

Faculty of Engineering

Precise Location Measurement in Indoor and Outdoor Performance Sports

Bertrand Perrat

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

September 2016

Abstract

Knowing precisely where an athlete is with regards to their environment and the motions involved is of critical interest from a sport performance perspective. Engineering and technology provide a significant opportunity for sport science to achieve this. The research investigates and develops state-of-the-art positioning technologies and techniques to be used for such sport tracking applications. Further addressing the real-world use of such techniques, the thesis additionally focuses on the development of sport specific metrics from such tracking data that can be used for performance analysis. Two separate case studies have supported this research, the first one focusing on precise location measurement indoors, the second one outdoors. Each case study has a targeted sport application.

The investigation of indoor precise location measurement for performance sports focuses on the use of a commercial Ultra-Wide Band (UWB) system for wheelchair court sports athlete tracking. UWB radio positioning systems are maturing very quickly and now represent a good candidate for indoor positioning. Several aspects have been investigated including system setup, calibration, sensor positioning, determination of sport performance indicators and quality assessment of the output. With a simple setup procedure, it has been demonstrated that athletes tracking can be achieved with an average horizontal positioning error of 0.37 m ($\sigma = \pm 0.24$ m). Distance covered can be computed after data processing with an error below 0.5% of the course length. It has also been demonstrated that the acquisition rate and the number of wheelchairs on the court does not affect significantly the positioning quality; however, for highly dynamic movement tracking, higher rates

are recommended for finer dynamic recording. To further support the sport science community, sport specific metrics derived from the UWB system output have been proposed and integrated into a software that has been used in a number of studies.

The investigation of outdoor precise location measurement for performance sports focuses on the use of GPS and inertial technologies for sprint canoeing and kayaking tracking. The requirements directing this research were the development of a solution that can achieve accurate sprint timing as well as provide new insights on the sprinting dynamics. The combined use of GPS and inertial technologies results in an inertial navigation solution producing position, velocity and orientation estimates of the boat at high rate. A novel approach is proposed to achieve the sensor fusion of GPS and IMU measurements through the use of factor graph modelling and efficient solvers. A key advantage of the proposed approach is to allow a complete re-linearisation of the system as opposed to traditional Kalman Filter approaches. As part of this research, a novel factor graph formulation of an advanced GPS positioning technique has been developed, allowing precise relative positioning without the need for a base station. In addition, a novel "tightly-coupled" approach is proposed to jointly optimise GPS and IMU measurements. The quality assessment of the proposed approach has demonstrated that horizontal positioning is achieved with a mean error of 4.9 cm ($\sigma = \pm 2.9$ cm) over the 30 minutes of an on-water trial. Boat velocity was estimated with an accuracy of a few mm/s. Roll and pitch angles reported an accuracy better than a $1/10^{th}$ of a degree and yaw a standard deviation of 0.336°. From a sport perspective, 10-metres split timing is robustly achieved within a $1/100^{th}$ of a second. Similarly, sprint timing for distances ranging from 100 to 500 metres is achieved with an accuracy of a few $1/100^{th}$ of a second. A stroke detection algorithm has been developed and demonstrated a detection rate of 99.5% over more than 20000 strokes. The high-rate and the precision of the solution obtained with the proposed approach combined with the accurate stroke detection allows the creation of a whole range of new metrics providing new insights on the sprinting dynamics. A selection of

sport specific metrics are proposed as part of this work, highlighting the potential of the developed approach.

Acknowledgements

This PhD has been undertaken at the Faculty of Engineering, University of Nottingham in collaboration with the English Institute of Sport. Funding for the project has been supplied by both the Engineering and Physical Sciences Research Council (EPSRC) and the English Institute of Sport (EIS) through an Industrial CASE award.

I would first like to thank my academic supervisors Dr Martin Smith and Professor Terry Moore for their support.

I would also like to thank the English Institute of Sport for their continuous support during this research. I would in particular like to thank Dr Barney Wainwright, Dave Thomas, Jono Byrne and Gavin Atkins for sharing their knowledge and interesting feedback from a sport perspective. I would like to extends my thanks to all the coaches and athletes who helped with the different practical trials.

I would also like to thanks Dr Barry Mason and Dr James Rhodes, from Loughborough University, for the enjoyable and successful collaborative work on the indoor tracking system.

I would like to say a very particular thank you to Dr Chris Hide who raised my interest for navigation problems before this PhD and all his very good advices.

I would finally like to extend my thanks to everyone else who helped with this work in any ways and in particular Gavin, James and Mark for their feedback on the thesis.

