Figure 7.21 and consists of four, square obstructions that are clustered together near the centre of the map. These obstructions take up approximately 25% of the total area of the map. The challenge with this terrain map is providing coverage of both the centre and the outside of the map

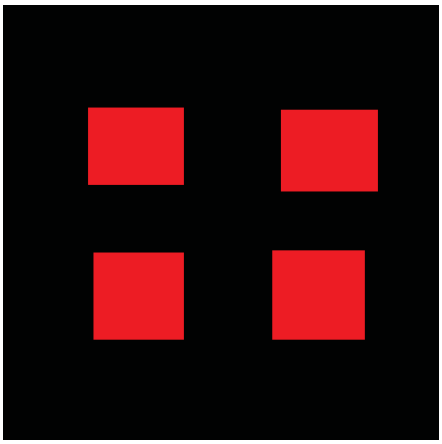

Figure 7.21: Close obstruction terrain map.

Figure 7.22 shows the field coverage (colour axis) for an arbitrary placement of receivers (again shown as black dots). It can be seen that these clustered obstructions prevent the receiving envelops of many of the receivers from overlapping. This is epitomised by the receiver at *(200,375)* which has no coverage near itself (This receiver is difficult to see in the plot and so has been marked in red). The centre of the cluster is covered by 5 out of the six receivers but most of the map is covered by three or less receivers, once again suggesting careful planning is necessary to optimise the performance of a tracking system in this environment.

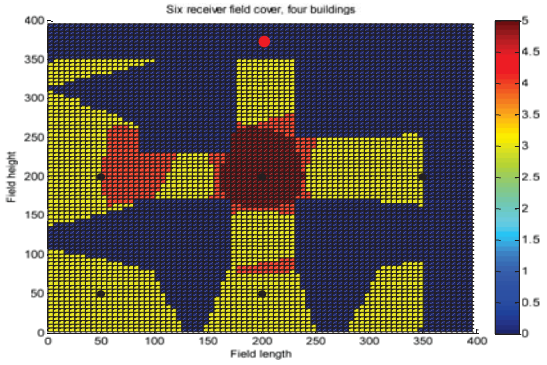

Figure 7.22: 6-receiver coverage of the close obstruction terrain map.

The plot of GDOP for this configuration of receivers is interesting in that it shows that even though there is coverage of the area, the precision of position measurements will be low in

some areas. Figure 7.23 gives the GDOP (colour axis) for this configuration.

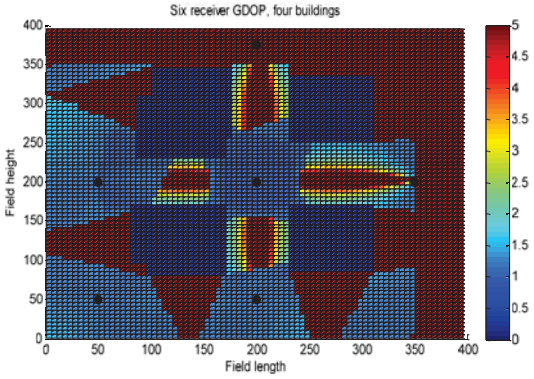

Figure 7.23: Six-receiver GDOP of the close obstruction terrain.

Note how in the spaces between the buildings where there is coverage by only three receivers the GDOP increases, this happens because the three receivers covering those areas are in line with one another which, from the literature review, is known to be the least precise configuration. Where there is greater coverage the GDOP is reduced and approaches 1 (the theoretical lower limit of GDOP). The reduction occurs because the extra receivers covering those areas are not on the same line as the other receivers and thus compensate for the imprecision. Also note that the receivers on the outside of the cluster are in areas of low GDOP even though there is coverage by only three receivers, again this happens because the receivers are not on the same line.

Finally, the static-obstruction model's ability to concentrate on a given area is tested by using the terrain map shown earlier in Figure 7.12. Recall that here only the areas coloured blue are analysed by the model. The results from an analysis with six, arbitrarily-placed receivers are shown in Figure 7.24. Note that the track here appears upside down compared with the map. This is owing to differences in the location of the origin between images and plots. It has no bearing on the results (note also that this was occurring in the previous trials but was more difficult to see due to the symmetrical nature of the images).

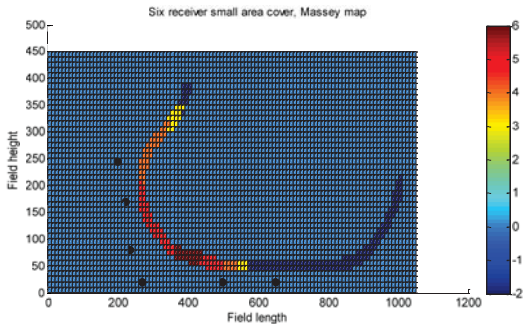

Figure 7.24: 6-receiver small area coverage on the Massey map.

Here the analysis has occurred only on the sections of the terrain map that were blue as was intended. This has resulted in most of the map being classed as zero coverage (this is the default value given that we are not interested in these areas). Any area of the track with less than three-receiver coverage has been set to -2 so that the remaining track can be visualised. The receivers here were placed in a 'lamp-post' fashion near to the known track of the objects. It can be seen that coverage is greatest near the centre receivers. By continuing this placement pattern it should be possible to provide coverage for the whole track. Extra receivers at each end would ensure that the entire track was covered. Once again, an optimisation process can be conducted to find the optimal spacing and placement of the receivers.

The GDOP plot shows what would be expected: Low GDOP where there is maximum coverage and because the track is outside the line of the receivers the GDOP is less than or equal to 2 wherever there is coverage.

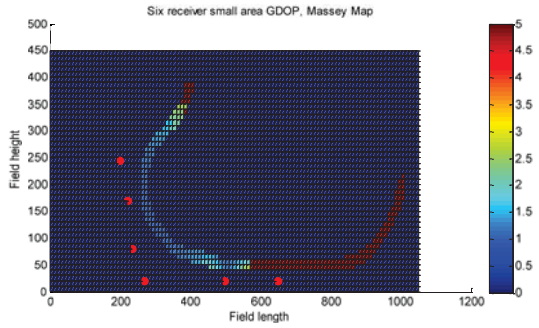

Figure 7.25: 6-receiver small area GDOP on the Massey map.

Analysing a map where the track of the objects is known greatly speeds up the computation. It will also have benefits in optimisation since it will concentrate the process on areas that must have coverage. The optimisation process is the topic of the next chapter.

7.4 Further development

The assumption that all players maintain an upright stance in the input data has been shown to be harmful to the validity of the results. A model that can predict the stance of the players based on their locations relative to one another would be useful here. As an example, if this model were to recognise a scrum from the locations of the players usually involved in such a formation then the stance data for those players could be adjusted to reflect their usual stance in a scrum. Similarly, a player who is moving quickly with no other players nearby will have an upright stance. Alternatively, player-position data could be obtained that records the stance of the player.

7.5 Visibility model conclusions

The idea and need for a visibility model has been presented in the context of a game of rugby tracked using an RFID, TOF-based system. A suitable, if conservative, model has been developed.

The data available for testing the model were discussed and the limitations and assumptions stated. It can be concluded that better data, which include information about the orientation and stance of the players, would make the model more accurate and reduce the amount of pre-processing required, but would not significantly change the results. Development of such a data set was not pursued, because it was outside the scope of the research and the required time to implement it would have been prohibitive.

Analysis of the model output suggests that the limited transmission arc of tags is the greatest cause of tag invisibility. Increasing the transmission arc of the tags will therefore have the greatest effect on performance. Increasing the transmission arc can be accomplished with the appropriate combination of tag antennae, placement on the player and padding to separate the tag from the shielding effects of the body (see chapter 6 for experiments on this shielding effect).

This visibility model was shown to be effective at calculating visibility in any open field situation.

The visibility model was further extended to handle static obstructions and known object-tracks. This generalises the model and allows it to be used in wide area sports (such as orienteering) or warehouse situations. The concept of a terrain map was introduced that allows the user to input extra data to the model (such as the height and composition of obstructions). This model was validated using a number of trial terrain maps and it was found that even small obstructions hindered both coverage and GDOP. This necessitates some kind of optimisation process in order to effectively use a minimalist tracking system.

With a valid visibility model it is now possible to optimise the placement of receivers in a tracking system to maximise tag visibility. This is the focus of the next chapter. It is also possi-

ble to develop and then evaluate the effectiveness of methods for estimating the missing position data based on the valid data.

8 Receiver Placement Optimisation

Methods for optimising the placement of receivers – rugby player tracking as an application – real game data for training – open field optimisations – measures of fitness – results of optimisations.

In order to meet the specifications for cost, performance and volume (see chapter 3), the tracking system must operate at peak effectiveness. This requires that it produce as many position outputs as possible with the highest accuracy possible which will, in turn, reduce the need for redundant equipment. This goal is valid for all tracking systems in any situation since minimising cost and volume while maximising performance is desirable in all tracking applications.

In previous chapters, the idea was presented that the placement of receivers around the area of interest contributes significantly to the correct operation of the system. Finding the optimum configuration of the receivers about the area of interest is, therefore, critical to the successful implementation and operation of the tracking system.

In this chapter, genetic algorithms (GA) used as a tool to find the optimum solution to the problem of maximising performance of a tracking system. Criteria are defined that gauge the effectiveness of a particular configuration of receivers; these are known as objective functions. Objective functions are critical to a successful optimisation and hence most of this chapter will describe their development. The GA optimisations are compared against more contemporary optimisation techniques.

8.1 Genetic algorithm optimisation

A genetic algorithm is a tool for optimising complex problems that have large search spaces and several local minima or maxima in the solution set. A GA does not search the solution space as such, instead it uses techniques inspired by biological evolution[151] to 'breed' the optimal solution.

The process and terminology for a GA is as follows. Each solution (or individual) is represented as a string of parameters, which are sometimes referred to as chromosomes. An initial population of individuals is generated (usually randomly) within predefined bounds. Each individual is evaluated using an objective function and is assigned a quantitative value relating to that individual's fitness as a solution to the problem. Once this is done, a set of breeding individuals is selected on a random basis, where individuals with higher fitness values have a higher probability of selection. The breeding individuals are then used to breed the next generation (called the children). The fittest solutions are passed unchanged to the next generation; these are called elite children.

Breeding is performed through a combination of crossover and mutation. In crossover, two 'children' individuals are generated from two 'parent' individuals by selecting at least one break point in the parents and juxtaposing the leading chromosomes from one parent to the trailing chromosomes from the other and vice

versa. Crossover always yields two children individuals that inherit most of their traits from the parents. Figure 8.1 illustrates the crossover breeding process.

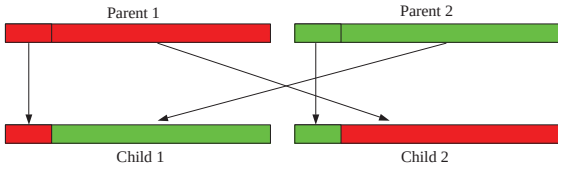

Figure 8.1: Illustration of crossover breeding.

With mutation, individual bits are randomly toggled; this requires only one parent and will yield only one child. The random numbers need not be uniformly random but could instead be part of a Gaussian distribution (for example). The use of non-uniform random distributions yields a 'controlled mutation' and prevents wild changes in the population. Mutation is a method of introducing new genetic data into the population. At appropriate levels, mutation prevents stagnation of the population and increases the probability that the optimum solution will be found (as opposed to a local minimum). An illustration of the mutation process is given in Figure 8.2 below. These breeding processes are further explained in the example of job-shop scheduling by Murata, *et al.*[152].

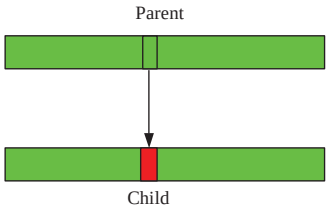

Figure 8.2: Illustration of mutation.

This evaluate-select-breed process is repeated until some termination condition is reached. In cases where the specific value of optimal fitness is known, this condition is a threshold of fitness. For situations where an optimal fitness is unknown, the termination condition is usually a predefined number of generations or a threshold of variance between individuals. A combination of maximum generations and variance threshold has been used in the optimisations discussed in this chapter; whichever condition is satisfied first will terminate the optimisation.

All GA optimisations described here have been performed using the Genetic Algorithm Optimisation Toolbox (GAOT) for Matlab 2008.

Of all of the methods for optimisation of a problem such as this, it could be asked, 'why use genetic algorithms?' To anyone new to the field, an obvious course for finding the optimum solution would be to try all possible combinations and see which yields the best results (known as the brute-force method).

The optimisation of any functional tracking system will be a hyper-dimensional problem and thus, a brute force search will be prohibitive. If a simple case of six receivers is considered, then to get an optimal solution, analysis must be made of every possible placement of those six receivers. In programming terms this requires six nested loops for the receivers alone. Additionally, the analysis must examine the conditions at a suitably large number of points on the field, which leads to several more degrees of nesting. Adding in capability to account for both static and dynamic obstructions will further nesting. The runtime on a brute force approach is unacceptably long; especially as the analysis becomes more complex, since increased complexity will exponentially increase the runtime.

In the case of a game of rugby, the journey to the fittest solution is a very shallow path. There is no single solution that will yield fitness far in excess of all others, and the probability of there being local minima in the solution set is high. This precludes the use of methods such as steepest-descent optimisation, which is prone to getting stuck at a local minimum. In the case of tracking in an environment of static obstructions there will be many more local minima that steepest-decent methods may get stuck in.

GAs are specifically designed to deal with the barriers to optimisation that arise from large numbers of input variable combinations and solution spaces that have some mix of shallow decent, local minima and discontinuities. Furthermore, GAs can reach the optimum solution via a shorter path than gradient methods [153].

Rastrigin's Function is used to test optimisation methods. The equation for this function is shown in equation (8.1)

$$Ras(x) = 20 + {x_1}^2 + {x_2}^2 - 10(cos2\pi x_1 + cos2\pi x_2) \tag{8.1}$$

The results of this function are shown in Figure 8.3 and it is clear that this is a multi-modal function with a global minimum (this occurs at the origin where this function has a value of 0).

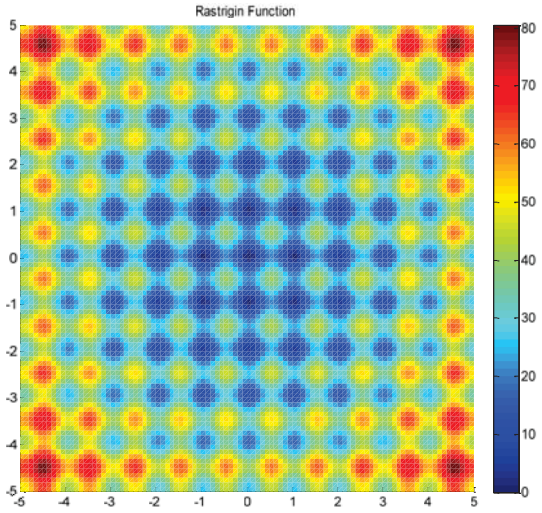

Figure 8.3: Rastrigin, optimisation test function.

A GA can always find this global minimum to within two decimal places. Note that GA optimisations have a random element and thus will not always produce the same solution for the same problem. Furthermore, this form of optimisation does not need to be supplied with a starting point because of the initial population process. These qualities of GA optimisations increase the probability of finding the optimum and give greater confidence in the results.

A Nelder-Mead-type simplex optimisation, in contrast, does require a pre-defined starting point and becomes stuck in a nearby local minimum. This type of optimisation does not have a random element and will thus always reach the same solution for a given start point. This makes it more difficult to confirm that a solution is optimal.

Given that it is expected that receiver-placement optimisation will have a multi-modal solution space and it is possible that plausible starting points may be difficult to visualise, GA optimisation has been selected for use in this problem. Further evidence will be presented that GA optimisations are the superior choice for this problem. The results will be compared and contrasted with other optimisation methods.

8.2 Optimisation configuration

The Matlab GA optimisation engine must be provided with several parameters that instruct it on how to perform the breeding process that yields new, and potentially better, solutions. Table 8.1 lists these parameters and the arguments that were used for the rugby-player, tracking problem.

Table 8.1: Summary of GA input parameters.

Parameter	Input argument
Bounds	[0 400] for each receiver
Initial population	50
Termination condition	Maximum generations reached
Crossover functions	Arithmetic, heuristic and simple
Mutation functions	Boundary, multi-nonuniform, nonuniform and uniform

The bounds define the allowable position of each of the receivers. In the Matlab GA, the bounds are defined in an array that contains one minimum and one maximum bound per row, where the number of rows defines the number of variables. The placement of receivers in all of these optimisations has been linearised to halve the number of variables that must be optimised. Instead of passing a pair of Cartesian co-ordinates to the objective function these optimisations pass a single variable that represents a distance along the allowable path of the receivers. This process can be likened to vectorising a matrix where successive rows are concatenated onto the end of the previous row and the entire matrix can then be indexed in one dimension. For the rugby-player tracking problem, the allowable path for the receivers is around the outside of the field, for the static-obstruction problem the allowable path is defined by the vectorised terrain map. Calculation is performed within the objective function to convert these values to useable Cartesian co-ordinates (which are easier to perform calculations on). The number of variables is user defined and is held constant for each execution of the optimisation.

The initial population parameter defines the number of solutions that are to be randomly generated for input into the optimisation. The size of the initial population controls how many executions of the objective function will be required per generation of the optimisation. Higher initial populations will result in longer runtimes but will be more likely to reach the optimal solution. Smaller initial populations will allow faster execution but run the risk of stagnating. Population stagnation occurs when all of the solutions become similar through the breeding process. In such situations, further breeding will not produce different solutions and therefore no further optimisation can occur regardless of how close the best solution is to the global optimum. The value for initial population can only be found experimentally and the risk of population stagnation must be weighed against excessive execution time.

The termination condition is set such that the optimisation will output the best solution after a given number of generations. A fixed number of generations is advantageous here, as the best possible solution is unknown and hence, if the GA were to terminate at a threshold of fitness there would be a possibility that the optimisation would either end before the ideal solution was reached (threshold set too low) or never terminate (threshold set too high). Again, this value can only be found through experimentation. The number of generations required for this optimisation can be as low as 100 for simple objectives up to 2000 for more complex functions.

A description of how each of the breeding functions works is given in the GAOT help file available in the Matlab GAOT documentation [154].

8.3 Training data

The objective functions described later in this chapter can be separated into two distinct categories: those that analyze the entire area of interest, and those that rely on a set of training

data for analysis. The first category makes several implicit assumptions, which are described in the following subsections.

8.3.1 Clear field

In analyzing the whole area, it is assumed that the area is empty except for the object to be tracked. Thus any receiver in range of that point on the field would see the signal. In a real situation, such as a game of rugby, this assumption is invalid. Many of the player formations that occur in a game of rugby involve tight clusters of players, and it is a certainty that at some time, one or more players will be in the line-of-sight path of a signal from a particular player. A game of rugby can be considered as an exercise in tracking objects in an environment of complex moving obstructions.

By definition, optimisation with static obstructions cannot be considered to be on a clear field.

It has been shown that obstructions such as people or buildings will absorb the signals from the tags, thus confirming that the clear field assumption is inappropriate.

8.3.2 Evenly distributed occupancy

When analysing the whole field there is also the assumption that there will be an evenly distributed occupation of the field by the tracked objects. In light of the occupancy signatures of the data provided by the Industrial Partner (see section 14.1) and knowledge that in static obstruction applications (warehouses for example) there are designated thoroughfares, this assumption can be seen to be invalid.

The effect of this assumption is that all parts of the area of interest are given equal consideration. Better results may be obtained by giving greater significance to areas where the objects appear most often (e.g. occupancy signatures or defined areas on a terrain map). Conversely, it maybe advantageous to have position measurements returned reliably at all locations within the area of interest, regardless of the probability of objects actually occupying those locations especially in sporting applications.

8.3.3 Real Game Data

Using data from an actual situation as training data is an alternative to analysing the whole field. With real data, both the clear field and evenly distributed occupancy assumptions are circumvented. The use of real data has the advantage that the number of points that must be analysed is reduced and the points are ones that objects have occupied. The disadvantage is that unlikely but possible configurations of objects may be missed. Situation data are required for those objective functions that use the player visibility model described in chapter 7.

For the rugby-player tracking problem, a suite of game data was obtained from the Industrial Partner. This suite contains 10 games, tracked manually from video. Each game contains approximately 3000 frames of positions for 30 players. A number of hypothetical situations have been defined for the static obstructions problem that involve mapping of actual building formations.

8.4 Fitness Measures

In order to perform a GA optimisation, some quantitative measure of fitness must be defined by which the individual solutions can be judged. Several fitness measures have been defined for the problem of tracking objects both on the field and among obstructions. These measures are independent but can be combined in the objective function as required.

8.4.1 Field coverage

This measure assumes that the fittest solution to the tracking problem will have a high proportion of the field in line of sight of at least three receivers. With three receivers a two-dimensional position can be calculated and more receivers will increase the precision of that measurement.

Furthermore, in order to achieve the minimalist system that performs as desired (see Specifications chapter), each receiver must be used to its full potential.

8.4.2 Position Measurement Accuracy

Accurate calculation of positions is one of the requirements of the player-tracking system (see section 4.6) thus estimating the positional precision that can be expected at a particular location within the field of interest is a useful measure of fitness of an individual solution.

Here, accuracy is defined in terms of positional error, where high accuracy corresponds with low positional error. Thus a position calculation that is 0.2m from the true position of the object is more accurate than a calculation that is 0.6m from the true position.

Positional error is found using the true location of the object and the estimate of position that the tracking system would output (based on the locations of the receivers and including any errors in timing); the difference between the two is the error.

8.4.3 Invisibility rate

Data are available from the visibility models that indicate when visibility to an object is lost. Minimisation of this number will signify a fitter solution.

This measure links in with field coverage since an area that is not covered will have invisibilities if objects move into them. Thus maximising coverage of areas where objects are is the first step in minimising the invisibility rate.

This measure directly contradicts the clear-field assumptions and, therefore, can only be employed in objective functions where the positions of objects and invisibilities are known. This knowledge can come from an actual game, as has been used here, or random generation.

8.4.4 Receiver separation

It has already been shown that the geometry of receiver placement is an important factor in the performance of the tracking system. The ideal geometry is for adjacent receivers to be evenly distributed around the object in a regular polygon (assuming all receivers have visibility to the object). A fitness measure based on adjacent receiver separation will help in the search for an optimum solution.

Furthermore, GDOP can be calculated directly from knowledge of which receivers are visible to a given point on the field. Constraints on receiver range and placement as well as static obstructions may make the solution with the lowest GDOP less than obvious.

The loss of visibility to a critical receiver will have serious repercussions in the minimalist system we aim for. Consider an example in which three receivers are in line and one is off that line. The off-line receiver is critical for lowering GDOP of the group. If that receiver looses visibility, the GDOP will rise considerably. Thus GDOP is, in part, dependent on visibility.

8.5 Objective functions

The objective function is the heart of any GA as it defines which individuals are suitable solutions to the problem and hence which solutions will be used to breed the next generation. During the course of the investigation it was planned that initial objective functions be simple to ensure that the workings of the GAOT were properly understood. Once these simple functions were working they could be expanded to increase their effectiveness. A description of each stage of planned development of the objective functions is given in this section. Objective functions are split into two groups; those relating to the rugby-player problem in the first and those relating to the static-obstruction problem in the second.

Note that effectiveness is defined here to mean the objective function's ability to find an opti-

mum solution to the problem. Effectiveness answers the question, '*how well does this function optimize the solution*?' Here the measure of effectiveness changes with the different objective functions as they each try to optimise for a key characteristic of the problem. There is a requirement for a standard quantitative measure to compare the results of each of the objective functions so that the most effective can be found. For the rugby-player problem this measure will be the number of invisibilities (calculated from the player visibility model in chapter 7) that occur in a fixed set of data from the pool of input data provided by the Industrial Partner. Higher values will denote more obstructions and hence a less effective solution. Average GDOP will be used for the static-obstruction problem.

The source code for each objective function is included on the attached DVD, see Appendix C: DVD index for the file names, descriptions and locations.

8.5.1 Solution by inspection

Though not an objective function, it is expected that the optimisations will yield an effectiveness value better than could be achieved by a configuration found through inspection.

From knowledge of the sources of error in positioning systems and inspection of the field of play, it is hypothesised that the best result will be obtained from a group of receivers that are spaced evenly along and several metres back from the sidelines of the field of play. This should minimise the positional error due to the geometry of the receivers and also be the most resilient to obstruction at any part of the field. The Figure 8.4 shows the solutions from inspection for 6, 8 and 10 receivers.

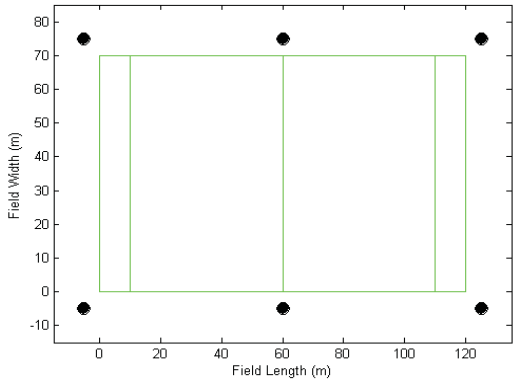

(a)

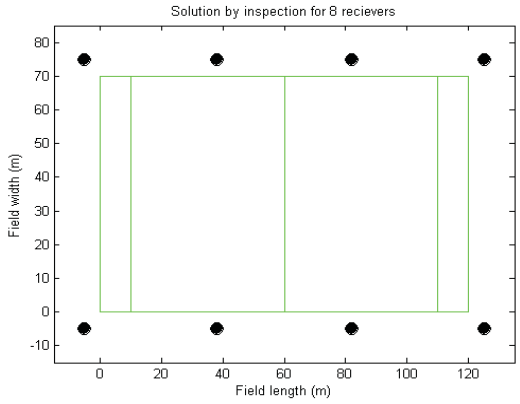

(b)

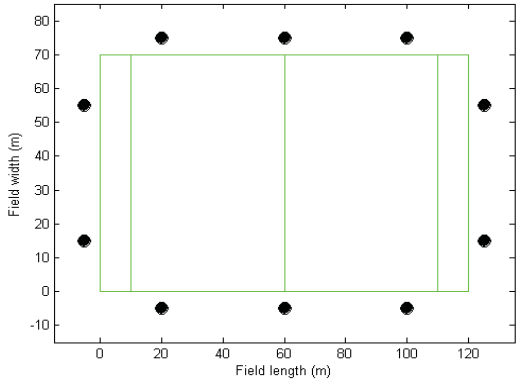

(c)

Figure 8.4: Solutions by inspection for (a) 6, (b) 8, and (c) 16 receivers.

These configurations will be used as the benchmark for solutions from other objective functions that aim to solve the rugby problem. Other solutions must perform better than these

in order that the objective functions are deemed effective.

The static-obstruction objective functions have SBIs dependent on the terrain map and as such will be discussed when those maps are introduced.

8.5.2 Maximise field coverage

In this objective function the entire area is analysed for each individual. The objective function populates an array, called the field-coverage array, which can be mapped to the area of interest. The dimensions of the field coverage array represent the resolution to which the analysis is performed. For example, for a 110x70m field of interest, if the field coverage array were of dimensions 110x70 then the resolution would be 1m in either direction. Higher resolution will give a more accurate fitness value but will take longer to calculate.

When this is applied to the rugby problem a distance is calculated between the receivers and each point on the field for every element in the array. For every distance that is less than the range of the receivers the value in that element of the field-coverage array is incremented. Thus if an element in the field coverage array contains the value 3, then that location is in range of three receivers (and will hence output a valid position measurement).

The direct-calculation approach taken in this objective function uses the clear field and evenly distributed occupancy assumptions so the orientations of the tags have no effect on the line of sight and there are no obstructions to transmission.

The overall sum of the field-coverage array gives the fitness of the current individual. Maximising the sum of the field coverage array will move the solution towards the optimum.

It is also assumed that the receivers have an isotropic receive arc. It is more likely that the receive arc will be lobe shaped but the assumption has been made to decrease complexity. See section 6.3 for the receiver lobe trials. The range for the receivers is set at 70m, as this is a conservative estimate of range found from the system characterisation.

8.5.3 Look-up table

The assumption of isotropic radiation pattern of the receivers in the previous objective function is false. It is far more likely that the receivers will have lobe shaped reception arcs. It is much simpler to define a look-up table with the lobe shape of the antenna than it is to calculate those ranges at each execution.

In the look-up table method, a map of the range of the receiver at various angles is created. This look-up table is navigated by calculating the angle between the transmitter and receiver then going to the nearest entry in the table.

The look-up table function can accommodate any antenna so long as the reception arc can be defined. This gives the look-up table objective function greater flexibility for handling different receiver types compared with the direct-calculation method used in the previous objective function.

The look-up table objective function moves to maximise the field coverage as before.

8.5.4 Rewards, penalties and non-linearity

Clearly, having coverage from only two receivers will not yield a position measurement. Yet areas of the field with two-receiver coverage still increase the fitness value of the individual. Conversely, field coverage by more than three receivers is advantageous. Rewards and penalties are introduced to rectify this. This functionality is added to the look-up table objective function described above.

A point on the field that is in range of less than two receivers will be penalized by up to 30[10] from the value at that point on the field. This penalises any individuals that have areas of the field covered by less than three receivers. By the same token, areas of the field that are in range of more than three receivers are rewarded with an increase in that element of the field coverage array. The Matlab code for rewards and penalties is shown in Figure 8.5.

```
if field(y,x) <= 2
    field(y,x) = (field(y,x)-2)* 30;
end
if field(y,x) >= 3
    field(y,x)=(field(y,x)-3)*0.1 + 3;
end
```

Figure 8.5: Matlab code fragment for introducing non-linearity.

The value of 30 for a penalty was selected as one that was in excess of the maximum that could be expected from receiver coverage. This was necessary so that solutions with any areas of the field that would fail to produce a position are calculated to be inappropriate. It is necessary for a useable system to be able to output a position for an object at any location within the field of interest.

This function also addresses the problem that the simple objective functions average receiver coverage across the field in order to get a fitness value. It can be seen that, in these simple objective functions, an individual that has every point on the field in range of three receivers will have the same fitness values as an individual that has half of the field covered by six receivers and half not covered at all. The penalties introduced by the code in Figure 8.5 will address this problem by severely penalizing the individual with half the field covered by six receivers and should have the effect of spreading coverage over the whole field.

8.5.5 Minimise Average Position Error

This is the first of the objective functions classified into the real data category. These objective functions use a selection of 100 frames at random from the real-game-data pool. These frames give the position data that allow the evaluation of fitness of the receiver configuration. The use of a small sample of random frames that change with each function evaluation prevented the optimisation from trending towards a solution that is optimal for only this relatively small sample set (overtraining). Analysis of too many frames would increase the runtime to an unacceptable level. The calculations of the required sample size to give the desired level of confidence in the data are well understood and several online calculators exist to aid in this process; however, experimentation with the objective functions is required in order to find a sample size that executes in an acceptable time.

Those players who have line of sight with the receivers are found by using the player visibility model described in chapter 7. Once these players have been identified, the position measurement that the tracking system could be expected to output is calculated and compared with the known position of the player. The straight-line difference between the known and measured positions is the positional error. The process for calculating the position is clarified below.

It is possible to obtain an indication of the positional accuracy that can be expected by using a position calculation method similar to that used by the tracking system[50].

Recall that the proposed tracking system is a TOF, tag-based system. The critical measurements are the travel times of signals from tag to

[10] The value 30 is chosen as it is 10 times the minimum number of receivers required. This imposes a penalty significantly greater than the gain of maximum coverage. This ensures that a solution with insufficient coverage is low down in the performance rankings.

receiver, where that signal is moving at the speed of light to reach the receivers. In fact, in this function it is only necessary to know the distance from the player to the receivers, since distance is proportional to time. Gaussian noise is added to these values to simulate errors in timing at the receivers[11].

The process of calculating a position is, in itself, an optimisation process. A Nelder-Mead[155] optimisation algorithm is used here and the objective function minimises the sum of the squares error between the measured times of flight (parameters to the problem), the calculated times for a given position and an epoch start. Fontana[50] describes the operation of this optimisation more fully. The best case is an error of zero, which would signify that the estimated position satisfies all of the measured times of arrival. The equation for this objective function was originally presented in section 2.1.4.

$$e = \sum_{n=1}^{N} \left[\sqrt{(x - Rx_n)^2 + (y - Ry_n)^2} - c(t_n - t_0) \right]^2 \tag{2.3}$$

Players whose signals have been lost are given penalty values of error. This penalty makes solutions with a high number of obstructions less fit, thereby forcing the solution towards one that minimises obstructions as well as position error. The penalty in this case is 1 m. This value should be well in excess of anything that we can expect from calculations of position based on the system characterisation already performed.

At the end of these calculations, the resultant two-dimensional array that contains the position errors of each player for all of the frames is averaged to give a fitness value that represents the average position error. The optimisation will minimise the average positional error across the field.

Using an objective function that minimised the maximum error was considered here but was discounted after it became apparent that the optimisation would always see the penalty value as the maximum error. This would lead to a constant fitness value independent of solution. The optimisation process must see change in the fitness value in order to operate correctly.

8.5.6 Minimise the Number of Obstructions

This objective function was conceived after analysis of system characterisation that measured the precision of data from the tracking system showed that the measurements of position were precise to 0.3m wherever the signal was present and not subject to high GDOP (see section 6.6.4). If the signal cannot be seen then there is no position output by the system. The upshot of this is that the tracking system has an approximately constant error in positioning across the entire field of interest and therefore the specific accuracy of any one measurement can be ignored.

This objective function uses the player-visibility model to count the number of times a loss of signal occurs in the input data for a given receiver configuration and uses this as the fitness value. No penalties are necessary here.

The optimisation moves to minimise the number of signal losses and hence achieves maximise player visibility.

8.5.7 Variable Height

Up to this point it has been assumed that the system must be completely portable and, hence, the receivers must all be mounted near to the ground. In the interests of thoroughness this objective function checks for any advantage in permanently mounting the system with at least some of the receivers at a higher altitude.

[11] This Gaussian noise is distributed about the mean of 0ns with a standard deviation of 3ns. These values were chosen to induce an error up to 1m in either direction at a given receiver and also include the possibility that no timing error occurred.

The hypothesis for this objective function is that by having receivers at a greater height the likelihood of signal being obstructed by the bodies of other players is reduced when the players are upright. If the players are involved in a tackle and several players are piled up there is still a high probability of signal occlusion.

In this function the receivers are placed around the perimeter of the field as usual but an extra input variable defines how far back from the field the receiver should be placed. A threshold distance is defined; any receiver on the field side of the threshold is at the previously defined ground height of 2.5m and receivers on the outside of the threshold are placed at 15 m off the ground. This value of 15 m represents the receivers being mounted to the roof of the stadium.

The optimisation is given the option of putting the receivers on the ground so that it can be determined if putting them on the roof is actually the best course of action for getting the best performance.

Once again the receivers are not allowed within 5m of the field of play, even if they are mounted on the roof. This decision was made because most stadiums in New Zealand are not completely enclosed permanently. Some have retractable roofs but these roofs are unlikely to have structures that will support the receivers and their associated power and data transmission equipment. The stadiums do have covers for the spectators; it is to these covers that the receivers would be mounted.

This objective function moves to maximise player visibility once more.

8.5.8 Weighted field coverage and receiver separation

To avoid the computational intensity of analysing a large number of individual frames to calculate the fitness value, the cumulative occupancy can be used as a weighting for a field coverage objective, albeit slightly modified.

The occupancy measure compresses the entire data suite into one frame and, in doing so, removes all temporal information from the data. This makes it impossible to infer anything about the visibility of any player to the tracking system.

In this function, each point on the field is given a value representing the number of receivers that are in range, also known as the field coverage value. The field coverage at each point is then weighted by the occupancy of that field location. Thus, areas that have a high occupancy are more heavily weighted than those with low occupancy. The sum-of-weighted-field-coverage is used as a component of the overall fitness.

In addition to the weighted field coverage, the maximum angle between adjacent receivers is calculated. This objective function minimises this maximum angle since it is known that the lowest GDOP will occur when the receivers are evenly spread. Calculating the angles of the receivers circumvents the need to directly calculate GDOP in an attempt to increase execution speed.

The significance of the maximum angle is controlled by its weighting. The weighting value required to produce optimum results is difficult to predict and will have to be found from a series of trials.

The fitness is the difference of the contributions from occupancy-weighted field coverage and the weighted, maximum receiver separation. Hence the optimisation will move to maximise field coverage while minimising the maximum angle between adjacent receivers.

8.5.9 Static obstructions

The static-obstructions objective function combines maximum-field-coverage and minimum GDOP fitness measures to achieve an optimum solution.

The input data for this objective function is a terrain map. Three terrain maps will be used.

These will be referred to as the Dot, Four-Building and Riddet maps (pictured in Figure 8.6, reproduced from section 7.2.5). The first two are hypothetical situations that are designed to be difficult to plan for and the last is the map of the Riddet section of the Massey University campus (Palmerston North, New Zealand). They are presented in order of increasing complexity.

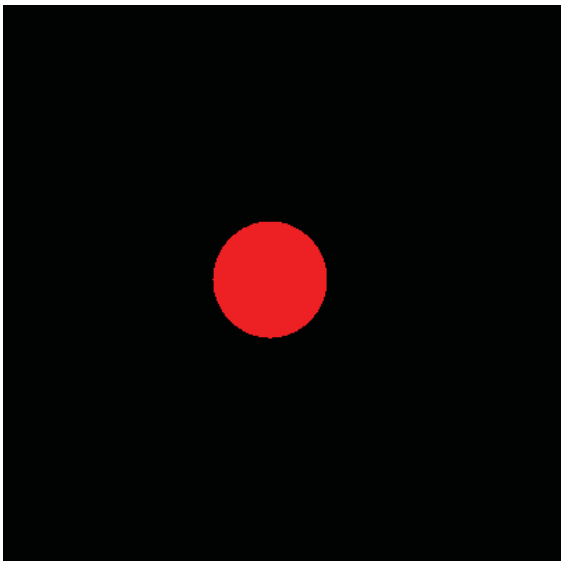

(a)

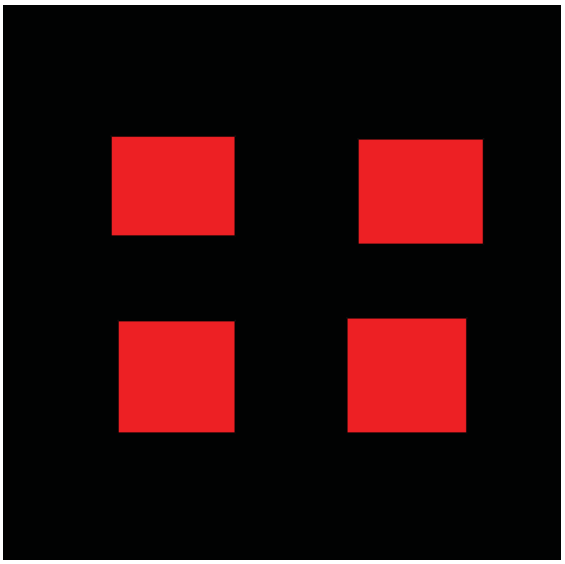

(b)

(c)

Figure 8.6: Terrain maps. (a) dot (Figure 7.18), (b) four-building (Figure 7.21) and (c) Riddet (Figure 7.11).

The Riddet terrain map will also be used in an optimisation where the track of the objects is known. Figure 8.7 shows this track in blue. This terrain map is an example of a hypothetical orienteering situation in which it is known that the competitors will take the track shown in blue and thus, coverage of that area is necessary.

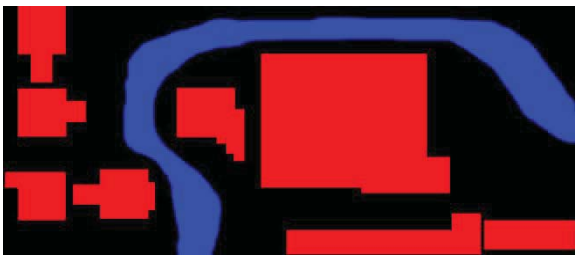

Figure 8.7: Riddet terrain map with known track included in blue.

The objective function uses the static obstruction visibility model described in section 7.2.5 to find those locations on the map (or sections of the map) where a position may be calculated then calculates the GDOP at those locations. A penalty value is applied where a position cannot be calculated and thus the objective function maximises coverage while minimising GDOP. Note that field coverage is implicitly included in the output since a location with coverage by less than three receivers is given the penalty value of GDOP.

Also, in this situation it is unacceptable for receivers to be placed within buildings and thus penalty fitness is applied when a receiver is located within a building (the receiver position corresponds with a red pixel in the terrain map).

As with the sports problem, configurations of six, eight and ten receivers will be optimised. The solutions by inspection for the terrain maps shown above are as follows in Figure 8.8 (receivers are shown as green dots). Only SBIs for six receivers will be used for comparison with the optimised results. It would be unnecessarily cumbersome to examine three SBIs for each of the four terrain maps.

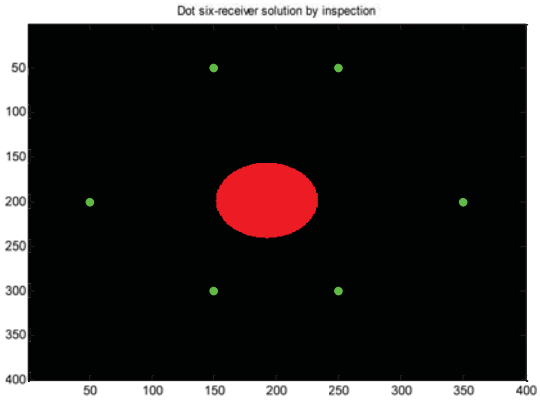

(a)

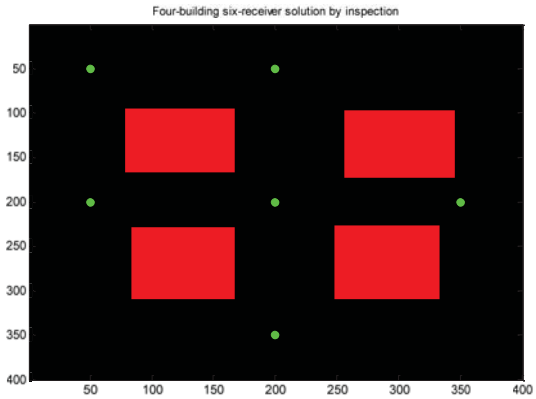

(b)

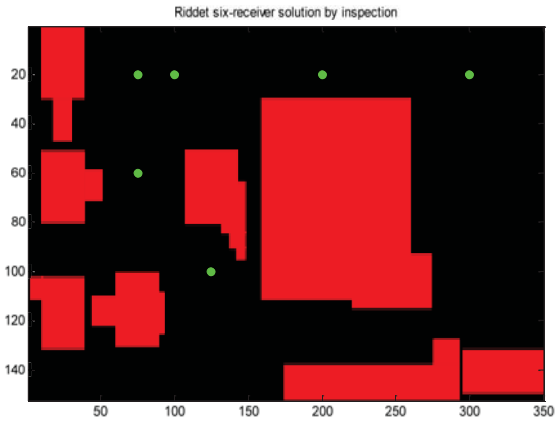

(c)

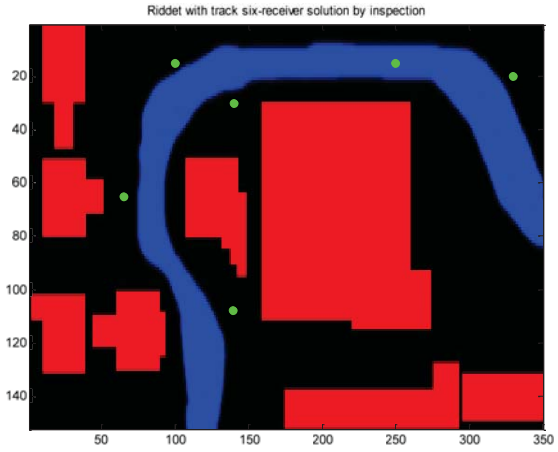

(d)

Figure 8.8: Solutions by inspection of the test terrain maps.

8.5.10 Validation and verification

Some tests are required in order to confirm that the optimisations are functioning as intended (verification) and are, in fact, calculating optimum solutions (validation).

The optimisations are verified by monitoring the progress of the optimisation after each generation. If working correctly, it would be expected that the fitness trends asymptotically to the optimal solution as generations increase; the greatest gains in fitness are made in early generations and later only 'fine tuning' is per-

formed. Significant deviations from this trend would be indicative of issues with the operation of the optimisation.

A truncated solution space, in which only two receivers are moved throughout the area of interest, will be used to map a portion of the solution space to give an indication of how fitness changes. The choice was made to move only two receivers because it would reduce computational intensity to a point where the solution space could be mapped and because two independent variables and one dependent variable are the most that can be visualised. The truncated solution space can further justify the use of GAs by identifying multi-modal solution spaces.

The random nature of the GA's initial population offers advantages over optimisation types requiring a specified start point (such as a Nelder-Mead simplex optimisation or pattern search[156]). A random initial population means that a GA can be implemented with little or no knowledge of the best location to begin optimisation. This concept will be shown with a cursory comparison of GA optimisations with simplex methods.

Validating each solution is more difficult. The only way to prove that a solution is optimal is to map the entire solution space, which is prohibitively time consuming. If the objective functions produce similar results after multiple runs and if their performance is better than solutions by inspection, when judged by the metric (visibility count for the rugby problem and average GDOP for the static obstruction problem), then the solution can be said to be optimum.

Validation and verification evidence will be given as is appropriate.

8.6 Results

After coding and testing the objective functions described, the following results were obtained. Comments and discussions will be given on each set of results and the general quantitative effectiveness measure will be applied to each solution.

Note that the less complex objective functions (those described in sections 8.5.2 through 8.5.4) have been run with a 6-receiver configuration only. Furthermore, due to the non-deterministic characteristics of the GA optimisation process, all optimisations were run several times.

For the rugby problem the number of invisibilities is used as a quantitative measure to compare the different objective functions. A fixed set of 3000 frames, selected from the data suite, is used for this measure. With 3000 frames and 30 players in each frame, the maximum number of invisibilities possible is 90,000.

8.6.1 Solution by inspection

Figure 8.4 shows the receiver placement configurations found by inspection of the problem. The number of invisibilities calculated for the 6, 8, and 10-receiver configurations were 7035, 4335 and 2810 respectively. As discussed, these values will be used as a benchmark for all of the objective functions.

8.6.2 Maximum field coverage

Figure 8.9 is a graphical representation of the best solution after 100 generations of a 6-receiver, direct-calculation genetic algorithm. The runtime of this function was approximately 1 hour.

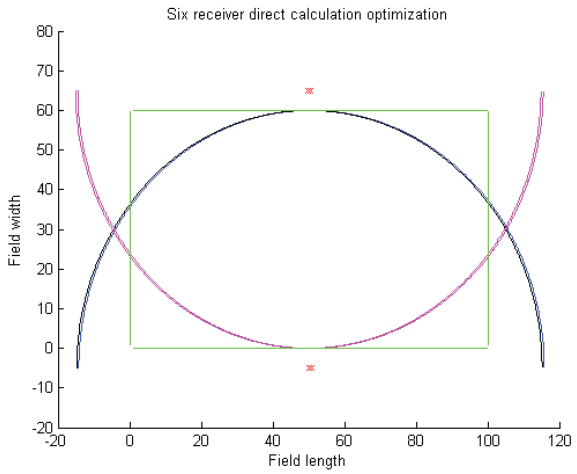

Figure 8.9: Results of a 6-receiver maximised field coverage optimisation.

The rectangular box within the graph represents the field of interest; the crosses represent the receivers placed around the field and the semi-circles represent the reception arcs for each receiver. Note that there are 6 crosses in the above figure but that they have been placed so as to almost overlap.

In this case, the best solution is one in which three receivers are placed very close together on either side of the field. This solution fulfils the requirement that the whole field be covered by at lease three receivers, and indeed, six receivers cover most of the field. Unfortunately, the clustered nature of this solution will have a high GDOP.

For this 6-receiver configuration, the number of invisible players was calculated to be 2707 for the test data set (recall that 90,000 is the maximum possible). While this does improve on the 6-receiver configuration from inspection, high GDOP make solutions found using this objective function inappropriate. This suggests that simply maximising the receiver coverage of the field is not sufficient to get a useable solution, especially when the clear-field assumption is in effect.

8.6.3 Look-up table

After running the look-up table type objective function the best solution after 100 generations was the same as that of the direct calculation (shown in Figure 8.9).

This was to be expected since using a look-up table does not change the fitness of any individual. Because the solution is the same as that of the direct-calculation objective function it has the same low level of acceptability as a solution to the player-tracking problem.

8.6.4 Rewards, penalties and non-linearity

Figure 8.10 below shows the best configuration from a six-receiver look-up table genetic algorithm optimisation. The field of interest, the receiver positions and the receive arcs are represented in the same way as in Figure 8.9.

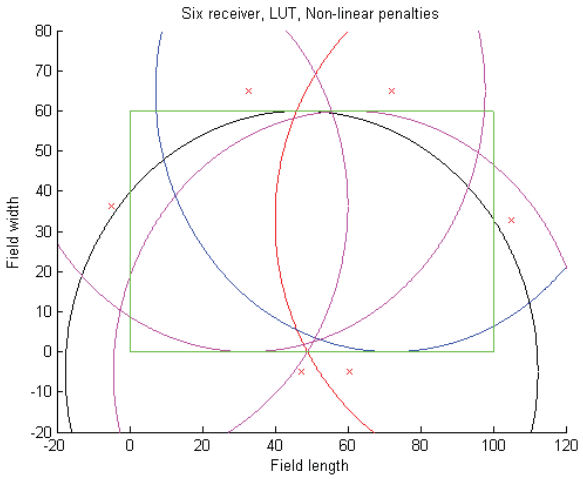

Figure 8.10: Position of receivers from a 6-receiver optimisation with rewards, penalties.

It can be seen here that using a system to penalize individuals that have areas of the field covered by less than two receivers forces the receivers apart, thus making the objective function more effective at optimising a TOF system. Most of the field is covered by at least three receivers except for the top two corners. From inspection, moving the top two receivers closer together, the minimum coverage of any point

on the field will be three and most of the field will be covered by more than three receivers. This shows that while this objective function is spreading out the receivers more and finding a solution that will give acceptable precision from a TOF system, it has not produced the optimum solution. It is interesting to note that this solution resembles the solution from inspection in Figure 8.4(a). The solution is rotated so that one receiver is located behind each goal post and the receivers on the sidelines are placed closer to the sideline than in the solution from inspection.

For the 6-receiver configuration shown in Figure 8.10 the signal obstruction rate was 2487. This improves on both the direct calculation and inspection solutions. Furthermore, given that the receivers have a significant spread around the field, the accuracy of position measurements with this configuration would be acceptable.

The truncated solution space for the maximised field coverage objective function is shown in Figure 8.11. Here four of the receivers from the six-receiver SBI are held constant (the first four counting clockwise from the top left of the field) and two are moved around the field. The plot shows the fitness at each receiver displacement, colour represents fitness (recall that receiver position is passed to the objective function as a displacement around the perimeter of the field where 0 is the origin and 400 is the same location after one revolution of the field).

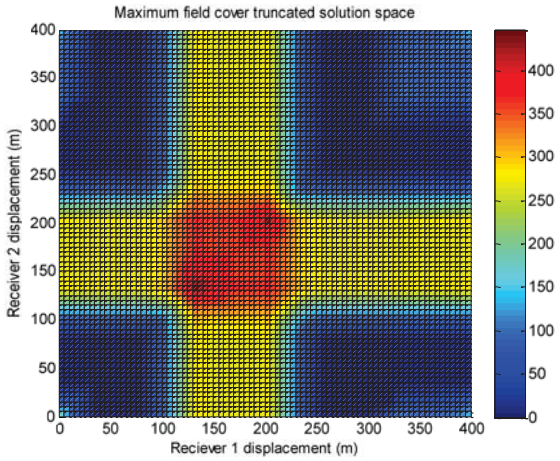

Figure 8.11: Maximum field cover truncated solution space.

The truncated solution space suggests that the problem is multi-modal but that these modes are located close to each other. The plot is symmetrical about the diagonal from the origin. This is because each receiver has the same effect on fitness and thus their positions can be reversed and the same fitness obtained.

There is an increase in field coverage when the position of the moveable receivers approaches the static receivers. This justifies the need for rewards and penalties since such solution will have a high GDOP. It also suggests that the optimisation is finding the global maximum for this objective function.

Remember that this is just one 'slice' of a seven-dimension space (six receivers and the fitness) and as such may not contain the global maximum for the problem. Instead this gives an indication of how the solution space is shaped.

8.6.5 Small sample interference

Through a series of trials it was found that the sample size of training data for the minimum average position error, maximum visibility, and variable height objective functions could not be more than 3% of the population data suite without resulting in an unreasonable execution time. Optimisations using a 3% sample took approximately 3 days to run.

The effect of using such a small sample is that the result of an objective function execution is based not only on the receiver configuration input (as is required) but is also heavily dependent on the random sample selected for that particular execution (recall that a new sample is selected at random at the beginning of each objective function execution). This means that the optimum solution may be missed if the optimisation is provided with particularly bad data. The converse is also true; a solution that will perform badly over the whole data set might have a high fitness value if it is provided with particularly good data.

This small-sample effect makes successfully running the optimisation very difficult. The fitness values do not relate directly to the receiver configuration being tested and so ranking of the solutions based on the fitness value is invalid.

Because of this random variation, the optimisations run with the minimum average position error, minimum number of obstructions and variable height objective functions did not produce any useable results. In almost all cases, the solutions returned as optimal were members of the initial population (Figure 8.12).

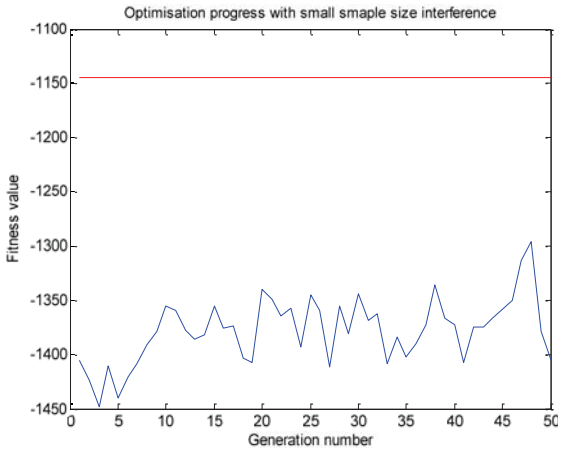

Figure 8.12: Best fitness (red) and mean population fitness (blue) in an optimisation suffering small sample interference.

Figure 8.12 shows that while the mean fitness of the solution population was changing, no solution ever exceeded the fitness achieved by one of the randomly generated initial population and thus, no optimisation was occurring (recall that in the first generation, all solutions are random).

An examination of the truncated solution space confirms this. Figure 8.13 illustrates the truncated solution space using the minimised-number-of-obstructions objective function and the same receiver configuration as in Figure 8.11. It can be seen that the problem is multimodal and that the minima are spread over approximately 60% of the solution space. Furthermore, the locations of these minima are different for each run of the calculation.

This large and flat solution space can be an advantage. If there is only a small variation between solutions then it can be said that the problem is insensitive to receiver configuration. That is, many receiver configurations will give similar results.

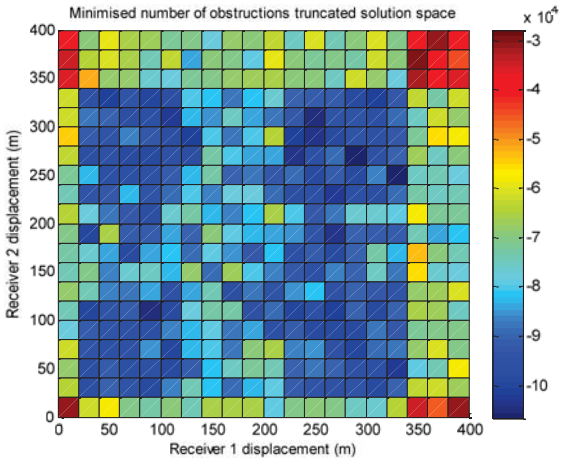

Figure 8.13: Minimised number of obstructions truncated solution space, suffering from small sample interference.

In general, the minimised average error, minimised number of obstructions and variable height objective functions cannot be satisfactorily solved until the execution time for running the optimisation with a sufficient sample size is more reasonable.

8.6.6 Combined field coverage and receiver separation

Since this objective function uses the full occupancy data in a single, two-dimensional array rather than a large, three-dimensional array, the optimisation is able to execute in approximately 30 minutes.

As expected, the selection of the angle weighting was very important to the results of this objective function. If set too low, the field coverage component of fitness has control of the optimisation and significant clustering can be observed about the centre line (Figure 8.14).

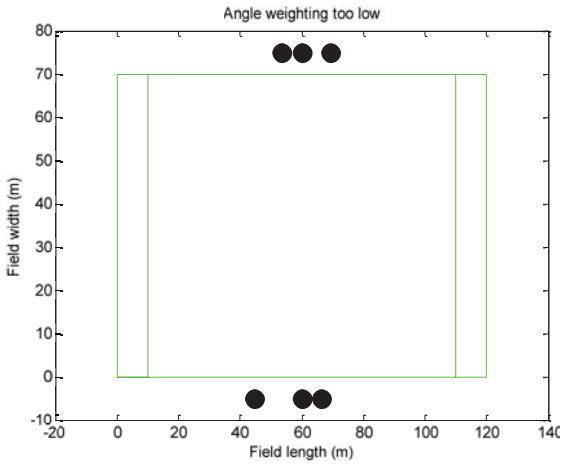

Figure 8.14: Optimisation in which the angle weighting is too low.

Placing the receivers near the centre line will result in the maximum field area being covered. Additionally, the highest occupancy occurs near the centre of the field and thus the greatest component of fitness will be achieved by covering this area with every receiver. Unfortunately, due to issues of geometry (GDOP), clustered receivers will not produce precise position measurements in the real system so the receiver separation must be given greater significance in the final fitness value. This is the same issue that was found in the maximised field coverage objective function, and it occurs for the same reasons (see Figure 8.9).

Increasing the angle weighting in the objective function caused a step change in the calculated optimum solution. With the weighting below $1x10^5$ the clustered configuration occurred, but at weightings higher than this, the optimum solution became much more suitable. Figure 8.15 shows the optimum solutions for 6, 8 and 10 receivers when the angle weight is higher than this threshold.

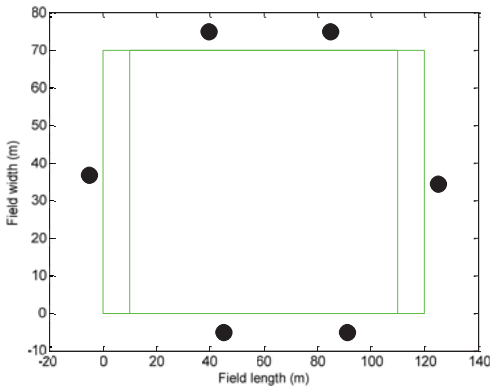

(a)

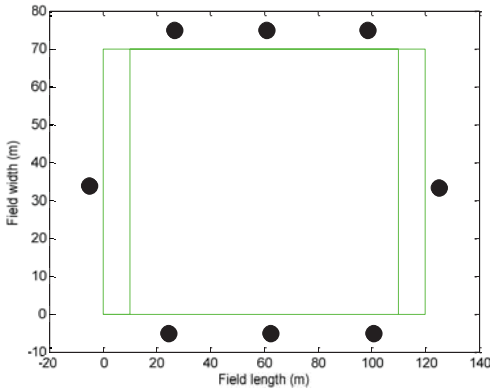

(b)

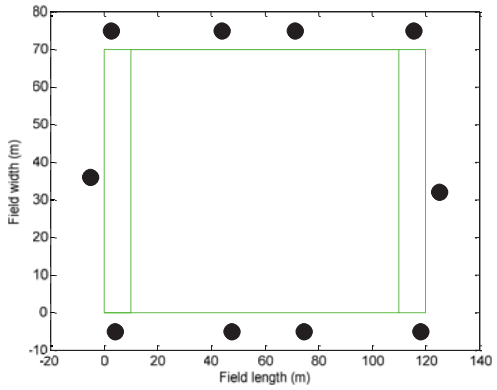

(c)

Figure 8.15: Optimum solution from (a) 6-receiver, (b) 8-receiver, (c) 10-receiver combination field coverage and receiver separation optimisation.