Contents

I. Introduction **19**

0.1. Background 20

0.2. Research Aims and Objectives 22

0.3. Thesis Outline 24

II. Indoor Athlete Tracking for Wheelchair Court Sports **26**

1. Introduction **27**

1.1. Aims 28

1.2. Objectives 28

2. Indoor tracking technologies **30**

2.1. Background 30

2.2. Ubisense Real-Time Location System 31

3. Research Methodology **34**

3.1. Trials and location facility 34

3.2. Surveying Equipment 35

3.3. Data collection 37

3.4. Data processing and smoothing 37

4. Trials, results and analysis **39**

4.1. System set-up 39

4.1.1. Set-up and configuration quality 39

4.1.2. Sensors spatial configuration 40
4.2. Positioning quality analysis 41
4.2.1. Stationary positioning 41
4.2.2. Dynamic positioning 42
4.2.3. Impact of wheelchair environment 46
4.2.4. Filtered positioning quality analysis for distance measurement 49
4.2.5. Tag mounting location 50

5. Sport specific metrics **53**
5.1. Speed and Speed Zones 53
5.2. Accelerations and Sprints 55
5.3. Court occupation 56
5.4. Analysis software 57

6. Conclusion **61**
6.1. Contribution to knowledge 62
6.2. Impact for elite sports 62
6.3. Discussion : Other Sports Applicability 63

III. Outdoor Athlete Tracking for Sprint Canoeing & Kayaking 64

7. Introduction **65**
7.1. Aims 65
7.2. Objectives 66
7.3. Sport requirements, constraints and goals 67

8. Review of existing systems and technologies available in the sport **69**
8.1. Existing commercial systems 69
8.1.1. Catapult - MinimaxX 69
8.1.2. Igtimi - YatchBot Tracker 70
8.1.3. McLaren Applied Technologies - Athena 70
8.1.4. Vaaka - Kayak Cadence Sensor 71
8.1.5. Limitations of the existing systems 72

8.2. Outdoor tracking technologies 76
8.3. Overview of the selected approach 80

9. Graph-based problem optimisation **82**
9.1. Probabilistic formulation 83
9.2. Graph representation 85
9.3. Non-linear Least Squares 87
9.3.1. Levenberg-Marquardt 89
9.3.2. Local Parametrisations 90
9.3.3. Loss Functions 92

10. GPS Positioning **93**
10.1. GPS Fundamentals 94
10.1.1. System description 94
10.1.2. GPS Observables 96
10.1.3. Errors and Correction Models 102
10.2. Time-Differenced Carrier Phase Positioning 109
10.2.1. Time-Differences of Phase Measurements 109
10.2.2. Positioning algorithm 111
10.2.3. Trajectory reconstruction 113
10.2.4. Errors after time-differentiation 114

11. Strapdown Inertial Navigation **118**
11.1. Introduction 118
11.2. Coordinate frames 119
11.2.1. Inertial Frame 119
11.2.2. Earth Frame 120
11.2.3. Navigation Frame 121
11.2.4. Body Frame 121
11.3. Rotations 122
11.3.1. Direction Cosine Matrix 122
11.3.2. Euler angles 123
11.3.3. Quaternions 123

11.4. Gravity compensation . . . 125
11.5. Navigation frame initialisation and mechanisation . . . 126
11.5.1. Initialisation and alignment . . . 127
11.5.2. Navigation Frame Mechanisation . . . 131

12. A Novel Graph-Based GPS and INS integration 134
12.1. Time-Differenced Carrier Phase Positioning . . . 134
12.1.1. A Novel Graph-based implementation . . . 134
12.1.2. Loss Function . . . 136
12.1.3. Weighting Scheme . . . 136
12.1.4. Ephemeris . . . 141
12.2. Graph-based Inertial Navigation . . . 142
12.2.1. A novel tightly-coupled integration for GPS and IMU sensor fusion . . . 143
12.2.2. Summarised IMU Factor . . . 145
12.2.3. Bias modelling . . . 147
12.3. Taking advantage of the application specificities . . . 148
12.3.1. Height constraint . . . 148
12.3.2. Loop closure . . . 150

13. Practical validation 151
13.1. GNSS - Static validation . . . 151
13.1.1. Equipment and location facility . . . 152
13.1.2. Methodology . . . 152
13.1.3. Results . . . 155
13.2. On-water sprint kayak validation . . . 165
13.2.1. Trials and location facility . . . 166
13.2.2. Commercial Equipment . . . 166
13.2.3. Developed system . . . 168
13.2.4. Data collection . . . 171
13.2.5. Methodology . . . 172
13.2.6. Results . . . 174