Note that the receiver configurations are not symmetrical about the centreline of the field. There is a definite clustering of receivers along the field length and an offset between receivers along the field width. This is an attempt by the optimisation to cover the most occupied sections of the field to reach the optimal solution. These irregular configurations could be difficult to reproduce in practice and so it is worth testing a configuration in which the positions of the receivers to the right of the centreline are duplicated from reflection of the receivers left of the centreline (referred to as a symmetrical configuration) to see if there is an insignificant drop in fitness, a drop that would be acceptable in exchange for a simpler setup. Averaged over all numbers of receivers, there was a 5% drop in fitness when using this symmetrical configuration. It is the decision of the operators as to whether this drop in performance is acceptable.

Examining the progress of the optimisation (Figure 8.16) it can be seen that an appropriately asymptotic trend is present, indicating that the optimisation is proceeding correctly. Figure 8.16 is typical of all optimisations that use this objective function.

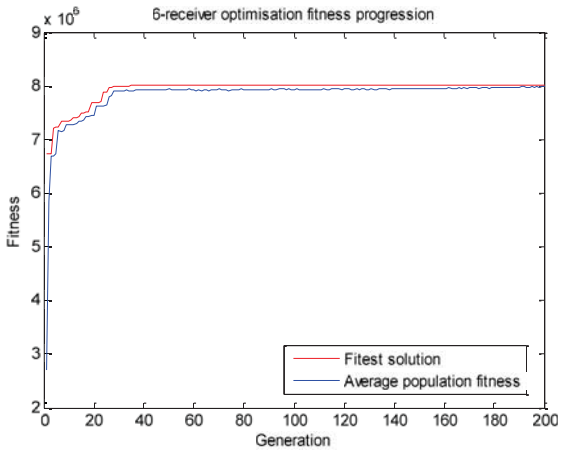

Figure 8.16: Typical optimisation progression.

The optimisation for each number of receivers was run 15 times as part of the validation process. The truncated solution space (using the same receiver configuration as previously) is shown in Figure 8.17. Here it can be seen that when each receiver approaches 300m displacement independently that fitness increases but that their both being at the same location is not the optimum solution. The optimum solution occurs when one is at 275m displacement and the other at approximately 350m displacement. Similar to the maximum coverage truncated solution space, this plot has symmetry since each receiver behaves in the same manner.

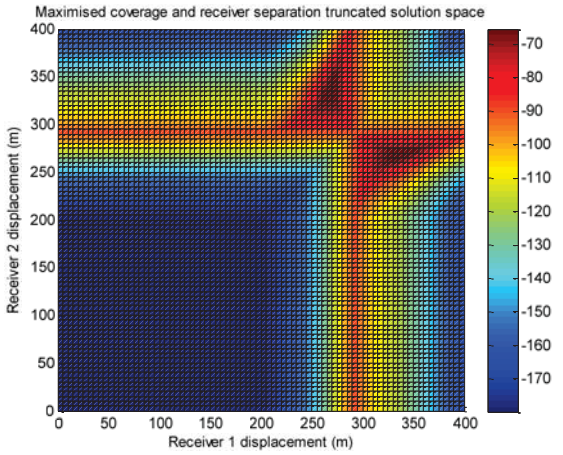

Figure 8.17: Maximised coverage and receiver separation truncated solution space.

For comparison with the solutions by inspection, the invisibility metric was applied to the optimum solutions. For 6, 8, and 10 receivers the numbers of invisibilities from the 90,000 test player positions were 6135, 4020 and 2820 respectively[12]. Figure 8.18 compares these results with the solutions by inspection.

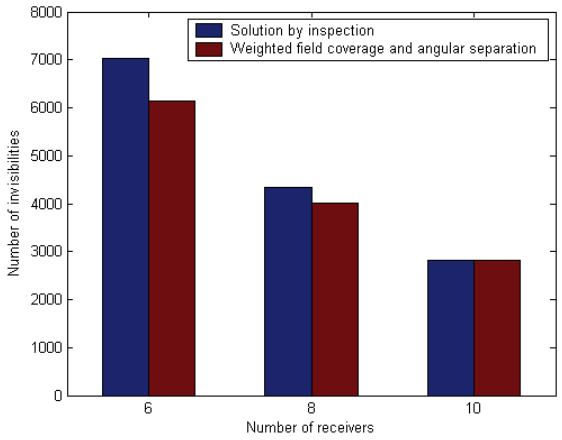

Figure 8.18: Comparison of results.

These results show that this objective function produces better results than those obtained by inspection in the 6 and 8-receiver trials. The 6-receiver configuration had a 1% reduction in invisibilities and the 8-receiver configuration had a 0.4% decrease[13]. With 10-receivers and the constraints of the objective function, the optimisation was not able to find a configuration that was better that the solution by inspection.

The configurations shown in Figure 8.15 are an improvement over the solution by inspection. However, the optimised configurations indicate a logical placement in general with some specialised tweaking to get optimum performance. The solution by inspection was the configuration that we felt would produce the best results based on inspection of the situation and knowledge of receiver coverage. It is entirely conceivable that a different operator would have suggested as the solution by inspection configurations similar to Figure 8.15 albeit more regular (i.e. a configuration symmetrical about the centre line). If that were the case then the optimisation would produce only a 5% increase in fitness. Scientifically, this shows that the optimisation can be used not only to find the optimal receiver configuration, but also to confirm that a given solution is optimal.

It should be noted that the field coverage component of fitness is important even though it appears that the angle component is in control of the optimisation. Without the field coverage component the receivers could be placed anywhere around the perimeter so long as they are equally spaced, resulting in a family of solutions where receivers are moved an arbitrary distance around the field perimeter but the angle-component of fitness is the same (Figure 8.19). These solutions have varying numbers of invisibilities. The field coverage component anchors the solution to cover the centre of the field as much as possible.

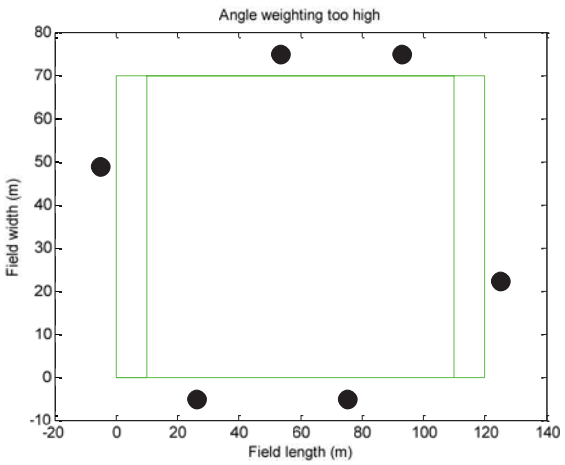

Figure 8.19: Solution with a random rotational component.

[12] Analysed with the fittest solution from the optimisation. The mean performance of each receiver configuration was 8059, 4440 and 3038 respectively with standard deviations of 588, 265 and 190 respectively

[13] These percentages are calculated as a proportion decrease over the whole 90000 positions.

8.6.7 Static Obstructions results

The results for minimised average GDOP in the Dot terrain map with the SBI receiver configuration are given in Figure 8.20. These results show that the entire field is covered by at least three receivers except for 4 locations (shown in the figure with a GDOP of 5, the penalty value). These results also illustrate the shadowing effect of the obstruction; there is a star-like pattern around the obstruction with slightly higher GDOP than the surroundings which is caused by the obstruction preventing LOS from some receivers and thereby reducing both coverage and GDOP. Furthermore, the GDOP is less than 2 in all locations where a position could be calculated and approaches 1 over more than half of the area. Overall, the average GDOP for this solution is 1.04 which approaches the Cramer-Rao lower bound for GDOP.

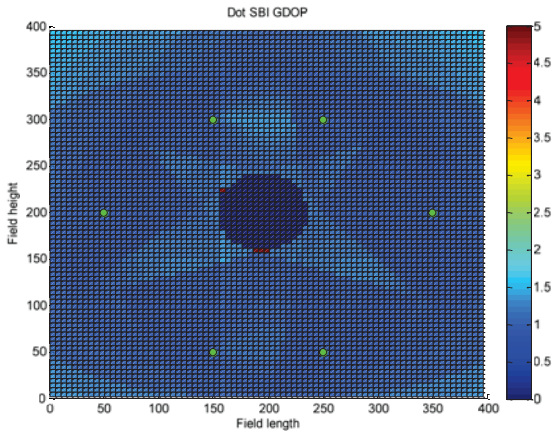

Figure 8.20: Six-receiver, minimum average GDOP, solution by inspection on the Dot terrain map.

The results for the GA optimisation with six receivers on the Dot terrain map are shown in Figure 8.21. The solution is less regular than the SBI but it performs as well with the same overall average GDOP. This suggests that the Dot terrain map does not require optimisation; sufficient, precise coverage can be obtained by spacing the receivers around the perimeter of the obstruction when a single, simple obstruction is present.

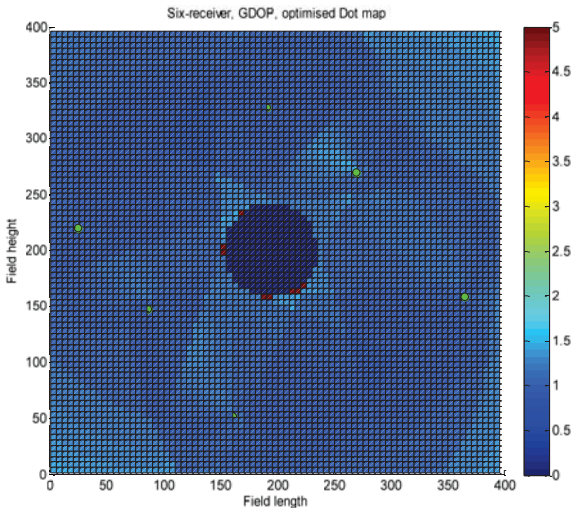

Figure 8.21: Six-receiver, minimised average GDOP on the Dot terrain map.

These results suggest that six receivers are sufficient to completely cover the Dot area with an average GDOP that approaches the Cramer-Rao lower bound.

Examination of the truncated solution space for this terrain map (Figure 8.22) shows distinct valleys caused by penalty values for receivers being placed within the obstruction (the red dot in this case). This is going to be a regular occurrence in the solutions spaces of static obstructions where certain receiver configurations result in an immediate and severe penalty (a receiver being placed within a building). Note that, as in previous situations, the truncated solution space uses the solution by inspection with two of the receivers allowed to move around; here the top two receivers from Figure 8.8. The fittest solutions with this static configuration occur when the mobile receivers are in the top section of the terrain map which, from the SBI results, is where the fitness is greatest. Recall that the receiver displacements shown in the truncated solution spaces are along the vectorised path as opposed to displacement in a single dimension (vectorising is where the two-dimensional map is converted into one single vector comprised of each row concatenated with the previous, thus a 400x400 map becomes a 1 x $16x10^4$ vector). The symme-

try about the major axis is to be expected here since the receivers are identical.

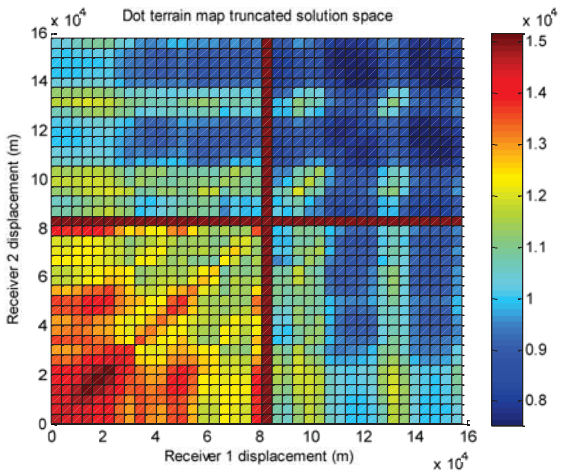

Figure 8.22: Six-receiver truncated solution space on the Dot terrain map.

The Four-Building terrain map offers a more complex scenario. The results for the minimised average GDOP in the Four-Building terrain map with the six-receiver SBI configuration are given in Figure 8.23. It is clear from these results that there are large areas of the field that are not covered by sufficient receivers and that the obstructions are a serious barrier to both field coverage and GDOP. Note especially the regions of high GDOP in the corridors between the obstructions. This increase is occurring because the only receivers with LOS to those areas are all on a straight line which has been shown to be detrimental to GDOP. A tracking system using this configuration would suffer from significant blind spots and low precision when the objects moved through badly covered areas. It is also worthy of note that where three or more receivers have LOS to an area there is GDOP less than 2. This once again confirms the premise that regular configurations of receivers do not produce the best results when complex obstructions are present.

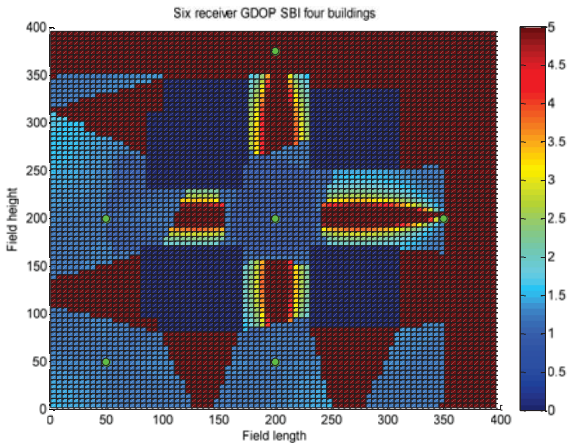

Figure 8.23: Six-receiver, minimum average GDOP, solution by inspection on the Four-Building terrain map.

The truncated solution space for this terrain map (Figure 8.24) again shows a number of ridges caused by incursion of the receivers onto the obstructions (which in turn causes a penalty). Optimum fitness in this case is achieved when one of the mobile receivers is in the middle of the map and the other is in the lower left corner. The mobile receivers are the top left and centre right receiver (with respect to Figure 8.23). This suggests that placement of another receiver near the centre of the map and one in the lower left corner will compensate for the high GDOP in the 'corridors' and will cover the blind spots in the lower left corner of the map.

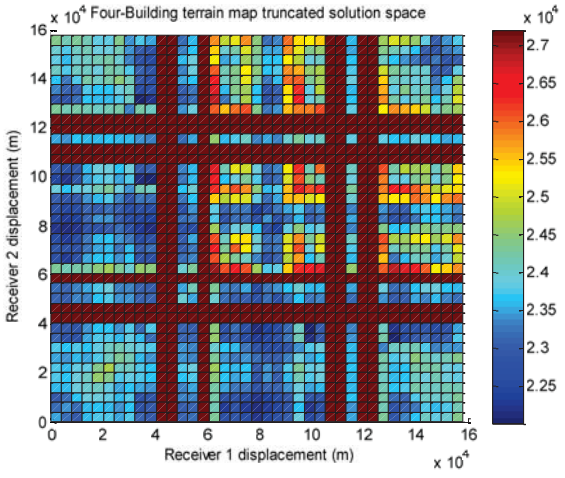

Figure 8.24: Six-receiver truncated solution space on the Four-Building terrain map.

Figure 8.23 suggests that six receivers are not sufficient to cover the entire map with the given obstructions. Experimentation with varying numbers of receivers showed that 10 receivers were required. The results for the GA optimisation of ten receivers on the Four-Building terrain map are shown in Figure 8.25. Here it can be seen that the receivers are spread in a much less regular pattern and there are less areas of no cover. It can also be seen that there are significantly fewer areas where GDOP is high on the connecting line of receivers. In all cases another receiver has been placed in a suitable location to correct the issues of geometry. The average GDOP for this solution is 1.65. This is well within the bounds of acceptability for most tracking systems and a significant portion of that average is from the areas of the map with no cover where the penalty GDOP applies.

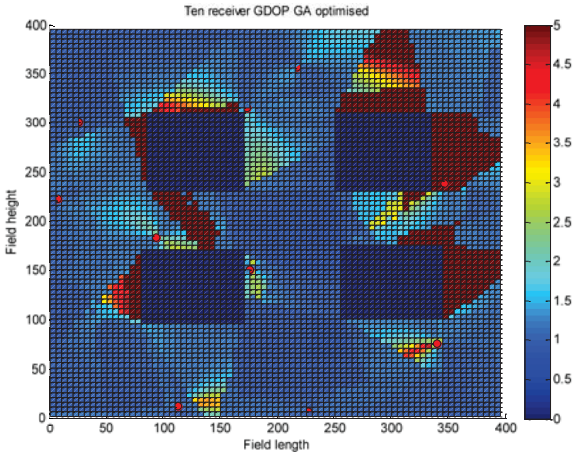

Figure 8.25: Ten-receiver, minimum average GDOP, GA optimisation on the Four-Building terrain map.

The results of the six-receiver solution by inspection on the Riddet terrain map are given in Figure 8.26. Once again it can be seen that six receivers are not sufficient to cover the whole area and that where three receivers have coverage, the GDOP approaches 1. Also note in the lower left of the plot that the high GDOP caused by receivers being regularly placed is apparent. The overall average GDOP for this area is 3.76.

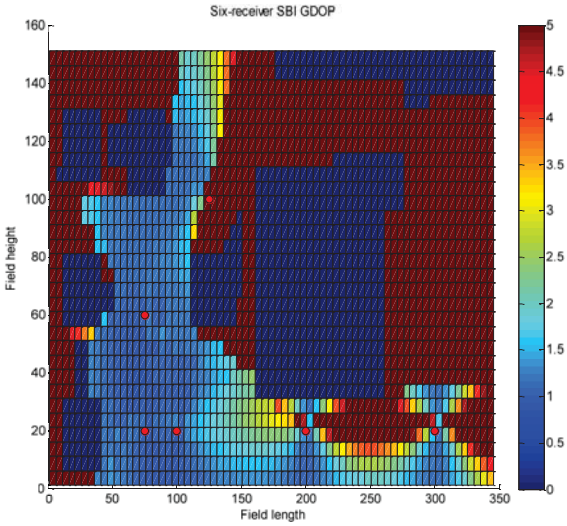

Figure 8.26: Six-receiver, minimum average GDOP, solution by inspection on the Riddet terrain map.

As a comparison with the GA optimisation method, this objective function and terrain map was also optimised using a Nelder-Mead simplex optimisation and a pattern-search optimisation. Recall that both require an explicit starting point and neither has a random element (and thus will produce the same output for a given input). Figure 8.27 shows the results of the Nelder-Mead optimisation.

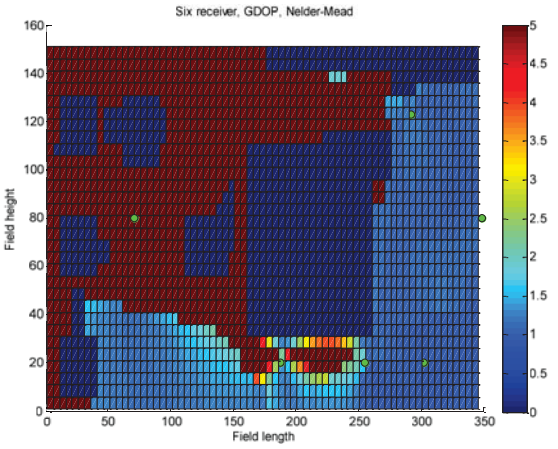

Figure 8.27: Six-receiver, minimised average GDOP Nelder-Mead simplex optimisation on the Riddet terrain map.

Here the optimisation was provided with the SBI receiver configuration as a starting point. The result has not moved significantly from the SBI but the overall average GDOP has dropped to 3.1. Once again, the lower portion of the plot shows three receivers in line which has caused an area of high GDOP along the connecting line. The returned configuration has also placed one receiver (centre left on the above figure) such that it is not covered by another two receivers and thus it is effectively wasted.

It is interesting to note that if the optimisation is given a starting point where the receivers are all near the origin then the optimisation does not progress and the returned 'optimised' solution is the same as the input. This test and the fact that simplex optimisation methods are not suitable for solving multi-modal problems suggest that Nelder-Mead optimisations are inferior to GA optimisations for this type of problem.

Figure 8.28 gives the results of a pattern search optimisation where the start point is given as the six-receiver SBI.

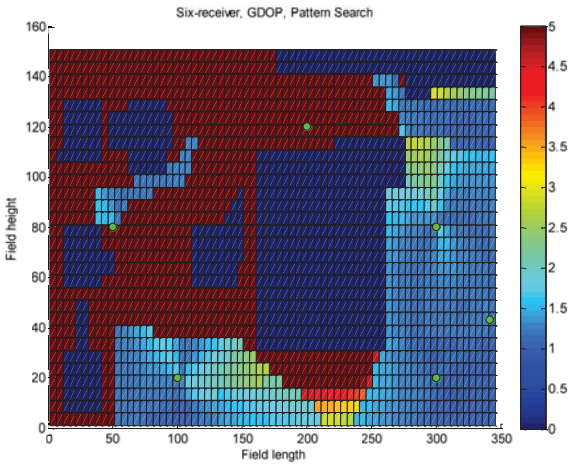

Figure 8.28: Six-receiver, minimised average GDOP pattern-search optimisation on the Riddet terrain map.

The pattern search method has produced a solution similar to that of the solution by inspection and the Nelder-Mead. The overall average GDOP is 3.29 which implies that the optimisation has been less effective than the Nelder-Mead optimisation. This method also appears to have placed two receivers in places where they are of little value (top-centre and centre left receivers). The slightly irregular placement of the receivers on the right side of the area has allowed the GDOP to be minimised.

Note here that the pattern search is considerably more robust in terms of starting point. When provided with all receivers near the origin for the starting point, the optimisation produced a result of similar fitness and configuration to that shown in Figure 8.28.

Once again, testing with various numbers of receivers suggests that ten receivers are required in order that the majority of the area is covered. Figure 8.29 shows the results of a GA optimisation when ten receivers are used. The optimisation has configured the receivers such that all but the areas close to the buildings and near the boundaries are covered. Where there is coverage the GDOP is generally less than 2 and it is only near the boundaries of coverage that the GDOP raises to 3. Overall, the average GDOP for this optimisation was 1.87.

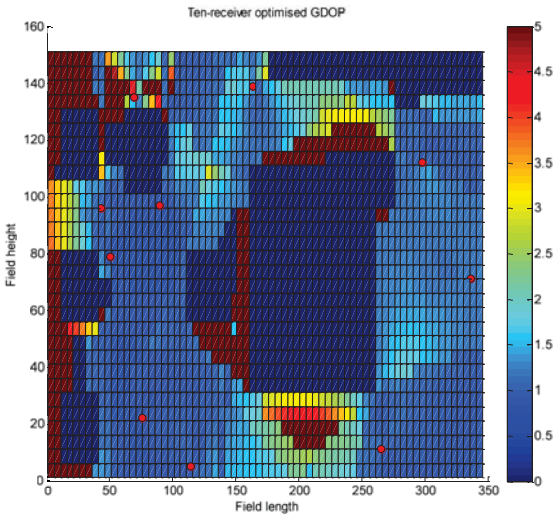

Figure 8.29: Ten-receiver, minimum average GDOP optimisation on the Riddet terrain map.

Given that this terrain map contains the most complex configuration of obstructions it would

be expected that it would have the most complex solution space. This is the case, as Figure 8.30 illustrates. Here again can be seen a turbulent space in which the peaks are caused by receivers being located within obstructions and that individual being penalised. There is a definite trend for the fitness to be best when both receivers are near the centre of their displacement range (near the centre of the terrain map). That the GA was able to find a suitable solution within this space confirms that this method of optimisation is a good choice for this type of problem.

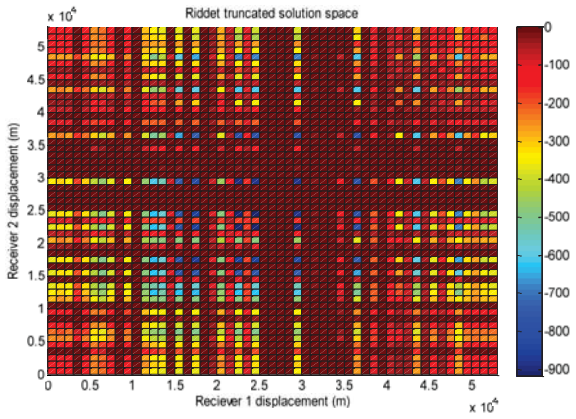

Figure 8.30: Truncated solution space for the Riddet terrain map.

Finally we analyse the results of the optimisation for a known object-movement track. Here the solution has allowed the centre of the arc of the track to be covered such that GDOP is low. The extreme ends of the track however, have been left uncovered or with a high GDOP. It would appear that the decision to place one of the receivers in the middle of the track *(250, 15)* was a mistake since it puts that receiver on the connecting line between two others which has been shown repeatedly to have an adverse affect on GDOP. A better solution would have been to stand that receiver off to the side of the track so that the connecting line between receivers does not intersect the track. The overall average GDOP for this solution is 2.9.

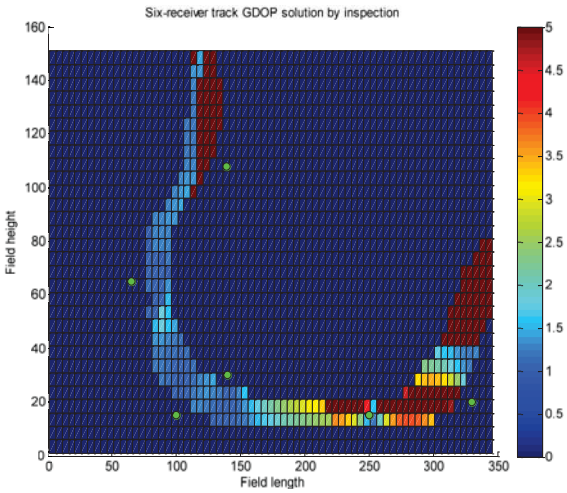

Figure 8.31: Six-receiver, GDOP solution by inspection for the Riddet-Track map.

The GA optimised results for the Riddet-Track map with six receivers is shown in Figure 8.32

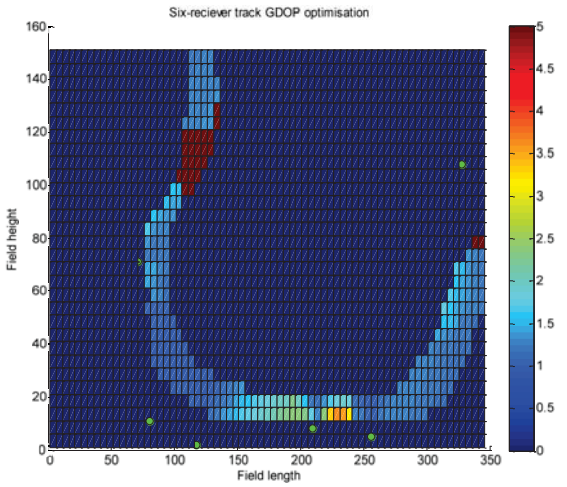

Figure 8.32: Six-receiver, minimised average GDOP for the Riddet-Track map.

The optimisation has correctly deduced that receivers need to be concentrated near the bottom of the area where the large building obstructs LOS from the other receivers. It has also found that by locating one receiver at *(330, 108)* both ends of the track can be covered by one receiver (in conjunction with receivers covering the centre of the track). There is only one section of the track that has significant loss of LOS; more receivers would be able to compen-

sate for this. Furthermore, the optimisation has minimised the GDOP in the lower part of the area such that there are no areas where the track intersects a connecting line between receivers. Once again the non-regular placement of the receivers has prevented valid GDOP from rising above 3. The overall average GDOP for this track is 1.7 which is a 41% reduction due to optimisation.

In an effort to achieve complete coverage of the area, the results for an optimisation using eight receivers are shown below in Figure 8.33. Here the overall average GDOP is further reduced to 1.2 and it can be seen that there is complete coverage and that the maximum GDOP is 3 which is still considered good for tracking applications. In all other areas the GDOP approaches the Cramer-Rao lower bound.

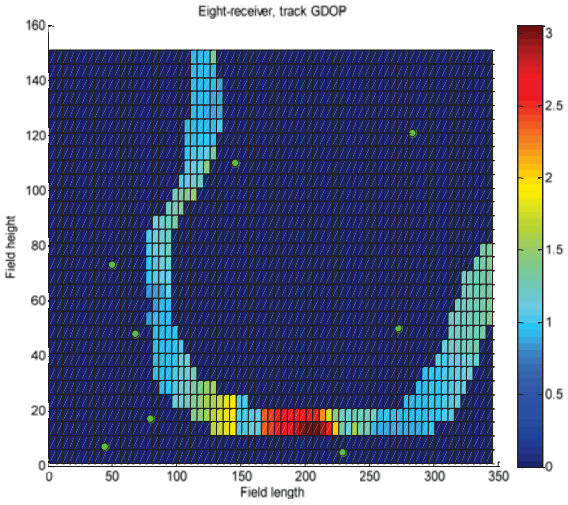

Figure 8.33: Eight-receiver minimised average GDOP for the Riddet-Track map.

8.7 Optimisation conclusions

The GA optimisations have been shown to be an effective tool for finding suitable solutions to the problem of receiver placement for a TOF player tracking system with a limited number of receivers. It has also been shown to be suitable for finding solutions to static-obstruction problems.

Objective function selection will be critical. It was shown that the three functions that required a sample of frames from the data suite could not execute in a timely manner when provided with a suitably sized sample. Reducing the sample size introduces an unacceptable random interference into the results.

The combination of weighted field coverage and receiver separation measures produced the best results for the rugby problem when compared with the solutions by inspection. Minimised average GDOP with penalties was effective for the static-obstruction problem and implicitly moved to maximise field coverage.

The development and testing of these objective functions has demonstrated a link between maximising coverage (to make the system effective for position tracking across an area of interest), obstructions (which cause a reduction of both coverage and precision) and GDOP (where the loss of LOS from a critical receiver has a significant impact on system performance). This link can be applied to all position-tracking systems regardless of application.

The static-obstruction problems especially have confirmed the premise that receivers configured in a regular pattern may not yield the best results (research question 2), particularly when there are complex formations of obstructions present. Even small deviations from the regular pattern can result in significant gains in performance (especially precision).

Using the information provided by the optimisation, it should now be possible to deploy any TOA position-tracking system to analyse the movements of any kind of object. Also, by utilising data-reconstruction techniques, the requirement for human intervention in the tracking process is minimised.

9 Position Reconstruction

Signal loss is inevitable – causes for signal loss both general and application- specific – methods for reconstruction – data processing and formatting.

It is highly likely, especially in sporting applications that at some time there will be a period of loss of signal from one or more player tags during the game. The likelihood is increased when a minimalist system (one that uses the minimum amount of equipment required to calculate a position) is implemented and two conflicting and mutually exclusive points of view occur. The first is that in the sporting application, a position is required for every player for every second of play (see Specifications chapter). The second is that the system will experience periods of signal loss for various reasons (see section 7.2). By using knowledge of the behaviour of tracked objects in the system (players in this case) it is possible to recover some of the missing positions based on known positions of the objects in the past and future, and the positions of other objects.

The performance of minimalist systems can be increased in this way. The model presented here uses software as an alternative to the brute-force method of adding more receivers to the system. The application of such a model is especially important in the rugby situation where each player acts as a moving obstruction for every other player.

There are two general methods for recovering missing positions: real-time and post-processing. In a real time recovery system, only the current and previous positions of all of the objects can be used for correction, as no future positions are available. This limits the potential effectiveness of the reconstruction. Also there is a time restriction; any required correction must be calculated and applied before the next position is measured or the corrected data stream will lag behind the current data. Post-processing generally involves taking a log of all of the position data and performing the position reconstruction at a later time. Post processing can generally yield better results, since the process will have access to both past and future positions. The decision was made to use a post-processing algorithm for this work, because more accurate results are desired over immediate correction.

Recall that position data are output when the signal from a tag is received by a sufficient number of receivers (see section 5.3). Thus the data are ordered in a first in first out fashion. For statistical analysis by the Industrial Partner's software it is necessary that the data be collected into frames (which can then be synchronised with a video feed). Within each frame (the duration of which is variable, but defined by the Industrial Partner) is a placeholder for each object to be tracked. Where position-data exist within a particular time period a position can be output in the relevant place in the relevant frame. Where no data exist the placeholder remains empty thus identifying a missing position. Where more than one position for a single object exists within the current time interval in the input, some form of averaging is required so that a single position can be recorded in the output. The code implementation of this initial data processing is included on

the attached DVD (see Appendix C: DVD index for the path to this code).

Separation of the data into frames with a single position for each object within each frame is also the format that would be most useful in the real world application of player tracking. Thus this data format is retained throughout the whole reconstruction process.

As has been mentioned, separation into frames gives an indication of the reliability of the position data. The reliability of the calculated position is proportional to the number of individual measurements that were averaged to find the position. For example, if the position of an object in a particular frame is found from 50 measurements in the input data then it can be said to be more reliable than a position found from only 5 measurements. Similarly if there are no measurements for an object within a specific frame, then the position is easily identified as being missing. With this reliability measure it is possible to set a threshold for reliability within the reconstruction algorithm. If the reliability of the positions of an object are below a threshold then some form of correction is required, and the opposite is also true; that positions with high reliability values need not be corrected.

In all situations where the position of an object is missing within a frame or series of frames, there is a limit to how far behind or ahead it is prudent to look and still be confident of getting an accurate reconstruction. The selection of window size will depend on the application in question, the interval over which the object was missing, and the corrections that have been used previously.

The methods of reconstruction or correction of the position data can be classified into two categories: those that are general and can be applied to any position-tracking situation, and those that are application specific. The methods are described in their relevant categories in the following subsections. The application-specific functions here relate to the real world application of player tracking in rugby union.

9.1 General position reconstruction/correction methods

These methods can be applied to any application of position tracking. Within each method will be described the logic behind the identification of the situation and the assumptions made that allow the position to be reconstructed or corrected. Discussions on the suitability of each method for different signal loss situations will be given.

9.1.1 Linear interpolation

The simplest method for reconstruction is to assume that the object moves in a straight line between signal loss and reacquisition (the signal loss interval or SLI). The reconstruction of the missing data is a case of interpolating the required number of positions across the SLI. For example, if the signal from an object is missing for three consecutive frames then the reconstruction consists of three positions spaced evenly between the known positions that bound SLI. Figure 9.1 illustrates this.

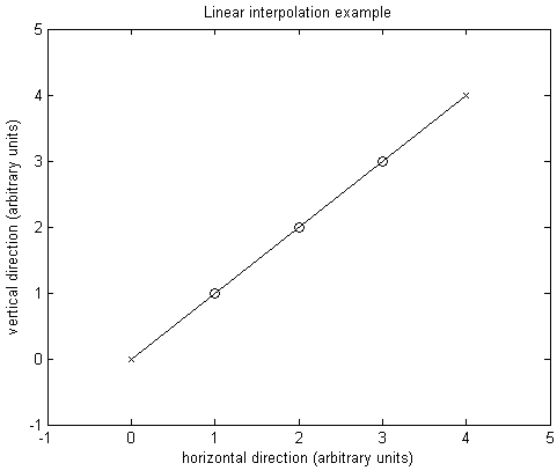

Figure 9.1: Positions of a single object reconstructed using first order linear interpolation.

Here the crosses represent the known positions that bound the SLI and the circles are the reconstructed positions.

This is an example of a first order interpolation where it is assumed that the player has a constant velocity.

Linear interpolation is appropriate only in situations where the SLI is small enough that the object is unlikely to have deviated significantly from the straight-line path. If linear interpolation is performed on a SLI that is too large, then the resultant reconstructions may miss critical movements of the object, thus reducing the accuracy of the reconstructed data. Figure 9.2 illustrates how linear interpolation can inaccurately reconstruct the data when the SLI is too large.

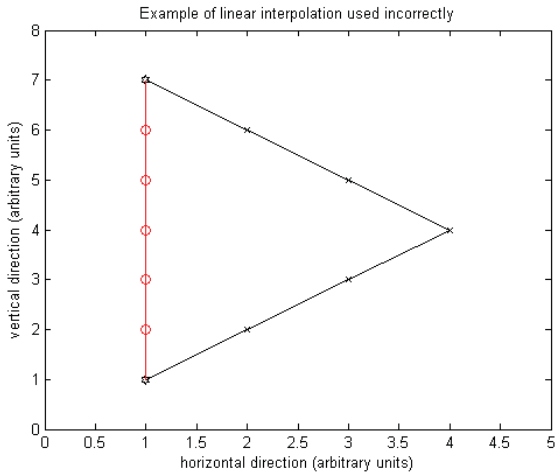

Figure 9.2: Linear interpolation used in a situation where the SLI is too high.

Here the hexagrams represent the boundaries of the SLI, the true – but unknown – path of the object is shown in black with crosses at each frame, and the reconstruction is shown in red with circles at each reconstructed position. Note that while the true path of the object deviated 4 units to the right and then back again, the interpolation process calculated the path as a straight line between the boundaries of the SLI, hence linear interpolation was inappropriate for this situation. The exact size of the maximum allowable interval is application dependant.

Linear interpolation is especially useful where the signal fades in and out quickly. That is, situations where there are many small SLI's. These situations are expected in rugby where the player is moving through a large group of other players (a precursor to a ruck/maul formation).

If the SLI occurs at the beginning or end of the data – and assuming that the interval is short enough – then extrapolation based on two known points can be used in place of interpolation. Again, the selection of appropriate points is application dependent.

It is possible to use higher order interpolations, however, these require more known positions. A second-order interpolation for example would use two points either side of the SLI to interpolate the velocity vector of the player and hence calculate that player position on a curved path. Figure 9.3 shows an example of a second order interpolation and compares it with the linear interpolation of the same data. Here the black marks indicate the known path of the player, the red marks indicate the recovered position from linear interpolation, and the blue marks are the recovered position from second-order interpolation. In this case the true position of the player is at *(5,5)* where the second order interpolation has calculated.

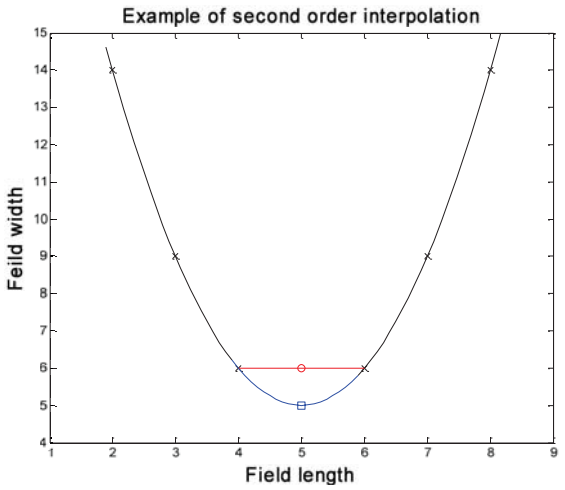

Figure 9.3: Example of second order interpolation.

It is clear from this example that the second order interpolation has more accurately calculated the position of the player as compared with linear interpolation (there is a 1 unit difference between the linearly interpolated position and the true position). However, this example is a worst case scenario when comparing accuracies; the signal was lost when the player was at the sharpest point of the curve (i.e. maximum acceleration). Had the SLI been in a different location, the inaccuracy would not have been so high.

Also, the example given in Figure 9.3 depicts a player moving along a smooth path. It is assumed from knowledge of the game that the players will move along such smooth curves only when in the open (i.e. dashing down field when there is no opposition) and that in these situations there will be either adequate coverage or the path will be sufficiently straight for linear interpolation to be used. It is recognised that there will be times when the player encounters obstacles (e.g. other players) and that the player will make evasive movements; however, it is likely that these evasive movements will be discontinuous and that second-order interpolations will not perform better than linear interpolations. Consider the following example. A player performs a sidestep (a momentary lateral movement to the general direction of motion) to avoid an obstacle. If the player's position is lost during this sidestep then neither linear nor second-order interpolation will be able to accurately recover the data. Figure 9.4 illustrates this example (assume the player is moving from left to right). Once again, the true path of the player is shown with black markings (the missing path is shown with a green line), the linearly interpolated path is shown in red and the second-order interpolation is shown in blue (note that the blue line has been overwritten by the red in this figure because they are both the same).

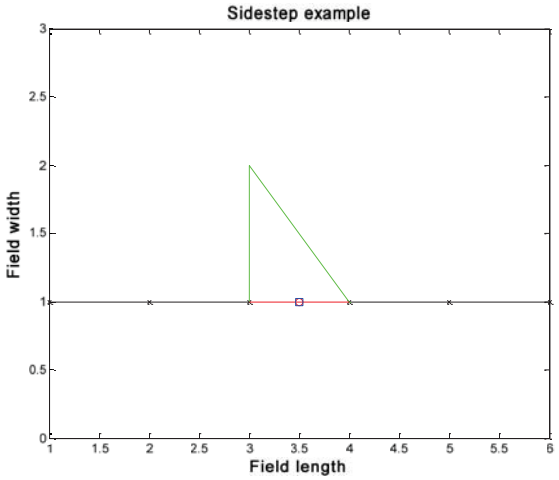

Figure 9.4: An example of interpolation when the player is on a discontinuous path.

In this example both interpolations have produced the same result. Here the missing datum is a discontinuous deviation from an otherwise straight path. Once again this is a worst case scenario; however, if the SLI were to either side of the sidestep manoeuvre then the second-order interpolation would still be outperformed by the linear interpolation. Regardless of which side of the sidestep was lost the linear interpolation would produce a point along the straight line; however, due to the discontinuity in velocity, the second-order interpolation would calculate that the player moved in an arc further away from the true position. Figure 9.5 illustrates this; the same symbols are used.

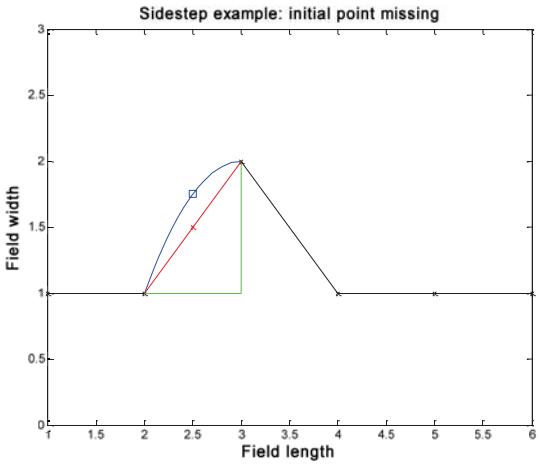

Figure 9.5: Sidestep interpolation with the initial point missing.

Here the point where the player began the sidestep manoeuvre is missing (once again, assume the player is moving from left to right). The second-order interpolation places the player further away from the true position than the linear interpolation does. This occurs as the second-order interpolation reconciles the velocities on either side of the SLI.

Finally, it has been noted through these examples that the second-order interpolation requires at least two positions either side of the SLI in order to function properly. In situations where position data is experiencing large numbers of low SLI invisibilities (data from the player visibility model suggests that most invisibility will be like this) this may result in there not being two contiguous, valid positions to either side of the invisibility. It would be possible to use two positions that are more spread but this may affect the accuracy of the results.

Linear interpolation was selected for use over higher order interpolations in the reconstruction process because it requires only one known position either side of the SLI and it was to be applied only to small SLI's to maintain reliability in the data. Rugby players tend not to follow gently sweeping paths in situations where they are likely to be invisible to the tracking system.

9.1.2 Speed

The objects to be tracked will have an expected maximum speed with which they can realistically travel. For example, with the exception of some world-class sprinters, humans cannot move faster than $10 m.s^{-1}$. Given that the positions of the objects are known across two frames and the time interval covered by a frame is known, the speed of the object can be calculated. Should the calculated speed exceed the set maximum then it is clear that some error in positioning has occurred and that position will be marked as erroneous.

The correction of this error will depend on which position caused the speed to exceed the maximum. Either the initial position is too far retarded of the true position or the second position is too far advanced (assuming the player is running in a straight line). In this situation it is most likely – but not always true – that the position with the lowest reliability value is at fault. Measurement of speed on the positions to either side of this suspect position can confirm this. Once the offending position has been identified it can be removed and a more reliable position calculated using linear interpolation.

9.1.3 Multiple tag data

In the situation where the players are wearing multiple tags, loss of visibility of just one tag will not result in a visibility loss for player. Data from the visible tags can be used to locate the player. However, in the event of all of the tags becoming invisible one of the other reconstruction methods must be applied to the data bounding the SLI.

Multiple tags per player is provided here as a way to recover position data on a particular player in the event of a single tag becoming invisible to the tracking system. While it is an option for increasing the visibility of a player, it contravenes the requirement for portability (see section 4.3) and it is assumed in the reconstruction that players wear only one tag each.

9.2 Application specific position reconstruction/correction methods

These specific methods are based on the behaviour of rugby players. There are several player formations that are common and likely to cause obstructions. Knowledge of the formation shapes can be used to reconstruct missing positions. As has been shown in the player-visibility model (chapter 7) the loss of signal from one player due to another being within the line-of-sight path is a major cause of signal loss. This is aggravated by the tight clustering of many rugby formations (see Figure 9.6 through Figure 9.10 for examples of these formation types), but suggests that a significant number of missing points could be recovered if these common formations could be identified.

From examination of video footage of rugby games and knowledge of the behaviour of the tracking system, it has been concluded that the situations that are most likely to cause a loss of signal are rucks, mauls, scrums and lineouts. The following subsections detail each formation and which players are likely to experience signal loss, the method for detecting the formation from the available data, and the recovery solution.

9.2.1 Rucks and mauls

The International Rugby Board (IRB) defines a ruck as follows:

> *"A ruck is a phase of play where one or more players from each team who are on their feet, in physical contact, close around the ball on the ground. Open play has ended."[157]*

Figure 9.6 shows a ruck formation; note how many of the players are off their feet. Note that while the rule above states that all players are on their feet, in almost every ruck there will be players who fall to the ground (allowable by the rules if it is unintentional).

Figure 9.6: A ruck

In terms of tracking, a ruck is a group of bodies in a tight formation around the ball. Some players involved may fall down and may have the signal from their tag directed towards the ground (this will depend on the placement and orientation of the tag), making the signal invisible to the receivers. Furthermore, if a large number of players become involved, the bodies of the players on the outside will shield signals from players near the centre.

From inspection, a maul appears similar to a ruck except that the players must remain on their feet at all times, and the formation moves around the field (see Figure 9.7). Should the formation collapse it ceases to be a maul and becomes a ruck. Once again, from the perspective of tracking, the same possibilities for obstruction exist in a maul as in a ruck. Given that the recovery of positions is being done on a frame-by-frame basis, there is no difference in the identification and handling of mauls compared with rucks.

Figure 9.7: A maul.

It is assumed that the players with the least probability of being detected will be nearest the centre of the formation. Also, it is assumed that a significant number of players on the periphery of the formation will be visible to the receivers. Thus, if a player is seen to be approaching a tight formation of players and the signal from that player disappears it can be assumed that the player has entered the formation. If the signal from the lost player does not reappear in subsequent frames it can also be assumed that the player has not left the formation. When the player does leave the formation, or the formation breaks up, the signal should become visible again.

In order to implement an identification procedure for rucks and mauls it is necessary to find tight formations of players near the last known position of the lost player. Specifically, formations near to the path along which the player was moving (found by extrapolation) are the most likely candidates. The formations themselves are found by mapping the distance of every player to every other player for that frame. Where those distances are small for three or more players a formation is defined. The definition of 'near' relates to the maximum speed with which a player can move in the time of a single frame; any formations outside this range are too far to be candidates. If a formation satisfies the requirements then a ruck or maul has been identified.

The calculation of a recovered position for the player entering the ruck presents a problem. The player could conceivably be anywhere within the ruck or maul and if the position of the player is missing the entire time spent within, then the achievable best in recovery is an arbitrary position for the missing player within the ruck, in this case the centre of the formation (defined by the visible players at the periphery). If any of the players become visible while in the formation then this datum can be used to define the position of that player. The linear interpolation recovery method would be effective when receiving intermittent signals from players within a ruck or maul.

9.2.2 Scrums

A scrum is a set-piece formation that is used at the beginning of a play after a minor infringement. Eight players from each side form into three rows each and, upon the command of the referee, engage to contest possession of the ball, which is injected into the formation at ground level between the front rows by a player outside the formation (usually the player in the half-back position). Figure 9.8 illustrates this description.

Figure 9.8: A scrum, about to engage.

Figure 9.9 gives the relative player positions within a scrum. This is the standard formation for a scrum. Substitutions and rule infringements leading to players being sent off the field may change the configuration of players within the formation. Furthermore, scrums are not limited to eight players and the number-eight player may sometimes move one position left or right. However, the arrangement of players within a scrum is always the same.

Figure 9.9: Locations of players from each team in a standard scrum formation[158].

Given the proximity of the players in this formation, it is reasonable to assume that one or more players will lose visibility, especially near the centre. The number of players that go missing may be limited with careful placement of the transmitter tags. The players involved in a scrum are bent at the waist during the formation, their backs facing up. If the tags were placed here, it is possible that the signals from the players would remain clearly visible to the receivers.

In order to identify that a particular player is missing owing to the formation of a scrum, that player must first be one of the players that is involved in a scrum (player numbers 1 to 8 from each team). Should the missing player not be a usual participant in a scrum it is reasonable to assume that they are not missing because of involvement in the scrum and hence cannot be dealt with under the scrum-reconstruction rules. If the other players that are usually involved in a scrum are either missing or fit within a bounding box that defines the approximate size of a scrum then it can be assumed that a scrum formation is occurring.

A template containing the relative positions of players in a scrum can be used to recover missing scrum players. The template can be deformed and translated to best fit the visible scrum players (a technique known as template matching). Given that there is negligible movement of scrum players relative to one another, the positions of the missing players are defined by the adjusted template. A minimum of two scrum players being present within the bounding box is required so that proper orientation and scaling of the template is possible. The position of missing players is taken from the deformed template.

The selection of template matching algorithm is not important for this research. The reconstruction algorithm is designed in a functional-block manner and so one technique can easily be substituted for another. The technique used here was chosen for expedience.

9.2.3 Lineout

A lineout in rugby occurs whenever the ball goes outside the bounds of the field. Lineouts occur at the start of play but will swiftly become a maul or break up. Figure 9.10 shows a lineout formation in progress.

Figure 9.10: A lineout formation.

The formation itself consists of several players from each side lined up perpendicular to the sideline where the ball was taken out of play and between 5m and 15m from the sideline. The teams have one row of players each, separated by no less than 1m of clear space. The exact number of players involved in the lineout is determined by the throwing team, with two players each side being the minimum. The players that form the lines in this formation will be packed front to back which could cause a number of player obstructions.

Identification of this formation is difficult because of its variable nature. There is no set number of players that constitute the formation, no set positions within the formation, and the positions of the players may change throughout its duration, and therefore the behaviour of players within a lineout is unpredictable. An interpolation method is the best way of dealing with lineouts.

9.3 Human intervention

Occasions may occur when the data are so severely degraded that it is impossible to use behavioural models to reconstruct the positions. In these situations it is best to mark the position – or interval – as requiring human intervention. Once all automated reconstruction is completed a human user can examine any flagged positions and correct them as necessary.

This form of reconstruction will require that the user have access to some other record of the objects' movement such as video footage. In the application to rugby this is not an insurmountable problem, as the footage of the game will exist and can be obtained with minor effort. It would not be appropriate in a factory setting, for example, as it would require that a video recording system be installed in addition to the position-tracking system.

In all cases, the requirement for a human to correct the missing positions manually is undesirable as it is both time and labour-intensive, thus breaking the requirement that the system be able to produce useable results as fast as possible (see section 4.8).

9.4 Reconstruction order

The order in which the reconstruction methods are applied will affect the performance of the whole reconstruction. It was decided that the methods would be applied in such a way that the earlier reconstruction types would deal with the majority of signal invisibilities and be the least computationally intensive (similar to the implementation of the player visibility model in chapter 7).

First, a check of the speed of the players is made to ensure that the data present are valid; positions which are not are marked as missing data. Next, interpolation of small SLI invisibilities is performed; this may include correction of positions calculated as being too fast. Detection and correction of players lost due to scrums is per-

formed next. Finally players who are invisible due to a ruck or maul formation have their positions reconstructed.

Each reconstruction method is applied once only per application of the recovery algorithm. The block diagram below shows the flow of execution of the algorithm.

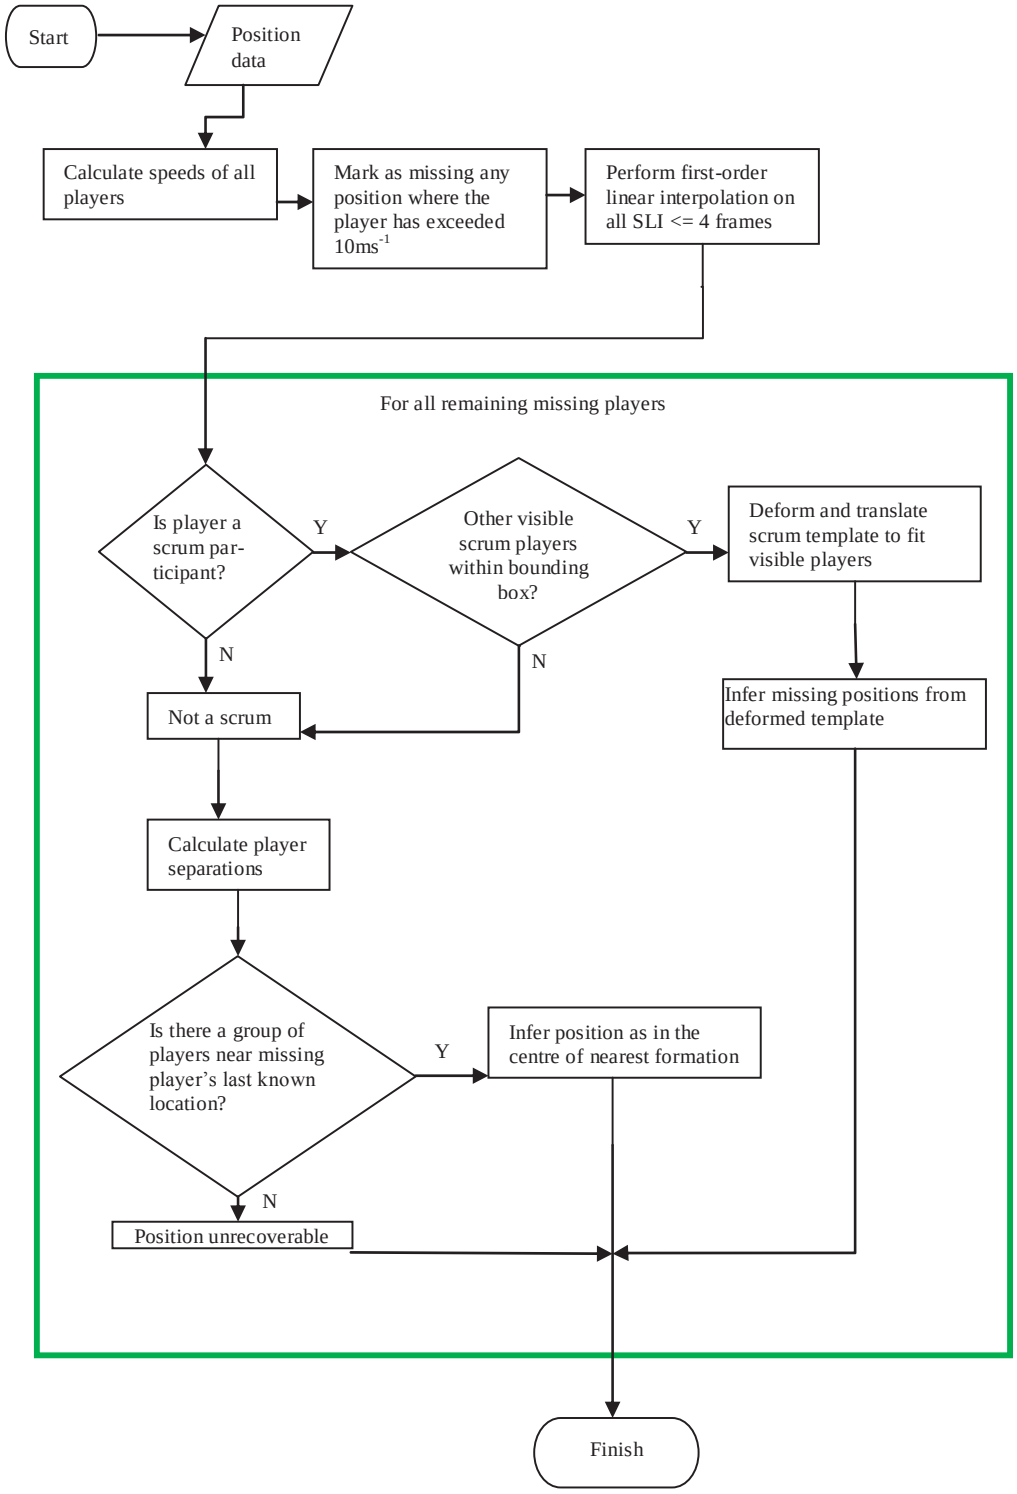

Figure 9.11: Block diagram for the reconstruction algorithm.

9.5 Reconstruction results and discussion

Data from field trials with the Multispectral tracking system and the pre-tracked game data from the Industrial Partner were used to test the effectiveness of each of the position reconstruction methods. For each application-specific reconstruction, ten groups of ten frames have been selected from visual analysis of the pre-tracked data. Each group is known to contain the desired formation to be reconstructed. The reconstruction algorithm must be able to locate and correctly recover the missing player positions.

Unless stated otherwise, each method was implemented on a 6-receiver system with the receivers located as in Figure 8.4(a), and at a height of 2.5m. Some statistics will be given on the effectiveness of these reconstruction methods when applied to a rugby game.

9.5.1 Linear interpolation

Data from field trials with the Multispectral tracking system show that this simple reconstruction method can be used effectively in the reconstruction of missing player positions.

Several trials were conducted with the tracking system in which a test player wearing a tag moved along the field lines of a rugby pitch. Figure 9.12 shows the raw data recorded by the tracking system. In the corners of the field it can be seen that the density of positions is not as great as near the centre line, indicating that there are some points missing.

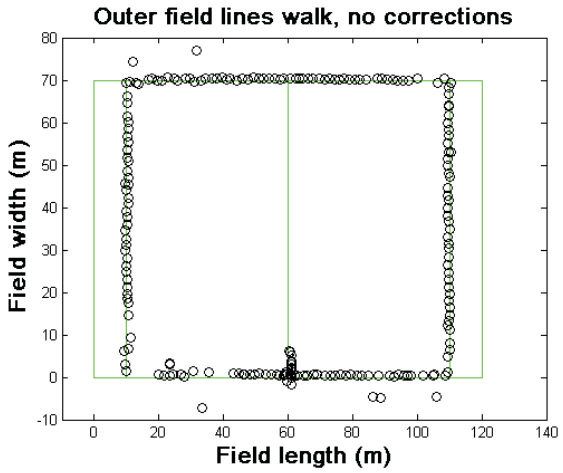

Figure 9.12: Player moving around the perimeter of the field. No corrections.

Applying linear interpolation reconstruction to these data fills in some of these missing areas. Figure 9.13 shows the effect of linear interpolation reconstruction. 4 seconds was set as the maximum SLI window, corrections cannot be made where the player is missing for longer than this.

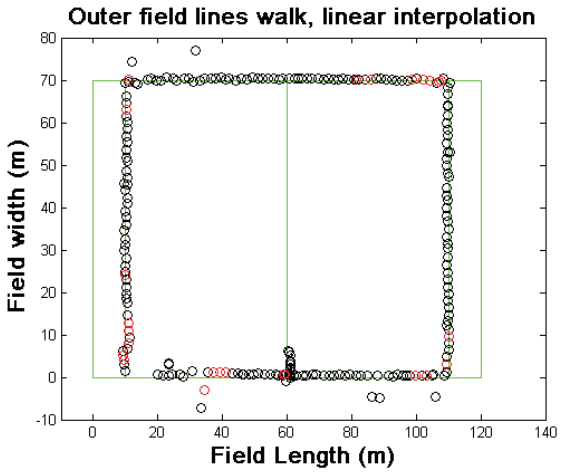

Figure 9.13: Applying linear interpolation to recover missing positions.

Comparing this with Figure 9.12 it can be seen that most of the areas that were missing positions have been filled in. Furthermore, the interpolated positions are on the field perimeter (the true path of the player). The only exception to this is in the lower left corner of the field. Here there is a discontinuity that spans 10m of

field length; this SLI could not be corrected because it was too large.

It can be seen in Figure 9.12 that the player has been located well outside the field on several occasions. Given that the player's track is known to have followed the field lines, these points are in error. Any erroneous position recorded in the data will have an adverse affect on an interpolation reconstruction process, especially if that position appears in the middle of an SLI. Such an event will cause the interpolation algorithm to calculate erroneous player positions, thus reducing the accuracy of the reconstruction as a whole. This can be seen to occur in Figure 9.13 in the region of *(30, -8)* where the error position has preceded an SLI. At other times when the tracking system output inaccurate positions, they have not occurred during signal loss.

The erroneous position calculations were caused by a combination of the tag signal reflecting off stadium light posts and the signal being visible to only three receivers. The position calculation then converges to an erroneous location. The test site at Massey University's Ruby Institute is a cramped environment in this sense. In a real game situation it is very unlikely that there will be light posts this close to the field. This does, however, highlight the fact that interpolation alone is not a complete solution for reconstruction.

From these test data it was calculated that there were 241 frames of 1-second duration. 86% of these data were recorded as valid positions (including the erroneous calculations), 11% of frames were reconstructed using interpolation and 3% of frames were unrecoverable. This implies that approximately 79% of the missing data were recovered (this percentage is calculated as the number of recovered positions with respect to the number of positions missing in the original data).

9.5.2 Speed

Remaining with the data shown in Figure 9.12, the erroneous points can be removed by applying the maximum speed limitation to the player. Figure 9.14 shows the result of applying a limit of 7ms^{-1} and then interpolating.

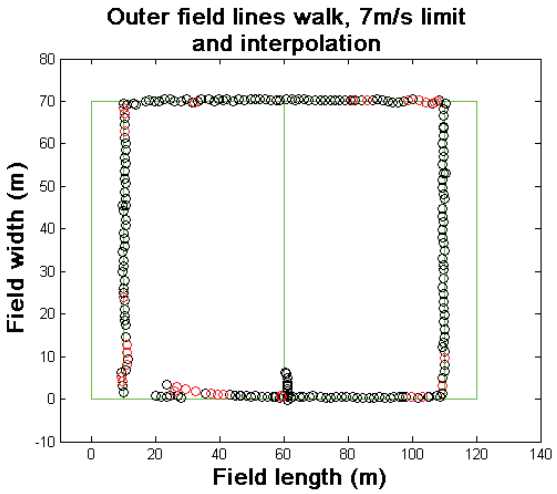

Figure 9.14: Positions filtered based on player speed and linear interpolations applied.

It can be seen that the path now much more closely follows the true path of the player and the majority of the erroneous data have been removed. Selection of the appropriate maximum speed is very important for correctly manipulating the data. When maximum speed is over estimated, some unwanted data will be included and an underestimate will exclude valid position data. The value of 7ms^{-1} was chosen here after consultation with the Industrial Partner and with reference to their accumulated statistical data. The value is the upper quartile of player speeds when sprinting (i.e. data relating to player speed when standing or taking part in a formation such as a ruck was not considered). The upper quartile was chosen to be inclusive of the top speeds achievable by the majority of players.

In Figure 9.14 there is still one position (located at approximately *(20, 5)*) that has deviated from the known track of the player. It is clear from the reconstructed data that this erroneous point

has occurred in the middle of an SLI and it does not contravene the maximum speed limit. Detection of such an error would be very difficult without prior knowledge of the true path of the player. Human intervention will be required to correct this inaccuracy.

The statistics for maximum speed limiting indicate that only 83% of the data contained valid positions (recall that previously the erroneous positions were considered valid data), 14% of frames were recovered using interpolation and 3% are without any reconstruction. Compare this with the results of pure interpolation and it can be seen that the proportion of frames that could not be recovered remains the same but some of the data recorded as valid have been defined as in error and have been corrected using interpolation. While there is no increase in data recovery, by imposing a maximum speed limit, the accuracy of the data was increased.

9.5.3 Multiple tag data

Figure 9.15 shows the track of a single player wearing two tags as they move in a circuit of one half of the field. Two symbols have been used to represent the tracks of individual tags.

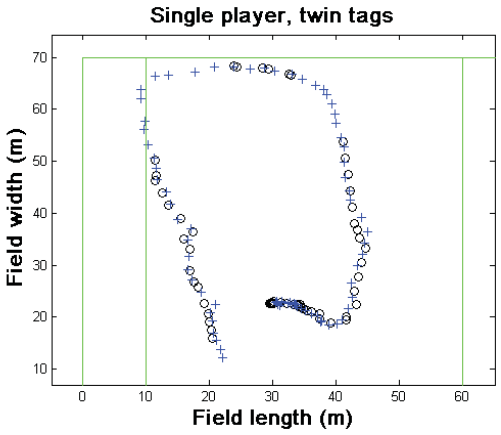

Figure 9.15: Single player wearing two tags moving around the field.

It can clearly be seen that there are regions where one tag (black circles) is not visible to the tracking system; and in many of these locations the second tag is visible. It is reasonable to assume that the system operators will have knowledge of which tag is on which player so it is a simple matter to combine the data from the tags into a single track for the player. Figure 9.16 shows the combined track.

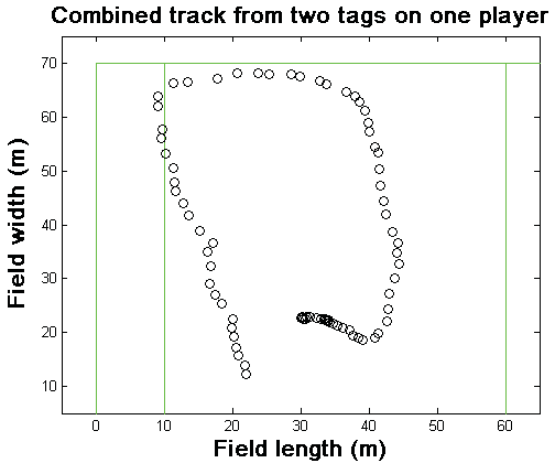

Figure 9.16: Combined track of a single player wearing two tags.

Examining the statistics from the combined tracks reveals that all of the missing positions from one tag were reconstructed using data from the other tag.

While using multiple tags is a viable option for increasing the visibility of a particular player, it will increase the cost to operate the system, which may interfere with the minimal cost requirement (see section 4.5) and more tags will have greater physical interference (see section 4.1). Furthermore, extra tags will need to be managed and attached to the players correctly; this will take more time and adversely effect the ease of use requirement (see section 4.7).

The effect of having data from multiple tags is weighed against the first research question. It may be that adding more transmitters (relatively small and cheap at approximately $20 each) is an acceptable way of realising the minimalist tracking system that performs satisfactorily. On the other hand, the software tools for the third research question will also help to realise the suitable minimal system.

9.5.4 Rucks and mauls

Figure 9.17 shows a ruck formation in which one player has become invisible to the tracking system owing to player obstruction (this player is highlighted in red). This section of a frame was taken from one of the ruck example groups.

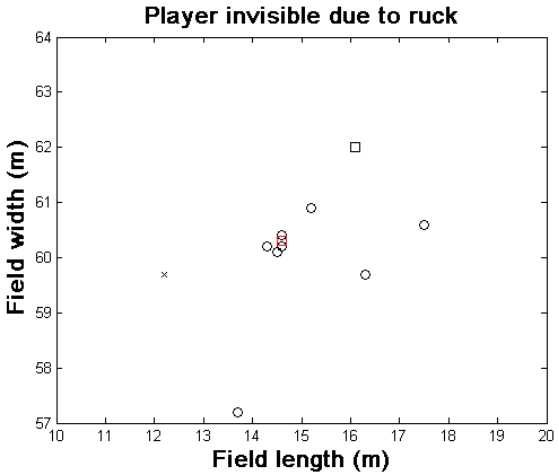

Figure 9.17: Player loses visibility when entering a ruck.

The reconstruction algorithm was able to successfully identify the formation and recover this player's position. Figure 9.18 shows the reconstructed player location as a blue circle. The true position of the player has been retained to give a visual indication of accuracy. The reconstructed position is 0.14m from the true player position.

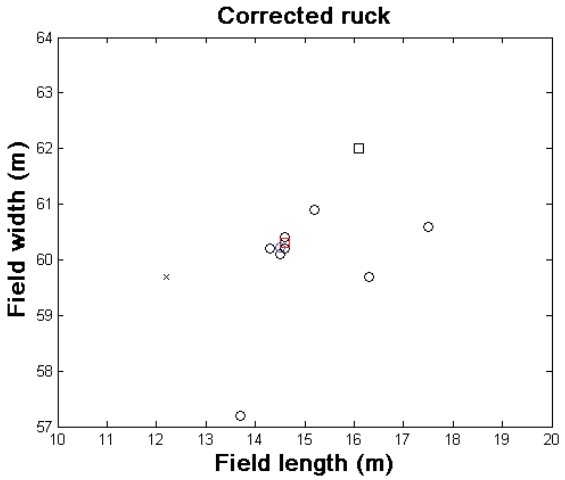

Figure 9.18: Missing position recovered through detection of a ruck.

Recall that the reconstruction algorithm assumes that the missing player is in the centre of the formation, thus the accuracy of ruck reconstruction in general will depend upon the number and locations of other players involved in the ruck. For example, if the ruck were more spread out and the player was only at the periphery when visibility was lost, the reconstructed position at the centre of the formation would be less accurate. However, if the players are too spread out (more than 1m from the centre) then the formation does not qualify as a ruck. Using this logic, the worst-case error in reconstruction will be 0.7m ($\sqrt{2}/2$, the diagonal distance from the centre of a square to one corner).

The ruck recovery algorithm was able to identify and correctly reconstruct player positions in a ruck formation in all ten of the test cases. The positions were reconstructed with a 0.20m average error and a standard deviation of 0.04m. As hypothesised, variation for this method of recovery is affected by the size of the formation; however, the reconstruction algorithm was able to recover positions at well under the theoretical worst case error and also with better accuracy than was required from the measurement accuracy specification (see section 4.6).

9.5.5 Scrums

An example of a scrum with several players invisible due to player obstruction is shown in Figure 9.19. In this figure, black circles are visible players and black squares are invisible players (recall that with the simulation data we have access to the true positions of the players even though they may be calculated as invisible). The scrum template as it would be applied is superimposed over this formation with blue triangles.

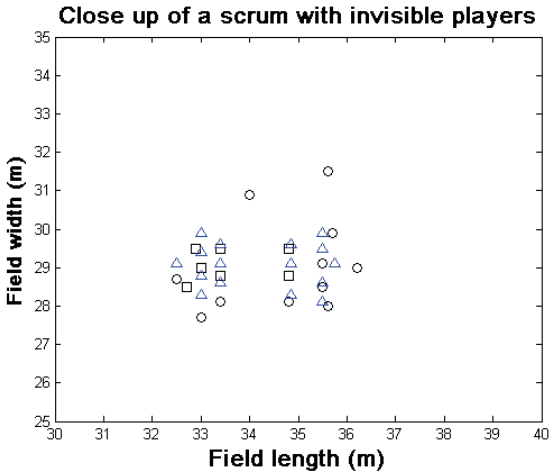

Figure 9.19: Scrum including some invisible players with the scrum template superimposed on the formation.

It can be seen that there is some difference between the actual positions of the players (black symbols) and the template positions (blue symbols). It is unreasonable to assume that the players will construct the formation with the exact same positions as the template for every scrum. Note that there are two visible players above the formation in Figure 9.19. These players are the halfbacks for each team; they inject the ball into the scrum but are not part of the formation.

The scrum detector was able to correctly identify the formation and recover the missing positions. The results of this reconstruction are shown Figure 9.20 as red squares overlaid on the original data (note that these have been taken from the translated and scaled template).

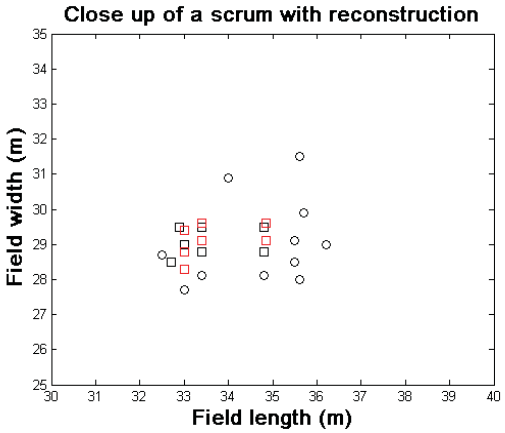

Figure 9.20: Scrum with reconstructed positions.

The error in reconstruction is 0.22m in this case. The true positions of the players appear quite different from the template. This is attributed to inaccuracies in the Industrial Partner's manual tracking system.

In general, the recovery algorithm was able to identify the scrums in each of the ten test groups and it was able to correctly recover the missing player positions. The average position reconstruction error across all ten examples was 0.26m with a standard deviation of 0.06m. This position error seems low given the number of players involved; however, a scrum is a set-piece formation that is performed the same way each time. Also when the scrum either collapses or ends it ceases to be classed as a scrum by the reconstruction algorithm. This is intended since a template match is not a suitable way to recover positions missing in that situation since the player positions will not conform to a template (in that case the formation will be classed as a scrum).

The 0.26m error is within the required limit for accuracy of the system of 0.3m from the true position.

9.5.6 Full game analysis

The final test of the reconstruction algorithm is how well it performs when used on data from a real game. The analysis of such a test will be considerable and is important to this thesis. For

this reason chapter 10 has been devoted to this analysis

9.6 Reconstruction conclusions

It was the aim in this chapter to develop a reconstruction algorithm that could accurately recover a significant number of invisible positions from a tracked rugby game and thus increase the effectiveness of a minimal system (research question 1). This algorithm was intended to be the final part of the software expert system (research question 3). This aim was achieved and this model can be applied to any position-tracking application given sufficient knowledge of the application-specific behaviour of the objects.

Both general and rugby-specific reconstruction methods were developed. General methods included first-order linear interpolation (second order was considered but rejected) and speed checks. The rugby specific reconstruction methods were developed to handle rucks/mauls and scrums. Lineouts were investigated as a rugby-specific method but were discounted based on the variable nature of lineout formations.

Reconstructions were applied in the following order: speed check, interpolation, scrums then rucks/mauls.

Multiple tags were investigated as an alternative way to mitigate invisibility. This approach was found in field trials to be an effective option but it may have adverse effects on meeting the specifications.

In single player trials of the Multispectral tracking system, simple interpolation was able to increase the visibility to 97% both with and without speed checking.

Ten example data sets for both rucks/mauls and scrums (each containing ten frames) were run through the reconstruction algorithm to test the identification and correction of special formations. The algorithm was able to correctly identify the formation in each data set. Positions within a ruck/maul were recovered with an average position error of 0.20m and scrums with 0.26m average (standard deviations were 0.04 and 0.06 respectively, both sufficiently low to give confidence in the results).

The reconstruction algorithm is ready to be tested in a real game situation. This testing and analysis will be performed in chapter 10.

10 Simulated tracking system performance

In which the various threads of research are implemented in a single simulation and performance is gauged.

All research to this point has been pursued in order to develop a tracking system capable of tracking players during a game of rugby to the specifications given in chapter 4 (the system required by the first research question).

In this chapter, all of the performance enhancements and measurement tools are applied in a single analysis of the game data suite (this is the expert system proposed by research question 3). The aim of this analysis was to see how the level of player visibility is affected by using the calculated optimum configuration for receivers (calculated for the second research question) and how well the reconstruction algorithm can recover the lost positions.

Explanation will be given for why this analysis was not performed in a real-game situation, and the advantages of analysing the game data suite are outlined.

The process for running the analysis on simulated data will be explained and reference made to previous sections of this thesis where required.

The results of the analysis will be discussed with special attention given to how the visibility level changes before and after reconstruction.

10.1 Simulation versus a real game

Ideally this analysis would be performed in a real game (two full teams wearing tags). With data from an actual game, the visibility level can be reported with confidence since there is no need to apply a model to find where positions are invisible.

It was not possible to use the Multispectral tracking system (see section 5.3) for testing in a real game owing to issues with the Industrial Partner relating to commercial sensitivity and insurance (see section 3.8). For these reasons this analysis was performed on the previously tracked data suite obtained from the Industrial Partner (see section 7.1). While simulations on these data are less desirable than analysis of a real game, the data are representative of player positions in a game of rugby and will allow calculations of accuracy.

10.2 Analysis process

The analysis was performed on the entire data suite of games. This equates to 10 games with approximately 30,000 frames; each frame is 2 seconds so this approximates to 17 hours of game play.

Using the player visibility model developed in chapter 7, the data in the data suite will be processed to find those players who are invisible to the tracking system. It is in this part of the analysis process that the optimum receiver placements found in chapter 8 will be used. These receiver configurations were shown to produce the least number of player invisibilities when compared with other configurations. The player visibility model will be run using the optimum 6, 8 and 10-receiver configurations.

Tag placement tests (section 6.4), indicated that low down on the leg or on the back of the shoe were the best places to mount the tracking tag for optimum performance. Originally, the model developed in chapter 7 was not designed to simulate the legs of the players and thus configuring the player visibility model with the tags 0.1-0.2m from the ground would cause significantly more invisibilities than would be expected from an actual game.

The model was modified slightly here to allow modelling of the tags low on the legs (tag height is set at 0.15m off the ground). If the line-of-sight path of the signal from the player to the receiver is below the waist of the candidate for obstruction (the waist is defined as being 0.8m off from the ground), then a check is made to see if the candidate's legs obstruct the signal. The legs are defined as having an average radius of 0.1m.

Once the visibility model has been applied, the number and type of invisibilities can be analysed. Of particular interest is the length of time that a player is invisible (referred to as contiguous invisibility) since this will affect the performance of the reconstruction algorithm (developed in chapter 9). Large contiguous invisibilities will be unrecoverable by the algorithm and will indicate that human intervention is required (see section 9.3). The level of human intervention that is required will be an important factor for gauging the performance of the system, since this will need to be minimised to make the system marketable by the Industrial Partner.

The final stage of the process is to apply the reconstruction algorithm to the data to recover a portion of the invisible data. Analysis of the number of invisible positions remaining after this reconstruction will give an indication of how effective the algorithm is and also how much time can be saved when compared with a fully manual tracking system. Furthermore, the change in the contiguous invisibility can be analysed. Remaining invisibilities will be examined and, where possible, explanations will be given for why they remain.

An investigation will be made as to the effect of applying the reconstruction algorithm multiple times.

10.3 Simulation results

After running the player visibility model on the data suite with the optimum 6, 8, and 10-receiver configurations the level of invisibility was calculated as a percentage of the number of frames. Figure 10.1 shows the level of invisibility. The field trials with the Multispectral system suggested that the invisibility level calculated by the player visibility model accurately represented what could be expected from the real system (see section 6.6.4).

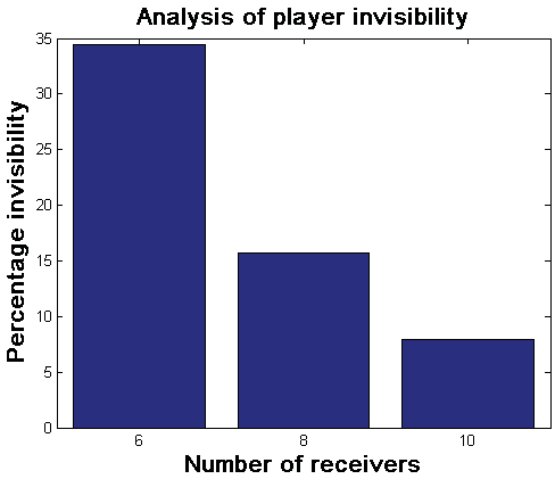

Figure 10.1: Percentage invisibility after application of the player visibility model.

From this it can be seen there are a high number of invisibilities present in the 6-receiver configuration (even with optimal placement). There is a definite improvement in visibility when more receivers are used. With 8 receivers the invisibilities drop below 17% while 10 receivers result in invisibilities slightly in excess of 5%. Increasing the number of receivers used will increase both the cost and volume of equipment and so the advantage of greater visibility must be weighed against the requirements

of the system (see chapter 3). Figure 10.2 shows the distribution of contiguous invisibilities in the data. For clarity, contiguous invisibility is the number of consecutive frames in which a player is missing; thus, being a histogram, Figure 10.2 shows the frequency of a particular contiguous invisibility length.

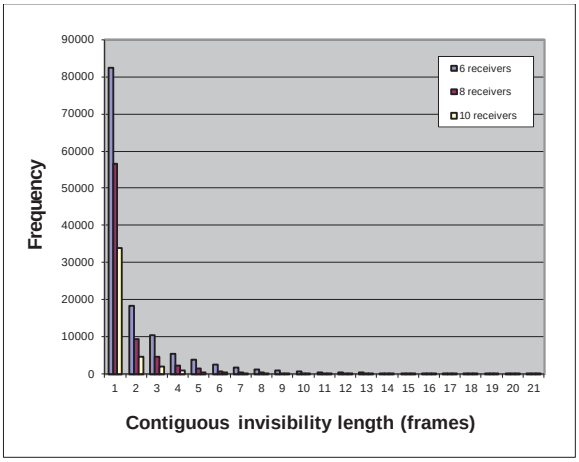

Figure 10.2: Histogram of contiguous invisibility.

It is clear from Figure 10.2 that the majority of invisibilities are over short intervals. This will be an advantage for the purposes of recovery, since short intervals can be interpolated with high accuracy. The number of receivers has little effect on the shape of the plots; in all cases over 90% of invisibilities are less than 5 frames in length and the trend is inverse exponential.

It was noticed that at times, especially near the end of each game, there was at least one player who remained stationary for long periods of time. These stationary players are artefacts in the data that represent a player being removed from the field and not replaced (i.e. not a substitution) and the Industrial Partner's analysis requires that there be 30 players on the field at all times. Reasons for such a removal would include a yellow card (10 minute 'sin bin'), a red card (sent off) or a 'blood bin' (a player temporarily sent off when they are bleeding from an injury sustained during the game). In all cases, the removed player is put in a stationary location at the periphery of the field of play. This could explain the small number of long contiguous invisibilities, especially if that player has been placed at a location calculated to be invisible. In terms of a real system, these long invisibilities could be removed or corrected manually.

After the reconstruction was performed the invisibilities for all numbers of receivers dropped significantly. Visibility after reconstruction is 93%, 97%, 99% for 6, 8 and 10-receivers respectively (calculated as total number of visible positions as a proportion of the total frames). This is shown in Figure 10.3. The invisibility level both before and after reconstruction is shown for comparison.

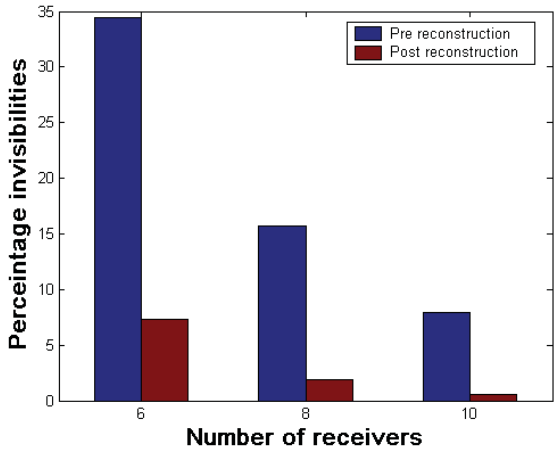

Figure 10.3: Comparison of invisibilities before and after reconstruction.

These results show that the reconstruction algorithm would be effective in increasing the performance of the tracking system and minimising the required human intervention level. These results also suggest that the recovery algorithm is more effective when higher numbers of receivers are used. This increased effectiveness is due to the greater mass of short-length contiguous invisibilities, which are easily reconstructed using interpolation.

Figure 10.4 shows how the contiguous invisibilities changed after the reconstruction was applied. In this figure, recovered positions are

considered to be valid data and thus do not appear.

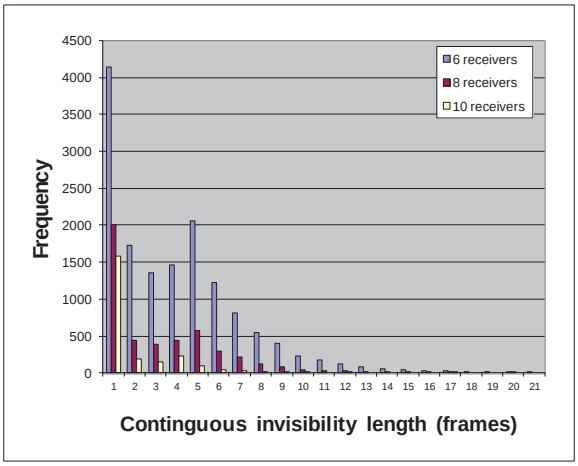

Figure 10.4: Histogram of contiguous invisibility, after applying the reconstruction algorithm.

There is a clear difference in the shape of the histogram when compared with Figure 10.2. This change in shape can be explained by the operation order of the reconstruction algorithm.

As described in chapter 9, interpolation is applied first to recover those positions that are short in duration. After this simple recovery has been performed there will be no contiguous invisibilities less than 4 frames (the threshold for a small interval invisibility defined after consideration of interpolation accuracy, see section 9.1.1) since they will all have been recovered. Next the situation-specific reconstruction methods are applied (see section 9.2). These will be able to recover the positions of some of the invisible players who are experiencing longer contiguous invisibilities, depending on the visibility of surrounding players. For example, if sufficient surrounding players become visible for a single frame, then the position of the player of interest can be recovered for that frame only. This would have the effect of splitting a large contiguous invisibility into two smaller invisibilities thus skewing the remaining contiguous invisibilities to the left of the histogram. Figure 10.5 illustrates this by showing the contiguous invisibilities after applying only the linear interpolation. Notice here that there are no invisibilities below 5 frames (as was expected). This confirms that the 1-4 frame contiguous invisibility lengths seen in Figure 10.4 are caused by corrections that are applied after interpolation.

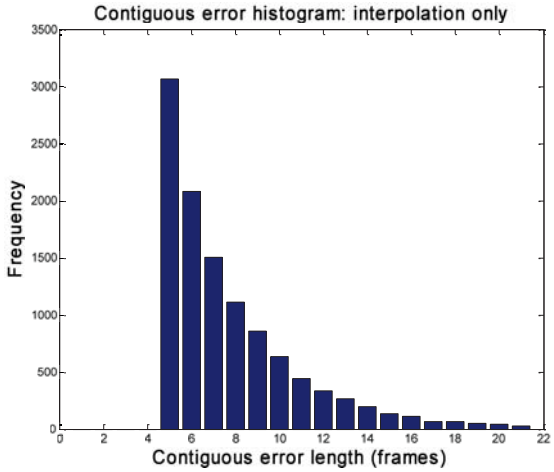

Figure 10.5: Histogram of continuous invisibility after applying linear interpolation only.

These results are after only a single pass of the reconstruction algorithm, therefore, the residual short contiguous errors are not corrected. The reconstruction algorithm can be applied *ad infinitum* to remove all but the worst-case invisibilities. However, making inferences on position from already inferred positions can lead to increasing reconstruction error and thus there is a limit to the number of times the reconstruction should be applied. The specification of this limit must be reached with reference to the requirements for accuracy in the system (section 4.6), and is left to the judgement of the operator. The reconstruction algorithm was applied again to gauge the effectiveness of a second pass. This will be discussed shortly.

Given that invisibilities greater than 4 frames are more difficult to recover, it is worth examining the contiguous invisibilities greater than this threshold. In Figure 10.4 it can be seen that invisibilities of greater length than 4 frames, for all numbers of receivers, trend towards 0 frequency as the contiguous visibility length increases. The only difference between the

numbers of receivers is the scale of the trend. This indicates that the recovery algorithm is effective at recovering positions when the contiguous invisibility length is greater than 4.

Figure 10.6 gives a summary of the type and number of corrections made in each receiver configuration.

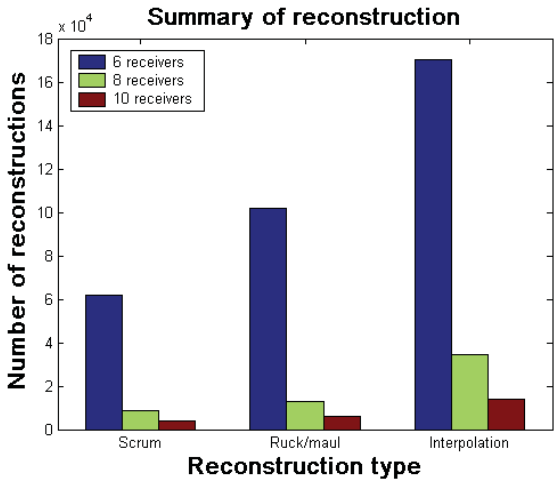

Figure 10.6: Summary of reconstruction.

Figure 10.6 shows that, for all numbers of receivers, interpolation is the most common reconstruction type employed by the algorithm followed by ruck/maul then scrum and the trend between reconstruction types is similar across different numbers of receivers. It also shows that the number of reconstructions is higher when fewer receivers are used. The number of reconstructions is significantly higher in the 6-receiver configuration compared with other configurations, which can be attributed to the significantly higher numbers of invisibilities in that configuration.

It was expected that most of the invisible positions would be recovered by interpolation, since it was seen that the majority of invisibilities were short in duration and the interpolation reconstruction is applied first. It was also expected that the number of reconstructions from interpolation would decrease with more receivers, since the number of invisible players decreased as receiver numbers increased. A similar trend can be seen for special case reconstructions for similar reasons.

The numbers of reconstructions within each category are partially dependent on the number and length of the particular situation in the invisible data. Given that the interpolation reconstruction type is applied first and that the majority of invisibilities are less than 4 frames in length, most reconstructions will be of interpolation type. The same rules influence the occurrence of ruck/maul and scrum reconstructions. During the course of a modern game of rugby it can be expected that in the order of 15 scrums will occur, involving 18 players (a number of whom can be expected to be invisible, depending on the number of receivers), and lasting up to 20 seconds. Rucks and mauls can occur upwards of 50 times in a game, contain anywhere from 4-15 players and be minutes in length (it can take some time for all of the players to untangle themselves after the formation proper). With this information, the relationship between the different reconstruction types is to be expected.

From visual inspection of the input data it was found that there were 452 scrum frames in the data set each lasting approximately 5 frames (this averages to approximately 9 scrums per game). The reconstruction algorithm was able to correctly identify initial frames (where the set piece is as standard, this identification rate is based on recovered frames marked as being corrected by scrum) from each scrum but was only able to maintain scrum correction for 3 frames. This is approximately 60% overall success at identifying scrums. However, from visual inspection, after 3 frames the scrum does not match the template and so is treated as a ruck/maul (this behaviour was expected). Because of the correct scrum identification prior to collapse the algorithms classification as a ruck/maul is correct and effective.

From visual inspection of a sample of 6000 frames (a 20% sample from the frame population, equivalent to two full games) it was

calculated that there were approximately 3000 ruck/maul frames in the dataset[14]. Each ruck could last between 6 and 20 frames each. It has been mentioned that it sometimes takes several seconds for players to disentangle themselves after the official formation is finished. The reconstruction was able to identify 2525 ruck/maul frames across the 10-game data set (this is based on frames where a ruck correction has occurred). This calculates as an 84% identification rate for rucks/mauls. The other 16% of rucks remain unidentified because the initial conditions for a ruck were not visible. This is either because not enough players in the ruck are visible to classify a ruck or the player in question was not visible immediately preceding the ruck and so could not be calculated to be inside the ruck.

Reconstruction error can be defined in this case as the radial distance between reconstructed position of an invisible player and their true position (recorded in the original data suite). Ignoring those positions that could not be recovered (and thus have no valid data in the output), the mean reconstruction error from the process were 0.58m, 0.44m, and 0.41m for 6, 8 and 10 receivers respectively (with standard deviations of 0.33m, 0.28m and 0.26m). These figures show that with reconstruction the accuracy of the system can be made to be close to that which is required (even when considering the mean plus one standard deviation).

In order to mitigate the effects of reconstruction error, it is possible to check the speeds of the players again (player speeds are calculated as part of the reconstruction process, see section 9.1.2). This would highlight potentially inaccurate reconstructions that could then be checked by a human operator. Alternatively, these can be dealt with automatically by more than one application of the reconstruction algorithm (assuming that the limit has not been reached).

[14] The time required to classify this formation across all 30000 frames was prohibitive.

Running the reconstruction a second time further reduced the level of residual invisibility in the dataset. Figure 10.7 shows the contiguous invisibility of residual invisibilities.

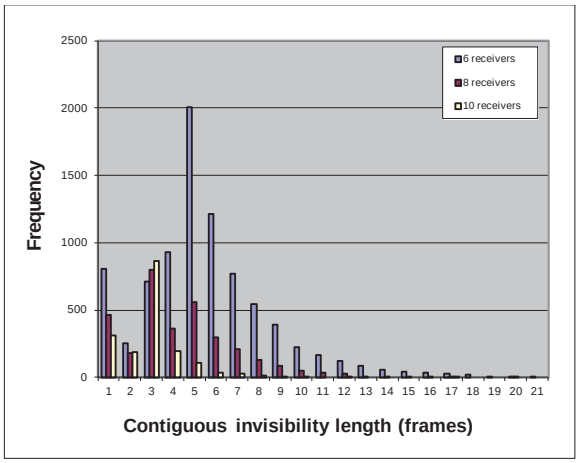

Figure 10.7: Contiguous invisibility after the second application of the reconstruction algorithm.

This figure shows a decrease in both short interval invisibilities and long interval invisibilities compared with what is shown in Figure 10.4. This reduction is especially noticeable in the 6-receiver configuration below 4 frames.

Here again it can be seen that a number of invisibilities remain below 4 frames (a result of situation specific reconstructions splitting up longer invisibilities). If the reconstruction algorithm were to be run again then these would be corrected by interpolation. This supports the supposition that applying the reconstruction multiple times will reduce the invisibility level.

Again, it is advantageous to examine how the contiguous invisibilities above 4 frames in length change with different numbers of receivers. It can be seen that all numbers of receivers have approximately the same trend but with different scale factors. Comparing Figure 10.7 with Figure 10.4 there is only a slight decrease at all lengths of contiguous invisibility. This is to be expected, since Figure 10.4 and Figure 10.7 are histograms in which values on the right hand side of the figure represent higher and

higher lengths of invisibility. Splitting up even a small number of these invisibilities would result in a large number of contiguous invisibilities shorter than 4 frames, which can then be corrected using interpolation.

Figure 10.8 compares the invisibility level after the second application of the reconstruction algorithm with previous levels of invisibility.

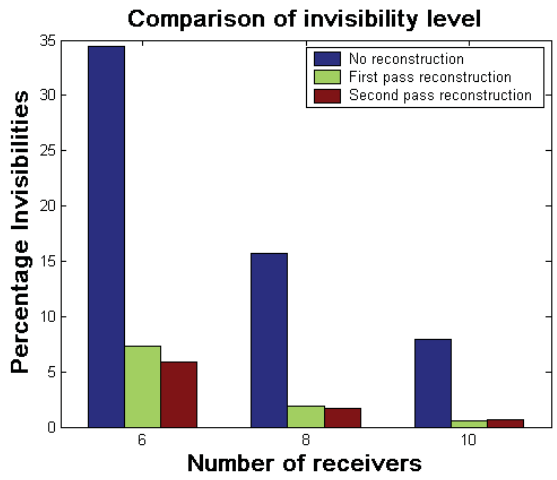

Figure 10.8: Comparison of invisibility levels.

From Figure 10.8 it is clear that there is a diminishing return from multiple passes of the reconstruction algorithm. In the case of 10-receivers there is actually a slight increase in the invisibility level (from 0.55% to 0.63%, calculated as percentage of invalid positions to total possible positions), which can be attributed to reconstruction errors in the first pass causing players to be logged as moving too fast in the second pass which then cannot be recovered.

The mean reconstruction errors after the second pass were 0.59m, 0.46m and 0.44m for the 6, 8 and 10-receiver configurations (with standard deviations of 0.29m, 0.26m and 0.24m). This is only a slight increase in error compared with the results of the first pass reconstruction, but still the figures conform to the accuracy requirement of the tracing system. It was concluded that running the reconstruction more than once would decrease the accuracy of recovered positions and that running the algorithm more than twice would decrease the accuracy beyond the specified limit for 8 and 10 receivers (as defined in section 4.6).

In general, the consequence of multiple applications of the reconstruction algorithm will be that the number of invisibilities will be decreased until only the unrecoverable player positions remain. At the same time, making inferences from already inferred positions will increase the reconstruction error.

The analysis of this simulation shows that the reconstruction algorithm can be used to increase the effectiveness of any tracking system (given that all systems output position data) and furthermore, can be used to increase the effectiveness of a minimalist system.

10.4 Tracking system inferences

In chapter 7, it was explained that the visibility model was conservative because of several limitations in the input data. This causes it to overestimate the invisibility level. Thus, it can be expected that in a real system the invisibility level will be less than calculated in Figure 10.1. Even so, the results of this simulation suggest that 8 receivers would be necessary to reduce the raw invisibility level below 20% (see section 4.6).

The effectiveness of the reconstruction algorithm means that invisibilities can be reduced below 7% even with 6 receivers, suggesting that the 6-reciever configurations may still be a viable minimalist option for use in tracking with respect to visibility; however, the inaccuracy of the 6-receiver system is slightly too high. The 8-receiver system is a better option for the minimalist system because it has a visibility better than the 6-receiver system and a more suitable accuracy. The mean accuracy plus one standard deviation is still above the maximum stated in the specifications; however, it was noticed that there was a diminishing return for both increasing visibility and measurement accuracy with

more receivers. An 8 receiver configuration represents the best balance between the requirements for measurement accuracy, equipment volume and ease of use (i.e. more receivers require more time to set up).

With eight receivers and one tag per player (including substitutions) the cost of the hardware for the minimalist system would be $40,000NZ for the receivers and $600NZ for the tags.

The effectiveness also means that the required time to track a rugby game can be significantly reduced (in accordance with the requirement for timely availability of data described in section 4.8). It has been stated previously that it requires 576 man-hours to track a rugby game manually from video (36 hours with 16 people). Given the post-recovery visibility of between approximately 93-99.5% (visibility percentage is 100% less the invisibility percentage), then it could be expected that the required time for tracking would be reduced to approximately 3-40 man-hours. These figures do not take into account that the checking process is required for both manual and automatically tracked games.

10.5 Simulation conclusions

Research results described in previous chapters were combined in one set of simulations, which were used to gauge the effectiveness of the tracking system as a whole.

Application of the research to a real game of rugby would have been preferred; however, the Multispectral tracking system was unavailable for such a trial.

The simulations were performed using the optimum configurations of 6, 8 and 10 receivers for a rugby game calculated in chapter 8. Using these configurations, the level of invisibility in the entire 10-game data suite was calculated. The reconstruction algorithm developed in chapter 9 was then applied several times to the invisible data.

The results of the simulations were analysed before reconstruction and after the first and second pass of the reconstruction algorithm. More receivers were found to produce less invisibility. The reconstruction algorithm was shown to be able to increase visibility in the data to between 93-99.5%.

The scrum detector was shown to be 60% effective at identifying scrum frames and the ruck/maul detector was 84% effective.

The mean reconstruction error was approximately 0.5m or better for all numbers of receivers after one application of the reconstruction algorithm; however, the spread suggested that only the 8 and 10-receiver systems were acceptable based on the accuracy specification. There was no improvement to the mean reconstruction error after the second application of the algorithm, however the spread did drop. Multiple applications of the reconstruction algorithm will increase both the visibility and the reconstruction error present in the data; more than three applications would likely increase reconstruction error to an unacceptable level in terms of the system specifications.

Using the Multispectral tracking system in combination with the recovery algorithm significantly reduces the required level of human intervention to track a game when compared with manual video tracking.

The simulations have shown that the reconstruction algorithm is suitable for realising a minimalist tracking system that is suitable in terms of the system requirements for timeliness and accuracy (sections 4.8 and 4.6). This answers the first research question.

This minimalist system is made possible with the use of the software tools for planning receiver placement and position reconstruction to increase visibility in the raw data (receiver placement, question two) and reconstruction of missing data (research question three).

The simulations have shown the Multispectral tracking system to be an effective way to track players during a rugby game.

11 Conclusions

In which the conclusions reached in all of the previous chapters are summarised. The achievements and original contributions are stated and the overall success of the project is gauged.

This research aimed to investigate the configuration and performance of position tracking systems. The primary application and example of this research has been the development of a minimalist system suitable for tracking rugby players which is useable by the Industrial Partner. By extension of knowledge gained, the application of the tracking system was applied, in concept, to any open-field sport and any tracking scenario with static obstructions. This research has been successful, and has produced original knowledge in several areas.

An in-depth analysis of current position-tracking techniques revealed several suitable techniques, the fundamental basis of which was usually measurement of time or angle from a reference.

A gap in knowledge was identified that suggested that the best configuration for receivers in any PLN was a regular pattern. This research challenged that notion based on the theory that complex environments would cause obstruction to critical receivers at some locations within the area of interest and thus decrease performance.

The candidate tracking system would be a distributed network of some kind and thus would require a level of time-synchronisation for correct operation. A study of synchronisation techniques found several suitable techniques.

Consultation with the Industrial Partner and research into the game of rugby union provided the specifications for the project and the constraints imposed by these specifications defined the direction of the research. These specifications were quantitative and measurable.

Several off-the-shelf position-tracking systems were identified, one of which met the project specifications. The decision was therefore made not to construct a tracking system in-house. The Sapphire Dimension system from Multispectral Solutions (Maryland, United States of America) was chosen. The parameters for further research were based on this and field trials were carried out using the Sapphire Dimension system. The research then moved to characterise the system so that its operation could be fully understood and improved upon. The research in this area extended the knowledge of the system beyond that of the manufacturer.

The system characterisation showed that human bodies act as shields to the signals transmitted by the tags which confirmed the premise that a player-tracking situation can be considered to be a complex environment and thus, obstructions caused by the tracked objects must be considered as a barrier to performance when the object density is high (as in rugby).

A visibility model was developed to predict the visibility of players to the tracking system for given positions of the players and system com-

ponents. The model was shown to be effective and successfully validated against the Multispectral tracking system. An extension was made to the model to allow it to calculate visibility in an environment containing static obstructions in an effort to generalise the concepts. The visibility model's development and testing further added to the general body of knowledge.

A genetic-algorithm based optimisation was used to find the optimal number and configuration of receivers. GA optimisations were found to be a superior way of optimising in the multimodal solution space in which this research is based. Performance of the system was shown to be dependent on coverage, invisibilities and GDOP. Several different objective functions were tested to produce a robust solution. Solutions were found for 6, 8 and 10-receiver systems with equal or better performance than those found by inspection. The optimisations, once again, implemented knowledge gained from the system characterisation. Further optimisations were performed using the static-obstruction model that confirmed the original premise that regular receiver configurations are not suitable in complex environments. The superior solutions moved receivers so that coverage zones were maximised and in such a way that the high GDOP areas that exist on the connecting lines of receivers were minimised.

Initial analysis and later experience indicated that objects would be intermittently invisible to all of the systems, resulting in missing data. A reconstruction algorithm was developed and shown to be effective in reducing the number of invisibilities. The reconstruction algorithm was also shown to be an effective way of significantly increasing the performance of a minimalist system. A small percentage of situations were too severely degraded to be recovered automatically and thus the need for human intervention cannot be removed but can be reduced by up to 90%.

A minimalist tracking-system configuration that employs only three receivers was shown to be insufficient for the purposes of rugby player tracking, based on the high level of invisibility. Results from field trials (in this case inference made from half of the field being covered by four receivers) showed that six receivers is close to the minimum required to track a game of rugby adequately (employing the reconstruction algorithm) but does not meet the requirement for accuracy (although it is close). Adding more receivers would improve the results but would also increase the volume of hardware. Eight receivers are recommended as the optimum number to satisfy the specifications of the system.

Data from real games – supplied by the Industrial Partner – were used in simulations to test the reconstruction algorithms after first using the visibility model to eliminate data from the suite. The analysis indicated that the proportion of visible players in a given set of raw data would be between 65-93%, depending on the number of receivers employed.

By applying the reconstruction algorithm a number of times, the proportion of visible players could be increased to 93-95.5%. The error involved in reconstructing the positions of invisible players had a mean of approximately 0.5m ±0.33m or better for all numbers of receivers. However, the standard deviation of the 6-receiver system meant that it did not meet the requirements for accuracy. The 8 and 10-receiver systems had acceptable mean errors and with plus one standard deviation are above the maximum acceptable error by no more than 0.22m. It was found that an 8-receiver configuration provided the best balance between the specifications of the system; partly because of a diminishing return from increasing numbers of receivers. It is recommended that 8 receivers be the basis for implementation of the minimalist system.

It was noted that the research was limited, owing to the denial of access to the equipment for a

trial in a game situation. It is suggested that any future work begins with a verification of this work on the real system.

A number of avenues for future research were identified. This thesis presents theoretical general models for tracking over wide areas with static obstructions. Testing of the concepts presented is another path for future research.

While it was found that the calf was the best place to mount the tags, little work was done regarding mounting the tags beyond suggesting that padding would be required. This work was deemed to be beyond the scope of this research; however, significant work can be done on ensuring that the tags are mounted in a way that is secure (i.e. will stay in place for the entire game, regardless of what situations the player is involved in) and comfortable to wear.

The optimisations that utilised a changing sample of player positions were limited by the computational requirements. By increasing the efficiency of the algorithms and with more computing power it may be possible to use some of these optimisations in the expert system. Similarly, the expert system can be applied to different tracking situations.

In the nature of technology, newer position tracking systems and platforms will have been produced since this thesis was published. It will be interesting to see if the results given here can be improved upon with higher technology equipment. It would be expected that newer systems will have better tag-to-receiver communications (better antennae, more efficient communication protocols) and more accurate timing (leading to more precise TDOA measurements). With newer technologies that have such characteristics it should be possible to increase performance and reduce the required volume of equipment for the minimalist system.

The modifications and developments applied to the Multispectral tracking system have made it acceptable with respect to the specified requirements from the Industrial Partner. Fur-Furthermore, the visibility model, optimisation data and the reconstruction algorithm collectively form an expert system for the tracking of rugby players and form the framework of an expert system for all position tracking applications.

12References

[1] Levanon, N., "Lowest GDOP in 2-D scenarios", *IEE Proceedings - Radar, Sonar Navigation*, 2000. 147(3): p. 149-155.

[2] Gillette, M.D. and H.F. Silverman, "A Linear Closed-Form Algorithm for Source Localization From Time-Differences of Arrival", *IEEE Signal Processing Letters*, 2008. 15: p. 1-4.

[3] Doherty, R.H., G. Hefley, and R.F. Linfield, "Timing Potentials of Loran-C", *Proceedings of the IRE*, 1961. 49(11): p. 1659-1673.

[4] Wu, N., et al. "TDMA Network Time Synchronization and Navigation Using DGPS", *IEEE International Frequency Control Symposium and Exposition* p. 287-290, (2006)

[5] Sugimoto, Y., et al., "Development of GPS positioning system `PRESTAR'", *Instrumentation and Measurement, IEEE Transactions on*, 1989. 38(2): p. 644-647.

[6] Zhu, J., "Calculation of geometric dilution of precision", *IEEE Transactions on Aerospace and Electronic Systems*, 1992. 28(3): p. 893-895.

[7] Yarlagadda, R., et al., "GPS GDOP metric", *IEE Proceedings - Radar, Sonar Navigation*, 2000. 147(5): p. 259-264.

[8] Lee, H.B., "A Novel Procedure for Assessing the Accuracy of Hyperbolic Multilateration Systems", *Aerospace and Electronic Systems, IEEE Transactions on*, 1975. AES-11(1): p. 2-15.

[9] Marchand, N., "Error Distributions of Best Estimate of Position from Multiple Time Difference Hyperbolic Networks", *Aerospace and Navigational Electronics, IEEE Transactions on*, 1964. ANE-11(2): p. 96-100.

[10] Bronk, K. and J. Stefanski. "Bad Geometry Influence on Positioning Accuracy in Wireless Networks", *EUROCON, 2007. The International Conference on "Computer as a Tool"* p. 1131-1135, (2007)

[11] Zhong, E.-J. and T.-Z. Huang. "Geometric Dilution of Precision in Navigation Computation", *2006 International Conference on Machine Learning and Cybernetics* p. 4116-4119, (2006)

[12] Bo, X. and B. Shao. "Satellite selection algorithm for combined GPS-Galileo navigation receiver", *ICARA 2009. 4th International Conference on Autonomous Robots and Agents.* p. 149-154, (2009)

[13] Chaffee, J. and J. Abel. "GDOP and the Cramer-Rao bound", *IEEE Position Location and Navigation Symposium* p. 663-668, (1994)

[14] Lahanas, M., "Measurements", http://www.mlahanas.de/Greeks/Measurements.htm, visited on 23 February 2009.

[15] Name Project, "Sine Rule", http://www.acts.tinet.ie/sinerule_673.html, visited on 23 February 2009.

[16] Charlton, D., M. Polkinghorne, and P. Robinson. "Development of an enhanced video location and tracking system for an XY servo table", *Fifth International Conference on Factory 2000 - The Technology Exploitation Process, (Conf. Publ. No. 435)* p. 69-75, (1997)

[17] Mikael Kais, et al. "Towards Outdoor Localization Using GIS, Vision System and Stochastic Error Propagation", *2nd International Conference on Autonomous Robots and Agents Palmerston North, New Zealand*, p. 198-205, (2004)

[18] Cevher, V., et al., "Target Tracking Using a Joint Acoustic Video System", *IEEE Transactions on Multimedia*, 2007. 9(4): p. 715-727.

[19] Chellappa, R., G. Qian, and Q. Zheng. "Vehicle detection and tracking using acoustic and video sensors", *IEEE International Conference on Acoustics, Speech, and Signal Processing, 2004. Proceedings.* p. iii-793-6 vol.3, (2004)

[20] Yang, C.-S., et al. "PTZ camera based position tracking in IP-surveillance system", *3rd International Conference on Sensing Technology* p. 142-146, (2008)

[21] Kakiuchi, H., et al. "An algorithm to determine neighbor nodes for Automatic Human Tracking System", *IEEE International Conference on Electro/Information Technology* p. 96-102, (2009)

[22] Li, B., et al. "Interval least-squares filtering with applications to robust video target tracking", *IEEE International Conference on Acoustics, Speech and Signal Processing* p. 3397-3400, (2008)

[23] Li, L., et al., "An Efficient Sequential Approach to Tracking Multiple Objects Through Crowds for Real-Time Intelligent CCTV Systems", *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 2008. 38(5): p. 1254-1269.

[24] Crowley, J.L. and K. Schwerdt. "Robust tracking and compression for video communication", *Proceedings. International Workshop on Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems* p. 2-9, (1999)

[25] Yu, M., S.M. Naqvi, and J. Chambers. "Fall detection in the elderly by head tracking", *IEEE/SP 15th Workshop on Statistical Signal Processing* p. 357-360, (2009)

[26] Snidaro, L., et al., "Quality-Based Fusion of Multiple Video Sensors for Video Surveillance", *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 2007. 37(4): p. 1044-1051.

[27] Snidaro, L., I. Visentini, and G.L. Foresti. "Multi-sensor Multi-cue Fusion for Object Detection in Video Surveillance", *Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance* p. 364-369, (2009)

[28] Akazawa, Y., Y. Okada, and K. Niijima. "Real-time video based motion capture system based on color and edge distributions", *IEEE International Conference on Multimedia and Expo, 2002. ICME '02. Proceedings.* p. 333-336 vol.2, (2002)

[29] Wang, H., L. Zeng, and C. Yin, *A video tracking system for measuring the position and body deformation of a swimming fish*, in *Review of Scientific Instruments*. 2002. p. 4381-4384.

[30] Brooks Zurn, J., J. Xianhua, and Y. Motai, "Video-Based Tracking and Incremental Learning Applied to Rodent Behavioral Activity Under Near-Infrared Illumination", *IEEE Transactions on Instrumentation and Measurement*, 2007. 56(6): p. 2804-2813.

[31] Schonfeld, D. and D. Lelescu. "VORTEX: Video retrieval and tracking from compressed multimedia databases-visual search engine", *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences, 1999. HICSS-32.* p. 10 pp., (1999)

[32] Mokhtarian, F. and F. Mohanna. "Content-based video database retrieval through robust corner tracking", *IEEE Workshop on Multimedia Signal Processing* p. 224-228, (2002)

[33] Jiang, H., et al. "A Fast and Effective Text Tracking in Compressed Video", *Tenth IEEE International Symposium on Multimedia.* p. 136-141, (2008)

[34] Choe, T.E., et al. "Image transformation for object tracking in high-resolution video", *19th International Conference on Pattern Recognition* p. 1-4, (2008)

[35] Li, Z., J. Chen, and N.N. Schraudolph. "An improved mean-shift tracker with kernel prediction and scale optimisation targeting for low-frame-rate video tracking", *19th International Conference on Pattern Recognition.* p. 1-4, (2008)

[36] Bouttefroy, P.L.M., et al. "Vehicle Tracking Using Projective Particle Filter", *Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance* p. 7-12, (2009)

[37] Jonathan R. Seal, Donald G. Bailey, and Gourab Sen Gupta. "Depth Perception with a Single Camera", *1st International Conference on Sensing Technology Palmerston North, New Zealand*, p. 96-101, (2005)

[38] Orchard, M.T. "On modeling location uncertainty in images", *2001 International Conference on Image Processing* p. 13 vol.1, (2001)

[39] Bikkannavar, S. and D. Redding, "The end of the blur", *Spectrum, IEEE*. 47(3): p. 46-53.

[40] H Jiang, et al. "Colour Vision and Robot/Ball Identification for a Large Field Soccer Robot System", *2nd International Conference on Autonomous Robots and Agents Palmerston North, New Zealand*, p. 323-327, (2004)

[41] Zhang, Y., H. Lu, and C. Xu. "Collaborate ball and player trajectory extraction in broadcast soccer video", *19th International Conference on Pattern Recognition* p. 1-4, (2008)

[42] Utsumi, O., et al. "An object detection method for describing soccer games from video", *IEEE International Conference on Multimedia and Expo, 2002. ICME '02. Proceedings* p. 45-48 vol.1, (2002)

[43] Mya, A.P.P. and M.M. Sein. "Tracking of the motion path of a person from video streams for the overlapping case", *IEEE Instrumentation and Measurement Technology Conference* p. 904-908, (2009)

[44] Zhu, G., et al. "Automatic Multi-Player Detection and Tracking in Broadcast Sports Video using Support Vector Machine and Particle Filter", *IEEE International Conference on Multimedia and Expo* p. 1629-1632, (2006)

[45] Intille, S.S. and A.F. Bobick. "Closed-world tracking", *Fifth International Conference on Computer Vision* p. 672-678, (1995)

[46] Han, J., et al. "An automatic analyzer for sports video databases using visual cues and real-world modeling", *2006 Digest of Technical Papers. International Conference on Consumer Electronics* p. 477-478, (2006)

[47] Saito, H., N. Inamoto, and S. Iwase. "Sports scene analysis and visualization from multiple-view video", *IEEE International Conference on Multimedia and Expo* p. 1395-1398 Vol.2, (2004)

[48] Li, H., et al., "Automatic Detection and Analysis of Player Action in Moving Background Sports Video Sequences", *IEEE Transactions on Circuits and Systems for Video Technology*, 2009. PP(99): p. 1.

[49] Kasuya, N., et al. "Automatic player's view generation of real soccer scenes based on trajectory tracking", *3DTV Conference: The True Vision - Capture, Transmission and Display of 3D Video* p. 1-4, (2009)

[50] Fontana, R.J., E. Richley, and J. Barney. "Commercialization of an ultra wideband precision asset location system", *IEEE Conference on Ultra Wideband Systems and Technologies Reston, V A.*, p. (2003)

[51] Runtz, K.J. and F. Yang. "Pulse shaping and signal modulation techniques to improve the multipath and noise performance of narrowband precision RF ranging", *Canadian Conference on Electrical and Computer Engineering, 2002. IEEE CCECE 2002* p. 358-363 vol.1, (2002)

[52] Grosicki, E. and K. Abed-Meraim. "A new trilateration method to mitigate the impact of some non-line-of-sight errors in TOA measurements for mobile localization", *IEEE International Conference on Acoustics, Speech, and Signal Processing.* p. iv/1045-iv/1048 Vol. 4, (2005)

[53] Li, J., et al. "NLOS error mitigation and mobile tracking", *7th International Conference on Signal Processing, 2004. Proceedings. ICSP '04.* p. 2453-2456 vol.3, (2004)

[54] Mathias, A., M. Leonardi, and G. Galati. "An efficient multilateration algorithm", *Tyrrhenian International Workshop on Digital Communications - Enhanced Surveillance of Aircraft and Vehicles.* p. 1-6, (2008)

[55] Fontana, R.J. and S.J. Gunderson. "Ultra-wideband precision asset location system", *2002 IEEE Conference on Ultra Wideband Systems and Technologies, 2002. Digest of Papers.* p. 147-150, (2002)

[56] D'Antona, G. and R. Ottoboni. "Time of flight measuring system for lakes acoustic tomography", *Instrumentation and Measurement Technology Conference, 2002. IMTC/2002. Proceedings of the 19th IEEE* p. 1671-1673 vol.2, (2002)

[57] Buzinski, M., A. Levine, and W.H. Stevenson. "Performance characteristics of range sensors utilizing optical triangulation", *Proceedings of the IEEE 1992 National Aerospace and Electronics Conference.* p. 1230-1236 vol.3, (1992)

[58] Axelsson, R. "Airborne Laser Surveying - Systems and Applications", *Conference on Lasers and Electro-Optics Europe, 1998. 1998 CLEO/Europe.* p. 387, (1998)

[59] Chang, C.C., C.Y. Chang, and Y.T. Cheng. "Distance measurement technology development at remotely teleoperated robotic manipulator system for underwater constructions", *2004 International Symposium on Underwater Technology, 2004. UT '04.* p. 333-338, (2004)

[60] Chadwick, J.B. and J.L. Bricker. "A vehicle location system (VLS) solution approach", *Position Location and Navigation Symposium.* p. 127-132, (1990)

[61] Manolakis, D.E., "Efficient solution and performance analysis of 3-D position estimation by trilateration", *IEEE Transactions on Aerospace and Electronic Systems*, 1996. 32(4): p. 1239-1248.

[62] Favre-Bulle, B., J. Prenninger, and C. Eitzinger. "Efficient tracking of 3D-robot positions by dynamic triangulation", *IEEE Instrumentation and Measurement Technology Conference, 1998. IMTC/98. Conference Proceedings.* p. 446-449 vol.1, (1998)

[63] Cai, S., et al. "Modified Improved Alternating Combination Trilateration Algorithm with a Variable m", *International Multisymposiums on Computer and Computational Sciences.* p. 98-101, (2008)

[64] Yu, Y.-B. and J.-Y. Gan. "Self-localization using alternating combination trilateration for sensor nodes", *International Conference on Machine Learning and Cybernetics.* p. 85-90, (2009)

[65] Yang, Z., Y. Liu, and X.-Y. Li. "Beyond Trilateration: On the Localizability of Wireless Ad-Hoc Networks", *IEEE INFOCOM 2009* p. 2392-2400, (2009)

[66] Tian, F., et al. "Robust Localization Based on Adjustment of Trilateration Network for Wireless Sensor Networks", *4th International Conference on Wireless Communications, Networking and Mobile Computing.* p. 1-4, (2008)

[67] Sun, Y., J. Xiao, and F. Cabrera-Mora. "Robot localization and energy-efficient wireless communications by multiple antennas", *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*
Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on VO - p. 377-381, (2009)

[68] Shin, D.-H. and T.-K. Sung, "Comparisons of error characteristics between TOA and TDOA positioning", *Aerospace and Electronic Systems, IEEE Transactions on*, 2002. 38(1): p. 307-311.

[69] Chaffee, J. and J. Abel, "On the exact solutions of pseudorange equations", *IEEE Transactions on Aerospace and Electronic Systems*, 1994. 30(4): p. 1021-1030.

[70] Geyer, M. and A. Daskalakis. "Solving passive multilateration equations using Bancroft's algorithm", *17th DASC. The AIAA/IEEE/SAE Digital Avionics Systems Conference Proceedings* p. F41/1-F41/8 vol.2, (1998)

[71] Duraiswami, R., D. Zotkin, and L. Davis, "Exact solutions for the problem of source location from measured time differences of arrival." *The Journal of the Acoustical Society of America*, 1999. 106(4): p. 2277.

[72] Yang, Y.E., J. Baldwin, and A. Smith. "Multilateration tracking and synchronization over wide areas", *Proceedings of the IEEE Radar Conference.* p. 419-424, (2002)

[73] Chan, Y.T. and K.C. Ho, "A simple and efficient estimator for hyperbolic location", *IEEE Transactions on Signal Processing*, 1994. 42(8): p. 1905-1915.

[74] Ho, K.C. and W. Xu, "An accurate algebraic solution for moving source location using TDOA and FDOA measurements", *IEEE Transactions on Signal Processing*, 2004. 52(9): p. 2453-2463.

[75] Chaoshu, J., L. Changzhong, and W. Xuegang. "GPS synchronized wide area multilateration system", *International Conference on Communications, Circuits and Systems.* p. 457-459, (2009)

[76] Braasch, M.S. and A.J. van Dierendonck, "GPS receiver architectures and measurements", *Proceedings of the IEEE*. 87(1): p. 48-64.

[77] Zhao, Y., et al. "A Hybrid Location Identification Method in Wireless Ad Hoc/Sensor Networks", *Tenth International Conference on Mobile Data Management: Systems, Services and Middleware* p. 465-473, (2009)

[78] Cheng, X., et al. "TPS: a time-based positioning scheme for outdoor wireless sensor networks", *Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies* p. 2685-2696 vol.4, (2004)

[79] Guizzo, E., "Bugged Balls for Tough Calls", *IEEE Spectrum*, 42(6), p. 12, (2005)

[80] Wood, M.L. "Multilateration system development history and performance at Dallas/Ft. Worth Airport", *Proceedings. DASC. The 19th Digital Avionics Systems Conferences.* p. 2E1/1-2E1/8 vol.1, (2000)

[81] Leonardi, M., A. Mathias, and G. Galati. "Closed form localization algorithms for mode S wide area multilateration", *EuRAD 2009. European Radar Conference.* p. 73-76, (2009)

[82] Wang, S.S., M. Green, and M. Malkawa. "E-911 location standards and location commercial services", *Emerging Technologies Symposium: Broadband, Wireless Internet Access, 2000 IEEE* p. 5, (2000)

[83] Spirito, M.A. "Mobile station location with heterogeneous data", *IEEE 52nd Vehicular Technology Conference,* p. 1583-1589 vol.4, (2000)

[84] Wylie-Green, M.P. and P. Wang. "GSM mobile positioning simulator", *Emerging Technologies Symposium: Broadband, Wireless Internet Access, 2000 IEEE* p. 5 pp., (2000)

[85] Wylie-Green, M.P. and S.S. Wang. "Observed time difference (OTD) estimation for mobile positioning in IS-136 in the presence of BTS clock drift", *IEEE VTS 54th Vehicular Technology Conference, 2001.* p. 2677-2681 vol.4, (2001)

[86] Laitinen, H., J. Lahteenmaki, and T. Nordstrom. "Database correlation method for GSM location", *IEEE VTS 53rd Vehicular Technology Conference, 2001.* p. 2504-2508 vol.4, (2001)

[87] Zhao, Y., "Standardization of mobile phone positioning for 3G systems", *IEEECommunications Magazine*, 2002. 40(7): p. 108-116.

[88] Digregorio, B.E., "Messenger Arrives at Mercury", *IEEE Spectrum*, 44(12), p. 11-12, (2007)

[89] Andrews, G.R., "Theory of the Amateur Radio Doppler", http://members.aol.com/BmgEngInc/Doppler.html, visited on 4 January 2008.

[90] Kirchner, N. and T. Furukawa. "Infrared Localisation for Indoor UAVs", *First International Conference on Sensing Technology Palmerston North, New Zealand*, p. 60-65, (2005)

[91] Niculescu, D. and B. Nath. "Ad hoc positioning system (APS) using AOA", *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies.* p. 1734-1743 vol.3, (2003)

[92] Curry, M., et al. "Indoor angle of arrival using wide-band frequency diversity with experimental results and EM propagation modeling", *2000 IEEE-APS Conference on Antennas and Propagation for Wireless Communications* p. 65-68, (2000)

[93] Wang, J., et al. "The impact of AOA energy distribution on the spatial fading correlation and SER performance of a circular antenna array", *IFIP International Conference on Wireless and Optical Communications Networks* p. 4, (2006)

[94] Lo, T.K.Y., H. Leung, and J. Litva, "Artificial neural network for AOA estimation in a multipath environment over the sea", *Oceanic Engineering, IEEE Journal of*, 1994. 19(4): p. 555-562.

[95] Dai, T.-Y. and T.-S. Lee. "Source localization in a multipath environment via beamspace cumulant-based neural processing", *International Conference on Acoustics, Speech, and Signal Processing, 1995. ICASSP-95* p. 3611-3614 vol.5, (1995)

[96] Cong, L. and W. Zhuang, "Hybrid TDOA/AOA mobile user location for wideband CDMA cellular systems", *Wireless Communications, IEEE Transactions on*, 2002. 1(3): p. 439-447.

[97] Li, W. and P. Liu. "3D AOA/TDOA emitter location by integrated passive radar/GPS/INS systems", *Proceedings of 2005 IEEE International Workshop on VLSI Design and Video Technology* p. 121-124, (2005)

[98] Kleine-Ostmann, T. and A.E. Bell, "A data fusion architecture for enhanced position estimation in wireless networks", *Communications Letters, IEEE*, 2001. 5(8): p. 343-345.

[99] Yang, C., Y. Huang, and X. Zhu. "Hybrid TDOA/AOA method for indoor positioning systems", *Location Technologies, 2007. The Institution of Engineering and Technology Seminar on* p. 1-5, (2007)

[100] Sikka, P., P. Corke, and L. Overs. "Wireless sensor devices for animal tracking and control", *29th Annual IEEE International Conference on Local Computer Networks.* p. 446-454, (2004)

[101] Ross, P.E., "Managing care through the air", *IEEE Spectrum*, 41(12), p. 14-19, (2004)

[102] Vermaak, J., S. Maskell, and M. Briers. "A unifying framework for multi-target tracking and existence", *8th International Conference on Information Fusion* p. 9 pp., (2005)

[103] Musicki, D., R. Evans, and S. Stankovic, "Integrated probabilistic data association", *Automatic Control, IEEE Transactions on*. 39(6): p. 1237-1241.

[104] Musicki, D. and R. Evans, "Joint integrated probabilistic data association: JIPDA", *Aerospace and Electronic Systems, IEEE Transactions on*. 40(3): p. 1093-1099.

[105] Craig, S., et al. "Synchronization strategies for GSM/EDGE networks", *Vehicular Technology Conference, 2001. VTC 2001 Spring. IEEE VTS 53rd*

Vehicular Technology Conference, 2001. VTC 2001 Spring. IEEE VTS 53rd VO - 4 p. 2670-2674 vol.4, (2001)

[106] Liu, Q., et al. "Time Synchronization Performance Analysis and Simulation of a kind of wireless TDMA Network", *International Frequency Control Symposium and Exposition, 2006 IEEE*

International Frequency Control Symposium and Exposition, 2006 IEEE VO - p. 299-303,

[107] Thomas, G. and J. Avula. "Signal Processing for a UWB Sensor Network for Position Location", *Global signal processing conference* p. (2003)

[108] Herzel, F. and B. Razavi. "Oscillator jitter due to supply and substrate noise", *Proceedings of the IEEE 1998 Custom Integrated Circuits Conference* p. 489-492, (1998)

[109] Shinagawa, M., Y. Akazawa, and T. Wakimoto, "Jitter analysis of high-speed sampling systems", *IEEE Journal of Solid-State Circuits*, 1990. 25(1): p. 220-224.

[110] Adams, C.A., J.A. Kusters, and C.M. Sousa. "Improved quality factor for quartz oscillators", *Proceedings of the 1996 IEEE International Frequency Control Symposium* p. 699-705, (1996)

[111] Zamek, I. and S. Zamek. "Crystal oscillators jitter measurements and its estimation of phase noise", *Proceedings of the 2003 IEEE International Frequency Control Sympposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum* p. 547-555, (2003)

[112] Lisowiec, A., A. Czarnecki, and Z. Rau. "A method for quartz oscillator synchronization by GPS signal", *European Frequency and Time Forum, 1996. EFTF 96., Tenth (IEE Conf. Publ. 418)* p. 96-99, (1996)

[113] Kim, S., et al. "An Improved Timing Synchronization Scheme for BOC Signals", *IEEE International Symposium on Signal Processing and Information Technology* p. 710-714, (2007)

[114] Gershenfeld, N., R. Krikorian, and D. Cohen, "The Internet of Things", *Scientific America*, 291(4), p. 46-51, (2004)

[115] Perillo, M.A. and W.B. Heinzelman, "Wireless Sensor Network Protocols", http://www.ece.rochester.edu/courses/ECE586/readings/perillo.pdf, visited on 23 February 2009.

[116] Elson, J., L. Girod, and D. Estrin, "Fine-Grained Network Time Synchronization using Reference Broadcasts", http://lecs.cs.ucla.edu/Publications/papers/broadcast-osdi.pdf, visited on 23 February 2009.

[117] Maróti, M., et al., "The Flooding Time Synchronization Protocol", http://portal.acm.org/citation.cfm?id=1031501, visited on 23 February 2009.

[118] Palchaudhuri, S., A.K. Saha, and D.B. Johnsin. "Adaptive clock synchronization in sensor networks", *Information Processing in Sensor Networks, 2004. IPSN 2004. Third International Symposium on* p. 340-348, (2004)

[119] Trimble, "How GPS works", http://www.trimble.com/gps/howgps.shtml, visited on 23 February 2009.

[120] Geppert, L., "Move Over, Quartz", *IEEE Spectrum*, 42(5), p. 10-11, (2005)

[121] Jones, W.D., "Chip-scale Atomic Clock", IEEE Spectrum, visited on 19 June 2012.

[122] Geier, G.J., et al. "Prediction of the time accuracy and integrity of GPS timing", *Frequency Control Symposium, 1995. 49th., Proceedings of the 1995 IEEE International Frequency Control Symposium, 1995. 49th., Proceedings of the 1995 IEEE International VO* - p. 266-274, (31)

[123] Yam, P., "Everyday Einstein", *Scientific American*, 291(3), p. 34-39, (2004)

[124] Martin, K.E. "Precise timing in electric power systems", *Proceedings of the 1993 IEEE International Frequency Control Symposium, 1993. 47th.* p. 15-22, (1993)

[125] Pflieger, K., P.K. Enge, and K.A. Clements. "Improving power network state estimation using GPS time transfer", *Position Location and Navigation Symposium, 1992. Record. '500 Years After Columbus - Navigation Challenges of Tomorrow'. IEEE PLANS '92., IEEE* p. 188-193, (1992)

[126] Dierks, A., I. Worthington, and P. Olivier. "GPS synchronised end-to-end tests of transmission line teleprotection schemes in the ESKOM network", *Developments in Power System Protection, Sixth International Conference on (Conf. Publ. No. 434)* p. 355-360, (1997)

[127] Houlei, G., J. Shifang, and H. Jiali. "Development of GPS synchronized digital current differential protection", *Proceedings Power System Technology, 1998.* p. 1177-1182 vol.2, (1998)

[128] Lewandowski, W., G. Petit, and C. Thomas, "Precision and accuracy of GPS time transfer", *IEEE Transactions on Instrumentation and Measurement*, 1993. 42(2): p. 474-479.

[129] Pollastri, F., "Network News: Network Time Protocol", http://toi.iriti.cnr.it/en/ntp.html, visited on 23 February 2009.

[130] Jun, S.-M., et al. "A time synchronization method for NTP", *RTCSA '99. Sixth International Conference on Real-Time Computing Systems and Applications* p. 466-473, (1999)

[131] Mills, D.L., "Internet time synchronization: the network time protocol", *IEEE Transactions on Communications*, 1991. 39(10): p. 1482-1493.

[132] Mills, D.L., "Improved algorithms for synchronizing computer network clocks", *IEEE/ACM Transactions on Networking*, 1995. 3(3): p. 245-254.

[133] Butner, S.E. and S. Vahey. "Nanosecond-scale event synchronization over local-area networks", *Proceedings. LCN 2002. 27th Annual IEEE Conference on Local Computer Networks, 2002.* p. 261-269, (2002)

[134] Gurewitz, O., I. Cidon, and M. Sidi. "Network time synchronization using clock offset optimization", *Proceedings. 11th IEEE International Conference on Network Protocols, 2003.* p. 212-221, (2003)

[135] Nakashima, T. and S. Ihara. "An experimental evaluation of the total cost of NTP topology", *Proceedings. 15th International Conference on Information Networking, 2001.* p. 240-245, (2001)

[136] Yang, L. and G.B. Giannakis. "Blind UWB timing with a dirty template", *Proceedings. (ICASSP '04). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2004.* p. iv-509-iv-512 vol.4, (2004)

[137] Zhang, H. and D.L. Goeckel. "Generalized transmitted-reference UWB systems", *2003 IEEE Conference on Ultra Wideband Systems and Technologies* p. 147-151, (2003)

[138] Tian, Z., L. Wu, and S.A. Zekavat. "Blind vs. training-based UWB timing acquisition with effective multipath capture", *Conference Record of the Thirty-Seventh Asilomar Conference on Signals, Systems and Computers, 2003.* p. 1771-1775, (2003)

[139] Semiconductor, N., "nRF24AP1 Ultra-low power 2.4GHz transceiver", http://www.nordicsemi.com/index.cfm?obj=product&act=display&pro=88, visited on 23 February 2009.

[140] Paric, "Description of P3", http://www.paric.co.nz/description.html, visited on 23 February 2009.

[141] Multispectral Solutions Inc, "Homepage", www.multispectral.com, visited on 23 February 2009.

[142] Multispectral Solutions Inc, "Sapphire Dart Datasheet", http://www.multispectral.com/pdf/Sapphire_DART.pdf, visited on 23 February 2009.

[143] Prosound, "Proel Mic and Speaker Stands", http://www.prosound.co.nz/index.htm?stands/proel-stands-page1.htm~mainFrame, visited on 23 February 2009.

[144] Bhole, Y., "Transmitter design for UWB communications", http://uwb.rediffblogs.com/, visited on 23 February 2009.

[145] Imada, S. and T. Ohtsuki. "Pre-RAKE diversity combining for UWB systems in IEEE 802.15 UWB multipath channel", *Conference on Ultrawideband Systems and Technologies* p. 236-240, (2004)

[146] IEEE, "IEEE 802.15 WPAN Low Rate Alternative PHY Task Group 4a (TG4a)", http://www.ieee802.org/15/pub/TG4a.html, visited on 23 February 2009.

[147] Commission of the European Comunities, *Commision decision of 21/11/2007 on allowing the use of the radio spectrum for equipment using ultra-wideband technology in a harmoised manner in the Community*. 2007: Brussels.

[148] Zasowski, T., et al., "UWB signal propagation at the human head", *IEEE Transactions on Microwave Theory and Techniques,*. 54(4): p. 1836-1845.

[149] See, T.S.P., Z.N. Chen, and X.M. Qing. "Proximity effect of UWB antenna on human body", *Asia PacificMicrowave Conference, 2009. APMC 2009* p. 2192-2195, (2009)

[150] See, T.S.P. and Z.N. Chen, "Experimental Characterization of UWB Antennas for On-Body Communications", *IEEE Transactions on Antennas and Propagation*. 57(4): p. 866-874.

[151] Whitley, D., "A genetic algorithm tutorial", *Statistics and Computing*, 1994. 4(2): p. 65-85.

[152] Murata, T. and H. Ishibuchi. "Performance evaluation of genetic algorithms for flowshop scheduling problems", *Proceedings of the First IEEE Conference on Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence* p. 812-817 vol.2, (1994)

[153] Chaiyaratana, N. and A.M.S. Zalzala. "Recent developments in evolutionary and genetic algorithms: theory and applications", *Second International Conference On Genetic Algorithms in Engineering Systems: Innovations and Applications* p. 270-277, (1997)

[154] MathWorks, T., "Genetic Algorithm and Direct Search Toolbox 2.0.1", http://www.mathworks.com/products/gads/, visited on 23 February 2009.

[155] Weisstein, E.W., "Nelder-Mead Method", MathWorld, http://mathworld.wolfram.com/Nelder-MeadMethod.html, visited on 23 February 2009.

[156] Hooke, R. and T.A. Jeeves, "'Direct search' solution of numerical and statistical problems", *Journal of the Association for Computing Machinery*, 1961. 8(2): p. 212-229.

[157] International Rugby Board, "Law 16: Ruck", http://www.irb.com/mm/Document/LawsRegs/0/070110LGLAW16red_712.pdf, visited on 23 February 2009.

[158] Wikipedia, "Scrum (rugby union)", http://en.wikipedia.org/wiki/Scrum_(rugby), visited on 23 February 2009.

[159] Weisstein, E.W., "Circle-Circle Intersection", MathWorld - A Wolfram Web Resource., http://mathworld.wolfram.com/Circle-CircleIntersection.html, visited on 23 February 2009.

[160] Weisstein, E.W., "Radical Line", MathWorld - A Wolfram Web Resourse, http://mathworld.wolfram.com/RadicalLine.html, visited on 23 February 2009.

[161] Bucher, R., "Exact Solution for Three Dimensional Hyperbolic Positioning Algorithm and Synthesizable VHDL Model for Hardware Implementation", http://ralph.bucher.home.att.net/project.html, visited on 23 February 2009.

13Appendix A: Equation Derivations

This appendix contains the full derivations of the positioning equations used in this thesis that are not the sole work of the author.

13.1 Geometric Dilution of Precision Equations

The following set of equations describes the process for calculating GDOP, for explanation of the derivation see the work of Levanon[1]. The equations are given for a three-dimensional space even though only two dimensions were used in this thesis. In a two-dimensional situation, z data can be ignored.

Let:

$$\boldsymbol{R_i = \sqrt{(x_i - x)^2 + (y_i - y)^2 + (z_i - z)^2}} \quad \textbf{(13.1)}$$

Where:

- x_i, y_i, and z_i are the coordinates of a receiver
- x, y and z are the coordinates of the object (or area) of interest.

The unit vector between object and receiver i is therefore:

$$\left(\frac{(x_i - x)}{R_i}, \frac{(y_i - y)}{R_i}, \frac{(z_i - z)}{R_i}\right) \quad \textbf{(13.2)}$$

A matrix can be constructed for every receiver in the network in the following manner:

$$\boldsymbol{A} = \begin{bmatrix} \frac{(x_1 - x)}{R_1} & \frac{(y_1 - y)}{R_1} & \frac{(z_1 - z)}{R_1} \\ \frac{(x_2 - x)}{R_2} & \frac{(y_2 - y)}{R_2} & \frac{(z_2 - z)}{R_2} \\ \frac{(x_3 - x)}{R_3} & \frac{(y_3 - y)}{R_3} & \frac{(z_3 - z)}{R_3} \\ \vdots & \vdots & \vdots \\ \frac{(x_N - x)}{R_N} & \frac{(y_N - y)}{R_N} & \frac{(z_N - z)}{R_N} \end{bmatrix} \quad \textbf{(13.3)}$$

Where:

- A is the matrix of unit vectors representing the current calculation.
- N is the maximum number of receivers with visibility to the object.

Calculate the covariance of A.

$$\boldsymbol{Q} = (\boldsymbol{A}.\boldsymbol{A}^{\boldsymbol{T}})^{-\mathbf{1}} = \begin{bmatrix} \boldsymbol{d}_{\boldsymbol{x}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{xy}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{xz}}^{\mathbf{2}} \\ \boldsymbol{d}_{\boldsymbol{xy}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{y}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{yz}}^{\mathbf{2}} \\ \boldsymbol{d}_{\boldsymbol{xz}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{yz}}^{\mathbf{2}} & \boldsymbol{d}_{\boldsymbol{z}}^{\mathbf{2}} \end{bmatrix} \quad \textbf{(13.4)}$$

Where:

- Q is the covariance matrix
- d_N^2 is the variance of a given dimension.

GDOP is found from the Q matrix as:

$$\boldsymbol{GDOP} = \sqrt{\boldsymbol{d}_{\boldsymbol{x}}^{\mathbf{2}} + \boldsymbol{d}_{\boldsymbol{y}}^{\mathbf{2}} + \boldsymbol{d}_{\boldsymbol{z}}^{\mathbf{2}}} \quad \textbf{(13.5)}$$

GDOP has no units and lower values denote greater precision. The Cramer-Rao lower bound for GDOP is 1.

13.2 Time of Flight Equations

This set of equations has been derived by Weisstein[159]. The equations define the intersection points of a pair of circles. While it is possible to calculate the locations of these intersections from circles located anywhere (be they real or imaginary), it is easier to define a coordinate system where one of the centres is at the origin and the other centre along the zero line of one dimension, translation and rotation of the solution can be done later to fit with the solution to a global coordinate system. The equations here are based on the following situation.

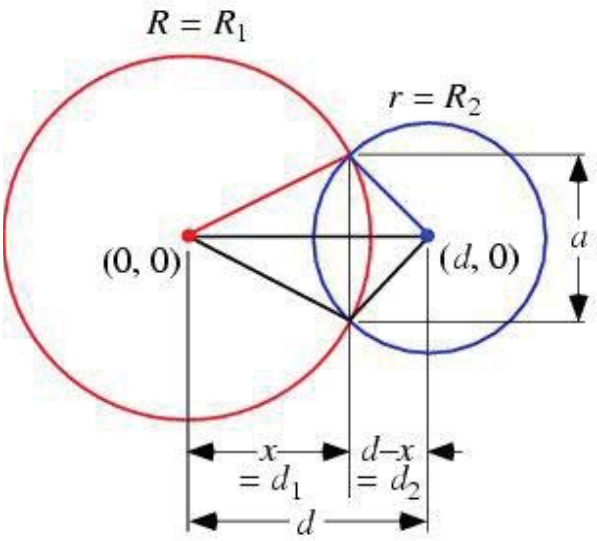

Figure 13.1: Circle-circle intersection situation.

From these two circles the following equations can be defined.

$$x^2 + y^2 = R^2 \quad \textbf{(13.6)}$$

$$(x-d)^2 + y^2 = r^2 \quad \textbf{(13.7)}$$

Rearrange 12.1 for y^2 and substitute into 12.2.

$$(x-d)^2 + (R^2 - x^2) = r^2 \quad \textbf{(13.8)}$$

Multiply through and rearrange

$$x^2 - 2dx - d^2 - x^2 = r^2 - R^2 \quad \textbf{(13.9)}$$

The x^2 values cancel and 12.4 can be rearranged to make x the subject.

$$x = \frac{d^2 - r^2 + R^2}{2d} \quad \textbf{(13.10)}$$

This is then substituted back into equation 12.1

$$y^2 = R^2 - \left(\frac{d^2 - r^2 + R^2}{2d} \right)^2 \quad \textbf{(13.11)}$$

$$y^2 = \frac{4d^2R^2 - (d^2 - r^2 + R^2)}{4d^2} \quad \textbf{(13.12)}$$

Take the square root of 12.7 to find the y ordinates of the intercepts.

$$y = \pm \frac{1}{2d} \sqrt{4d^2R^2 - (d^2 - r^2 + R^2)^2} \quad \textbf{(13.13)}$$

It can clearly be seen that in order to find a region of uncertainty perpendicular to the connecting line between centres, it is necessary to find the straight-line distance between the two intersection points (noted as the distance a in figure 12.1). This distance can be seen to be $2y$. Thus the perpendicular distance is as follows.

$$a = \frac{1}{d} \sqrt{4d^2R^2 - (d^2 - r^2 + R^2)^2} \quad \textbf{(13.14)}$$

With the equations for three circles it is possible to use the radical line[160] (a is a segment of the radical line) from each pair of circles to find the point of intersection for all three circles and hence the single position for the origin of a TOF signal.

13.3 4-Receiver Time Difference of Arrival Equations

This set of equations has been derived by Bucher[161]. The equations define the exact two-dimensional position of an object based on the time difference of arrival of a signal from the object to four fixed receiving stations of known position.

Starting from first principles, assume we have four synchronized receivers (i, j, k and l) all with known positions ((x_i, y_i), (x_j, y_j), (x_k, y_k) and (x_l, y_l) respectively). The distance (R_n) from an object at position (x, y) to the receivers is the time of flight (t_n) multiplied by the speed of light (c). Using this we can define equations 12.1-12.4 as follows.

$$ct_i = R_i = \sqrt{(x_i - x)^2 + (y_i - y)^2} \tag{13.15}$$

$$ct_j = R_j = \sqrt{(x_i - x)^2 + (y_i - y)^2} \tag{13.16}$$

$$ct_k = R_k = \sqrt{(x_k - x)^2 + (y_k - y)^2} \tag{13.17}$$

$$ct_l = R_l = \sqrt{(x_l - x)^2 + (y_l - y)^2} \tag{13.18}$$

Equations 12.5-12.8 are the time difference between the stations.

$$R_{ij} = \sqrt{(x_i - x)^2 + (y_i - y)^2} - \sqrt{(x_j - x)^2 + (y_j - y)^2} \tag{13.19}$$

$$R_{ik} = \sqrt{(x_i - x)^2 + (y_i - y)^2} - \sqrt{(x_k - x)^2 + (y_k - y)^2} \tag{13.20}$$

$$R_{kj} = \sqrt{(x_k - x)^2 + (y_k - y)^2} - \sqrt{(x_j - x)^2 + (y_j - y)^2} \tag{13.21}$$

$$R_{kl} = \sqrt{(x_k - x)^2 + (y_k - y)^2} - \sqrt{(x_l - x)^2 + (y_l - y)^2} \tag{13.22}$$

Move one of the square roots to the LHS.

$$R_{ij} - \sqrt{(x_i - x)^2 + (y_i - y)^2} = -\sqrt{(x_j - x)^2 + (y_j - y)^2} \tag{13.23}$$

$$R_{ik} - \sqrt{(x_i - x)^2 + (y_i - y)^2} = -\sqrt{(x_k - x)^2 + (y_k - y)^2} \tag{13.24}$$

$$R_{kj} - \sqrt{(x_k - x)^2 + (y_k - y)^2} = -\sqrt{(x_j - x)^2 + (y_j - y)^2} \tag{13.25}$$

$$R_{kl} - \sqrt{(x_k - x)^2 + (y_k - y)^2} = -\sqrt{(x_l - x)^2 + (y_l - y)^2} \tag{13.26}$$

Square both sides.

$$R_{ij}^{\ 2} - 2R_{ij}\sqrt{(x_i - x)^2 + (y_i - y)^2} + (x_i - x)^2 + (y_i - y)^2 = (x_j - x)^2 + (y_j - y)^2 \tag{13.27}$$

$$R_{ik}^{\ 2} - 2R_{ik}\sqrt{(x_i - x)^2 + (y_i - y)^2} + (x_i - x)^2 + (y_i - y)^2 = (x_k - x)^2 + (y_k - y)^2 \tag{13.28}$$

$$R_{kj}^{\ 2} - 2R_{kj}\sqrt{(x_k - x)^2 + (y_k - y)^2} + (x_k - x)^2 + (y_k - y)^2 = (x_j - x)^2 + (y_j - y)^2 \tag{13.29}$$

$$R_{kl}^{\ 2} - 2R_{kl}\sqrt{(x_k - x)^2 + (y_k - y)^2} + (x_k - x)^2 + (y_k - y)^2 = (x_l - x)^2 + (y_l - y)^2 \tag{13.30}$$

Expand terms to the right of the square root.

$$\begin{aligned} &R_{ij}^{\ 2} - 2R_{ij}\sqrt{(x_i - x)^2 + (y_i - y)^2} + x_i^{\ 2} - 2x_i x + x^2 + y_i^{\ 2} - 2y_i y + y^2 \\ &= x_j^{\ 2} - 2x_j x + x^2 + y_j^{\ 2} - 2y_k y + y^2 \end{aligned} \tag{13.31}$$

$$\begin{aligned} &R_{ik}^{\ 2} - 2R_{ik}\sqrt{(x_i - x)^2 + (y_i - y)^2} + x_i^{\ 2} - 2x_i x + x^2 + y_i^{\ 2} - 2y_i y + y^2 \\ &= x_k^{\ 2} - 2x_k x + x^2 + y_k^{\ 2} - 2y_k y + y^2 \end{aligned} \tag{13.32}$$

$$R_{kj}{}^2 - 2R_{kj}\sqrt{(x_k - x)^2 + (y_k - y)^2} + x_k{}^2 - 2x_k x + x^2 + y_k{}^2 - 2y_k y + y^2 = x_j{}^2 - 2x_j x + x^2 + y_j - 2y_j y + y^2 \quad (13.33)$$

$$R_{kl}{}^2 - 2R_{kl}\sqrt{(x_k - x)^2 + (y_k - y)^2} + x_k{}^2 - 2x_k x + x^2 + y_k{}^2 - 2y_k y + y^2 = x_l{}^2 - 2x_l x + x^2 + y_l - 2y_l y + y^2 \quad (13.34)$$

Eliminate x^2 and y^2 terms

$$R_{ij}{}^2 - 2R_{ij}\sqrt{(x_i - x)^2 + (y_i - y)^2} + x_i{}^2 - 2x_i x + y_i{}^2 - 2y_i y = x_j{}^2 - 2x_j x + y_j{}^2 - 2y_j y \quad (13.35)$$

$$R_{ik}{}^2 - 2R_{ik}\sqrt{(x_i - x)^2 + (y_i - y)^2} + x_i{}^2 - 2x_i x + y_i{}^2 - 2y_i y = x_k{}^2 - 2x_k x + y_k{}^2 - 2y_k y \quad (13.36)$$

$$R_{kj}{}^2 - 2R_{kj}\sqrt{(x_k - x)^2 + (y_k - y)^2} + x_k{}^2 - 2x_k x + y_k{}^2 - 2y_k y = x_j{}^2 - 2x_j x + y_j{}^2 - 2y_j y \quad (13.37)$$

$$R_{kl}{}^2 - 2R_{kl}\sqrt{(x_k - x)^2 + (y_k - y)^2} + x_k{}^2 - 2x_k x + y_k{}^2 - 2y_k y = x_l{}^2 - 2x_l x + y_l{}^2 - 2y_l y \quad (13.38)$$

Move all but the square root to the RHS and collect similar terms

$$\sqrt{(x_i - x)^2 + (y_i - y)^2} = \frac{R_{ij}{}^2 + x_i{}^2 - x_j{}^2 + y_i{}^2 - y_j{}^2 + 2x_j x - 2x_i x + 2y_j y - 2y_i y}{2R_{ij}} \quad (13.39)$$

$$\sqrt{(x_i - x)^2 + (y_i - y)^2} = \frac{R_{ik}{}^2 + x_i{}^2 - x_k{}^2 + y_i{}^2 - y_k{}^2 + 2x_k x - 2x_i x + 2y_k y - 2y_i y}{2R_{ik}} \quad (13.40)$$

$$\sqrt{(x_k - x)^2 + (y_k - y)^2} = \frac{R_{kj} + x_k{}^2 - x_j{}^2 + y_k{}^2 - y_j{}^2 + 2x_j x - 2x_k x + 2y_j y - 2y_k y}{2R_{kj}} \quad (13.41)$$

$$\sqrt{(x_k - x)^2 + (y_k - y)^2} = \frac{R_{kl} + x_k{}^2 - x_l{}^2 + y_k{}^2 - y_l{}^2 + 2x_l x - 2x_k x + 2y_l y - 2y_k y}{2R_{kl}} \quad (13.42)$$

Simplify equations with dual subscript simplifications ($x_{ji}=x_j-x_i$ etc)

$$\sqrt{(x_i - x)^2 + (y_i - y)^2} = \frac{R_{ij}{}^2 + x_i{}^2 - x_j{}^2 + y_i{}^2 - y_j{}^2 + 2x_{ji} x + 2y_{ji} y}{2R_{ij}} \quad (13.43)$$

$$\sqrt{(x_i - x)^2 + (y_i - y)^2} = \frac{R_{ik}{}^2 + x_i{}^2 - x_k{}^2 + y_i{}^2 - y_k{}^2 + 2x_{ki} x + 2y_{ki} y}{2R_{ik}} \quad (13.44)$$

$$\sqrt{(x_k - x)^2 + (y_k - y)^2} = \frac{R_{kj}{}^2 + x_k{}^2 - x_j{}^2 + y_k{}^2 - y_j{}^2 + 2x_{jk} x + 2y_{jk} y}{2R_{kj}} \quad (13.45)$$

$$\sqrt{(x_k - x)^2 + (y_k - y)^2} = \frac{R_{kl}^{\ 2} + x_k^{\ 2} - x_l^{\ 2} + y_k^{\ 2} - y_l^{\ 2} + 2x_{lk}x + 2y_{lk}y}{2R_{kl}} \tag{13.46}$$

Set equations 12.29 and 12.31 equal to equations 12.30 and 12.32 respectively.

$$\frac{R_{ij}^{\ 2} + x_i^{\ 2} - x_j^{\ 2} + y_i^{\ 2} - y_j^{\ 2} + 2x_{ji}x + 2y_{ji}y}{2R_{ij}} = \frac{R_{ik}^{\ 2} + x_i^{\ 2} - x_k^{\ 2} + y_i^{\ 2} - y_k^{\ 2} + 2x_{ki}x + 2y_{ki}y}{2R_{ik}} \tag{13.47}$$

$$\frac{R_{kj}^{\ 2} + x_k^{\ 2} - x_j^{\ 2} + y_k^{\ 2} - y_j^{\ 2} + 2x_{jk}x + 2y_{jk}y}{2R_{kj}} = \frac{R_{kl}^{\ 2} + x_k^{\ 2} - x_l^{\ 2} + y_k^{\ 2} - y_l^{\ 2} + 2x_{lk}x + 2y_{lk}y}{2R_{kl}} \tag{13.48}$$

Separate knowns and unknowns.

$$\frac{R_{ij}^{\ 2} + x_i^{\ 2} - x_j^{\ 2} + y_i^{\ 2} - y_j^{\ 2}}{2R_{ij}} - \frac{R_{ik}^{\ 2} + x_i^{\ 2} - x_k^{\ 2} + y_i^{\ 2} - y_k^{\ 2}}{2R_{ik}} = \frac{x_{ki}x + y_{ki}y}{R_{ik}} - \frac{x_{ji}x + y_{ji}y}{R_{ij}} \tag{13.49}$$

$$\frac{R_{kj}^{\ 2} + x_k^{\ 2} - x_j^{\ 2} + y_k^{\ 2} - y_j^{\ 2}}{2R_{kj}} - \frac{R_{kl}^{\ 2} + x_k^{\ 2} - x_l^{\ 2} + y_k^{\ 2} - y_l^{\ 2}}{2R_{kl}} = \frac{x_{lk}x + y_{lk}y}{R_{kl}} - \frac{x_{jk}x + y_{jk}y}{R_{kj}} \tag{13.50}$$

Multiply 12.35 by $R_{ij}R_{ik}$ and 12.36 by $R_{kj}R_{kl}$.

$$\frac{R_{ik}[R_{ij}^{\ 2} + x_i^{\ 2} - x_j^{\ 2} + y_i^{\ 2} - y_j^{\ 2}] - R_{ij}[R_{ik}^{\ 2} + x_i^{\ 2} - x_k^{\ 2} + y_i^{\ 2} - y_k^{\ 2}]}{2} = R_{ij}[x_{ki}x + y_{ki}y] - R_{ik}[x_{ji}x + y_{ji}y] \tag{13.51}$$

$$\frac{R_{kl}[R_{kj}^{\ 2} + x_k^{\ 2} - x_j^{\ 2} + y_k^{\ 2} - y_j^{\ 2}] - R_{kj}[R_{kl}^{\ 2} + x_k^{\ 2} - x_l^{\ 2} + y_k^{\ 2} - y_l^{\ 2}]}{2} = R_{kj}[x_{kl}x + y_{kl}y] - R_{kl}[x_{kj}x + y_{kj}y] \tag{13.52}$$

Separate RHS into x and y coefficients.

$$\frac{R_{ik}[R_{ij}^{\ 2} + x_i^{\ 2} - x_j^{\ 2} + y_i^{\ 2} - y_j^{\ 2}] - R_{ij}[R_{ik}^{\ 2} + x_i^{\ 2} - x_k^{\ 2} + y_i^{\ 2} - y_k^{\ 2}]}{2} = [R_{ij}x_{ki} - R_{ik}x_{ji}]x + [R_{ij}y_{ki} - R_{ik}y_{ji}]y \tag{13.53}$$

$$\frac{R_{kl}[R_{kj}^2 + x_k^2 - x_j^2 + y_k^2 - y_j^2] - R_{kj}[R_{kl}^2 + x_k^2 - x_l^2 + y_k^2 - y_l^2]}{2} = [R_{kj}x_{kl} - R_{kl}x_{kj}]x + [R_{kj}y_{kl} - R_{kl}y_{kj}]y \quad (13.54)$$

To get equations 12.39 and 12.40 in the linear form y=Ax+B and y=Cx+D respectively, define the constants A, B, C and D as follows.

$$A = -\frac{R_{ij}x_{ki} - R_{ik}x_{ji}}{R_{ij}y_{ki} - R_{ik}y_{ji}} \quad (13.55)$$

$$B = \frac{R_{ik}[R_{ij}^2 + x_i^2 - x_j^2 + y_i^2 - y_j^2] - R_{ij}[R_{ik}^2 + x_i^2 - x_k^2 + y_i^2 - y_k^2]}{2[R_{ij}y_{ki} - R_{ik}y_{ji}]} \quad (13.56)$$

$$C = -\frac{[R_{kj}x_{kl} - R_{kl}x_{kj}]}{[R_{kj}y_{kl} - R_{kl}y_{kj}]} \quad (13.57)$$

$$D = \frac{R_{kl}[R_{kj}^2 + x_k^2 - x_j^2 + y_k^2 - y_j^2] - R_{kj}[R_{kl}^2 + x_k^2 - x_l^2 + y_k^2 - y_l^2]}{2[R_{kj}y_{kl} - R_{kl}y_{kj}]} \quad (13.58)$$

Solving the equations 12.39 and 12.40 simultaneously using the above simplifications.

$$x = \frac{D - B}{A - C} \quad (13.59)$$

$$y = Ax + B = Cx + D \quad (13.60)$$

13.4 3-Receiver Time Difference of Arrival Equations

This set of equations is based on the derivation presented in the previous section. These equations define the exact two-dimensional position of an object based on temporal information arriving at three fixed receiving stations of known position. This solution uses the smallest number of receiving stations possible to get a unique solution.

Assume that we have an object at position (x, y) and three stationary receivers (i, j and k) at known positions (x_i, y_i), (x_j, y_j) and (x_k, y_k) respectively.

Starting with equations 12.29, 12.30 and 12.31 derived in the same fashion as in section A.1, we can pair 12.29 with 12.30 in the same fashion and arrive at equation 12.39 and use the simplifications 12.41 and 12.42. Substituting the simplified linear equation into equation 12.31 we get the following equation.

$$\sqrt{(x_k - x)^2 + (y_k - Ax - B)^2} = \frac{R_{kj}^2 + x_k^2 - x_j^2 + y_k^2 - y_j^2 + 2x_{jk}x + 2y_{jk}Ax + 2y_{jk}B}{2R_{kj}} \quad (13.61)$$

Factorize the top line of the RHS.

$$\sqrt{(x_k - x)^2 + (y_k - Ax - B)^2}$$
$$= \frac{[R_{kj}{}^2 + x_k{}^2 - x_j{}^2 + y_k{}^2 - y_j{}^2 + 2y_{jk}B] + 2[x_{jk} + y_{jk}A]x}{2R_{kj}} \quad (13.62)$$

Simplify 12.48 by defining the following constants.

$$Y = [R_{kj}{}^2 + x_k{}^2 - x_j{}^2 + y_k{}^2 - y_j{}^2 + 2y_{jk}B] \quad (13.63)$$
$$X = 2[x_{jk} + y_{jk}A] \quad (13.64)$$

Equation 12.48 can now be written as follows.

$$\sqrt{(x_k - x)^2 + (y_k - y)^2} = \frac{Y + Xx}{2R_{kj}} \quad (13.65)$$

Square both sides of equation 12.51 to remove the square root.

$$(x_k - x)^2 + (y_k - Ax - B)^2 = \frac{Y^2 + YXx + X^2x^2}{4R_{kj}{}^2} \quad (13.66)$$

Define constants D, E and F as follows for ease of reading.

$$D = \frac{Y^2}{4R_{kj}{}^2} \quad (13.67)$$
$$E = \frac{YX}{4R_{kj}{}^2} \quad (13.68)$$
$$F = \frac{X^2}{4R_{kj}{}^2} \quad (13.69)$$

Equation 12.52 can now be written as follows.

$$(x_k - x)^2 + (y_k - Ax - B)^2 = D + Ex + Fx^2 \quad (13.70)$$

Expand the LHS of 12.56.

$$x_k{}^2 - 2x_kx + x^2 + y_k{}^2 - 2y_kAx - 2y_kB + A^2x^2 + 2ABx + B^2 = D + Ex + Fx^2 \quad (13.71)$$

Factorize the LHS

$$[x_k{}^2 + y_k{}^2 - 2y_kB + B^2] + [2AB - 2x_k - 2y_kA]x + [1 + A^2]x^2 = D + Ex + Fx^2 \quad (13.72)$$

Define coefficients G, H and I for ease of reading.

$$G = x_k{}^2 + y_k{}^2 - 2y_kB + B^2 \quad (13.73)$$
$$H = 2[AB - x_k - y_kA] \quad (13.74)$$
$$I = 1 + A^2 \quad (13.75)$$

Equation 12.58 can now be written as follows.

$$G + Hx + Ix^2 = D + Ex + Fx^2 \tag{13.76}$$

Combining like terms

$$G - D = [F - I]x^2 + [E - H]x \tag{13.77}$$

Divide through by F-I

$$\frac{G - D}{F - I} = x^2 + \frac{E - H}{F - I}x \tag{13.78}$$

Complete the square

$$x^2 + \frac{E - H}{F - I}x + \left[\frac{E - H}{2[F - I]}\right]^2 = \frac{G - D}{F - I} + \left[\frac{E - H}{2[F - I]}\right]^2 \tag{13.79}$$

Factorize the LHS

$$\left[x + \frac{E - H}{2[F - I]}\right]^2 = \frac{G - D}{F - I} + \left[\frac{E - H}{2[F - I]}\right]^2 \tag{13.80}$$

Rearrange to make x the subject

$$x = \pm\sqrt{\frac{G - D}{F - I} + \left[\frac{E - H}{2[F - I]}\right]^2} - \frac{E - H}{2[F - I]} \tag{13.81}$$

This value can be used to find y by using $y=Ax+B$.

14Appendix B: Plots

The collected plots from simulations and testing performed throughout the research.

14.1 Occupancy plots

In the text, only the cumulative occupancy was presented. Here the occupancy of each of the games in the data suite provided by the Industrial Partner is presented individually. They can also be viewed on the attached DVD. The electronic versions of these plots can be rotated and zoomed as desired. All plots are shown to 2m accuracy in both the length and width dimensions.

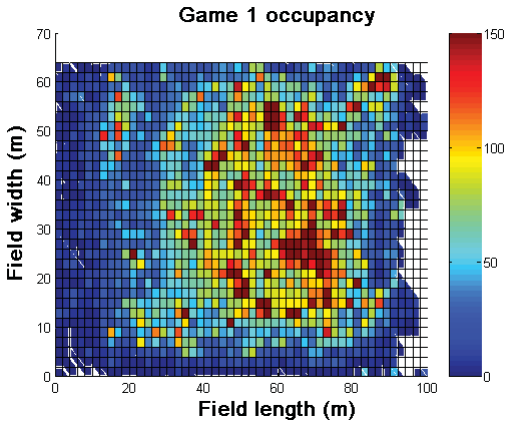

Figure 14.1: Game 1 occupancy.

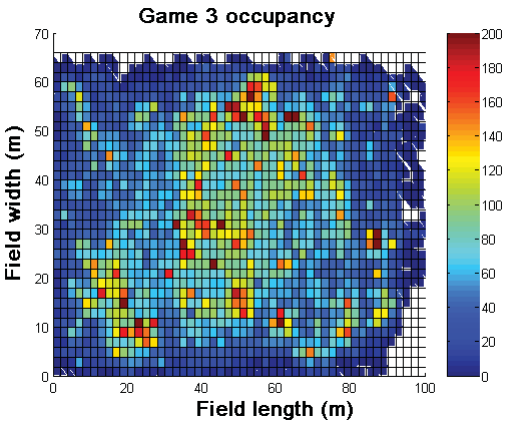

Figure 14.3: Game 3 occupancy.

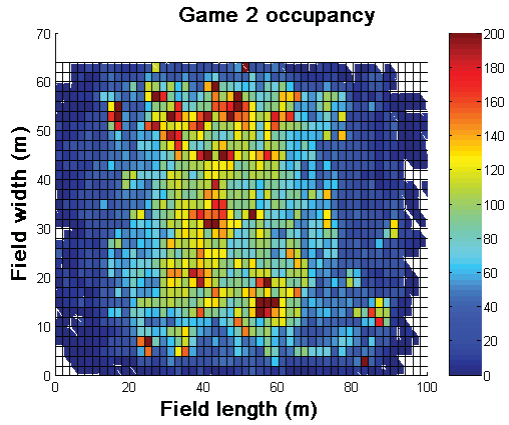

Figure 14.2: Game 2 occupancy.

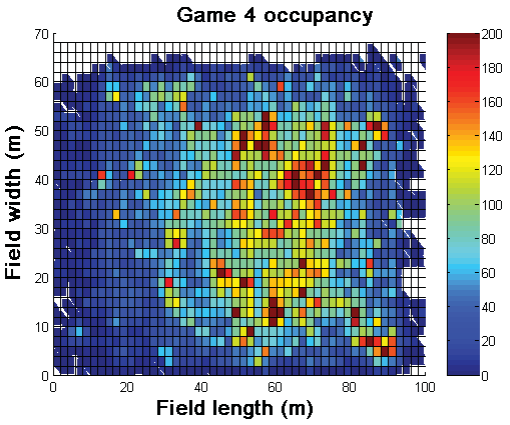

Figure 14.4: Game 4 occupancy.

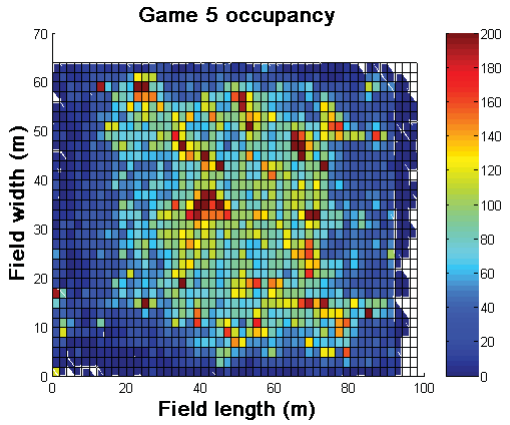

Figure 14.5: Game 5 occupancy.

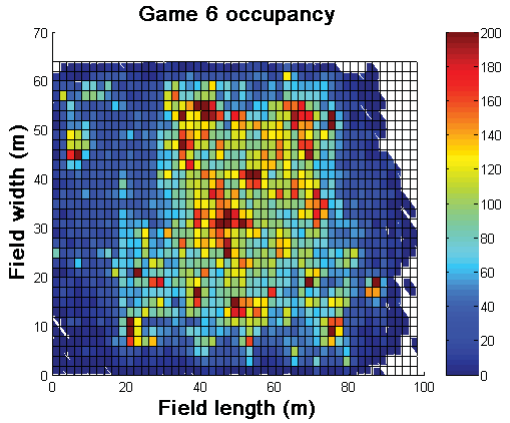

Figure 14.6: Game 6 occupancy.

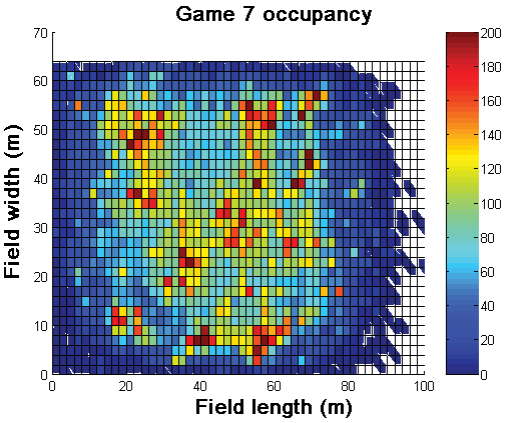

Figure 14.7: Game 7 occupancy.

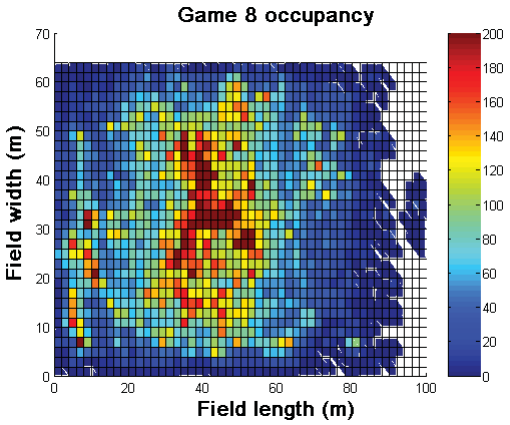

Figure 14.8: Game 8 occupancy.

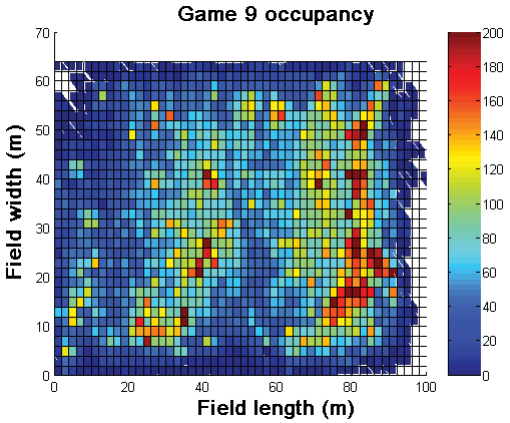

Figure 14.9: Game 9 occupancy.

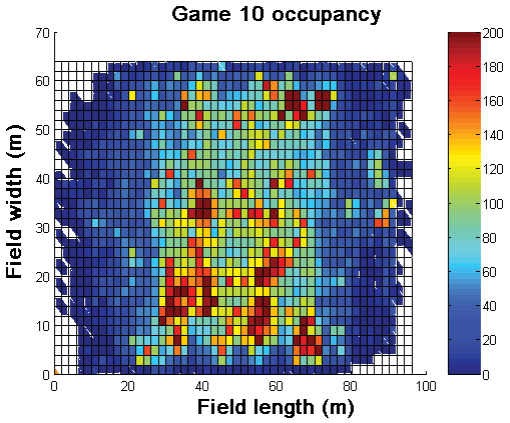

Figure 14.10: Game 10 occupancy.

15 Appendix C: DVD index

In which a map is given to the files included on the attached DVD.

The attached DVD contains electronic copies of this thesis, the source code for the developed algorithms and a number of the figures that were used in the thesis. The figures are included to allow the reader to zoom and rotate the plots, since some of the presented concepts are easier to view in three dimensions. The files are arranged by chapter where appropriate. A basic description of the files and their path is given in Table 15.1.

Table 15.1: DVD index.

File name	Description	Path
Local Area Positioning of Multiple Moving Objects.pdf	Electronic copy of this thesis.	Root directory
Destroyer4_2.m	The player visibility model, used to find which players will not be visible to the tracking system.	\6 Player visibility model\
Figure5_2.fig	Matlab plot of the full data suite occupancy.	\6 Player visibility model\
Visibility model animation.mpeg	Animation of the game data suite after application of the player visibility model.	\6 Player visibility model\
MaxCover4Eval.m	The objective function used in section 8.5.2.	\7 Receiver placement optimisation\
MaxCover4Ga.m	Main function for the max cover optimisation.	\7 Receiver placement optimisation\
MaxCover4LUTEval3.m	Objective function used in section 8.5.3.	\7 Receiver placement optimisation\
MaxCover4LUTGa3.m	Main function for the look-up-table optimisation.	\7 Receiver placement optimisation\
OccupancyEval3.m	Objective function used in section 8.5.8.	\6 Receiver placement optimisation\

Perimtoxy.m	Function to convert single-dimension, receiver perimeter values (input to objective function) into Cartesian coordinates.	\7 Receiver placement optimisation \
PerturbSol.m	Function to compare the fitness of the optimum solution with that of randomly perturbed solutions.	\7 Receiver placement optimisation \
ConvertGame.m	Main function for converting the Multispectral log file into a form useable by the Industrial Partner. This includes application of the reconstruction algorithm.	\8 Position reconstruction \
CsvGenerator.m	Converts any given Matlab array into a comma separated value file.	\8 Position reconstruction \
FindSpeed.m	Calculate the speed of each player.	\8 Position reconstruction \
FixErrorFn.m	The reconstruction main function.	\8 Position reconstruction \
GameInterface.m	A function to perform the quantisation of player positions into frames.	\8Position reconstruction \
GetTags.m	Finds what tag IDs are present in the log file.	\8 Position reconstruction \
SituationErrorFix.m	Performs the application specific position reconstruction.	\7 Position reconstruction \
SmallErrorFix.m	Performs the interpolation reconstruction	\8 Position reconstruction \
TagFileReader.m	Function to import and parse Multispectral log file data	\8 Position reconstruction \
TwoDOnly.m	Function to filter the data such that only two-dimensional position data are passed.	\8 Position reconstruction \
Figure13_x.fig	Individual occupancy plots for each game in the data suite. Where x is the game number (1-10).	\13 Appendix B \
Occupancy.m	Function for generating an occupancy plot	\13 Appendix B \

16Appendix D: Nomenclature

In which the acronyms and symbols are summarized.

Table 16.1 provides a glossary to the acronyms used in the thesis.

Table 16.1: Table of acronyms used in this thesis.

Acronym	Definition
AOA	Angle of Arrival
BS	Base station
COTS	Commercial-off-the-shelf
DC	Direct Current
FCC	Federal Communications Commission
FDMA	Frequency Division Multiple Access
FIRA	Federation of International Robot-soccer Association
GA	Genetic algorithm
GAOT	Genetic algorithm optimisation toolbox
GDOP	Geometric Dilution of Precision
GPS	Global Positioning System
ID	Identification
IEEE	Institute of Electrical and Electronic Engineers
IRB	International Rugby Board
NGZC	Negative Going Zero Crossing
NTP	Network Timing Protocol
PAA	Phased Array Antenna

PLN	Position Location Network
POE	Power over Ethernet
PSD	Power Spectral Density
RBS	Reference Broadcast Synchronisation
RFID	Radio Frequency Identification
RGB	Red, green, blue. A colour-encoding scheme for images.
Rx	Short hand for receiver
SLI	Signal Loss Interval
TDMA	Time Division Multiple Access
TDOA	Time Difference of Arrival
TOA	Time of Arrival. Synonymous with TOF
TOF	Time of flight. Synonymous with TOA
Tx	Short hand for transmitter
USB	Universal serial bus
UWB	Ultra Wideband

Table 16.2 gives a glossary of the symbols used in this thesis.

Table 16.2: Table of symbols.

Symbol	**Definition**	**Units**
α_r	AOA value calculated from an iteration of an optimisation	deg
a	In equation (2.1): Angle of a triangle opposite to side A In chapter 13: The length of the chord defined by two, overlapping circles	deg m
A	In equation (2.1): Distance between the right receiver and the object In chapter 13: Substitution to simplify the reading of equations	m no units
a_r	AOA value measured at receiver *r*	deg

B	In equation (2.1): Distance between the left receiver and the object In chapter 13: Substitution to simplify the reading of equations	m no units
b	In equation (2.1): Angle of a triangle opposite to side B	deg
C	In equation (2.1): Distance between the receivers In chapter 13: Substitution to simplify the reading of equations	m
c	Speed of light (approximated as $3x10^8 ms^{-1}$)	ms^{-1}
D	In chapter 13: Substitution to simplify the reading of equations	no units
d	Distance	m
E	In chapter 13: Substitution to simplify the reading of equations	no units
e	Error, usually the difference between a measured and calculated value in an objective function.	m
F	In chapter 13: Substitution to simplify the reading of equations	no units
G	In chapter 13: Substitution to simplify the reading of equations	no units
H	In chapter 13: Substitution to simplify the reading of equations	no units
h	Total number of receiver pairings	no units
I	In chapter 13: Substitution to simplify the reading of equations	no units
N	Total number of receivers	no units
Px	X ordinate of a player in the data suite. Usually referenced by frame and player.	
Py	Y ordinate of a player in the data suite. Usually referenced by frame and player.	
r	Receiver number	no units
range	The maximum range of a given receiver-transmitter pair	m
rolloff	The probability that a signal will be visible to a receiver as range increases	no units
R	Radial distance	m
R_{tr}	The range between transmitter *t* and receiver *r*	m

t	Time	s
t_0	Optimisation variable representing the epoch start	s
T_r	TOF measured by receiver *r*	m
v	Velocity	ms^{-1}
x	The x-ordinate of an object. Either the true value or an iteration.	m
X	In chapter 13: As a collection of constants to simplify the reading of equations.	no units
X_r or x_r	The x-ordinate of receiver *r*. Where *r* can be any subscript appropriate to the equation (e.g. letters *i-l*).	m
y	The y-ordinate of an object. Either the true value or an iteration	m
Y	In chapter 13: As a collection of constants to simplify the reading of equations	no units
Y_r or y_r	The y-ordinate of receiver *r*. Where *r* can be any subscript appropriate to the equation (e.g. letters *i-l*).	m