14. Sport specific metrics **188**

14.1. Distance estimation . . . 188

14.1.1. Haversine . . . 189

14.1.2. Vincenty's formulae . . . 190

14.1.3. Local tangent plane . . . 190

14.1.4. Comparison . . . 191

14.2. Sprint timing . . . 193

14.2.1. Potential distance error . . . 196

14.3. Aligned boat velocity . . . 198

14.4. Stroke analysis . . . 200

14.4.1. Kayaking dynamics . . . 201

14.4.2. Stroke detection . . . 203

14.4.3. Stroke rate . . . 205

14.4.4. Stroke distance . . . 206

14.4.5. Stroke efficiency . . . 206

14.4.6. Stroke velocity change . . . 208

14.4.7. Stroke phases . . . 209

15. Conclusion **212**

15.1. Contribution to knowledge . . . 216

15.2. Impact for elite sports . . . 217

15.3. Discussion : Other Sports Applicability . . . 218

IV. Summary and Conclusions **219**

List of Figures

2.1. Ubisense real-time location system 32

3.1. Viewing gallery showing the Leica TS-30 surveying equipment and one UWB sensor on an elevated stand 34
3.2. Interpolation benefits to more closely follow the track 36
3.3. Tags and prism mounted on a wheelchair 36
3.4. Workflow of the processing used in this study 38

4.1. Trajectory pattern used for sensors spatial configuration analysis . 41
4.2. Stationary tag positioning 42
4.3. Example trajectory during a quarter of a simulated match obtained with the Leica TS-30 43
4.4. Spatial distribution of positioning error for a particular configuration and setup of sensors 44
4.5. Cumulative distribution function of the positioning error 45
4.6. Distribution of the data gaps with one wheelchair on the court . . 47
4.7. Distribution of the data gaps during a real match with 8 athletes + 1 referee on the court 47
4.8. Example of an athletes trajectory during one quarter of a real match 48

5.1. Example output of the developed software offering a report of activity for athletes and coaches 58
5.2. Example output of the developed match report software offering a full match report for an athlete 58

5.3. Example output of the developed software for enhanced analysis displaying the time spent in different areas of the court 59
5.4. Example output of the developed software for enhanced analysis displaying the distance covered in different areas of the court . . . 59

8.1. Horizontal positioning error in metres during a 500 metres sprint . 72
8.2. Boat velocity during the first part of a 500 metres sprint 73
8.3. Boat velocity error reported by the Athena device during a 500 metres sprint 74

9.1. Factor graph representation of a navigation problem 87

12.1. Factor graph representation of time-differenced carrier phase positioning 134
12.2. Comparison of the different elevation based weighting schemes . . . 138
12.3. Factor graph representation of strapdown inertial navigation 143
12.4. Factor graph representation of strapdown inertial navigation integrating GPS positions in a loosely coupled form 143
12.5. Factor graph representation of the proposed tightly-coupled integration of strapdown inertial navigation and time-differenced carrier phase positioning 144
12.6. Factor graph representation of the developed tightly-coupled integration with the use of pre-integrated IMU factors. 145
12.7. Factor graph representation of strapdown inertial navigation with bias modelling factors 147
12.8. Factor graph representation of strapdown inertial navigation integrating height constraint factors 149
12.9. Factor graph representation of strapdown inertial navigation integrating a loop closure factor 150

13.1. Positioning error in metres obtained with the u-blox LEA-6T and the Leica GS10 receivers 155

13.2. Drift rates of the positioning error in millimetres per second obtained with the u-blox LEA-6T over different periods during the static validation trial 156
13.3. Positioning error in metres with and without the loop closure constraint 158
13.4. Elevation Mask 160
13.5. Positioning error in metres obtained with different weighting schemes 161
13.6. 3D Positioning error in metres obtained with height constraint factors inserted at different rates 162
13.7. Vertical Positioning error in metres obtained with and without height constraints 163
13.8. Positioning error in metres obtained with and without the use of loss function 164
13.9. Positioning error in metres obtained with and without the use of precise satellite ephemeris and clocks 165
13.10. Fully equipped K1 boat for the on-water trials 168
13.11. Overview of the developed system installed in the K1 boat 170
13.12. Fully equipped K1 boat during the on-water validation trial 173
13.13. Position trace of the on-water validation trial 174
13.14. Positioning error in metres obtained by the u-blox LEA-6T with the developed time-differenced carrier phase positioning technique and the RTK solution 175
13.15. Drift rates of the positioning error in millimetres per second obtained with the u-blox LEA-6T over different periods during the on-water validation trial 176
13.16. Positioning error in metres with and without the loop closure constraint 179
13.17. Positioning error in metres obtained with different weighting schemes 180
13.18. Elevation Mask 181
13.19. 3D Positioning error in metres obtained with different height constraints 181

13.20. Vertical Positioning error in metres obtained with and without height constraints . . . 182
13.21. Positioning error in metres obtained with and without the use of loss function . . . 183
13.22. Positioning error in metres obtained with and without the use of precise satellite ephemeris and clocks . . . 184
13.23. Histogram of the 3 dimensional velocity error over the 30 minutes trial . . . 185
13.24. Histogram of the planar velocity error over the 30 minutes trial . . 186
13.25. Attitude error in degrees over the 30 minutes trial . . . 186

14.1. Histogram of the timing error on 10 metres splits using 110 splits . 196
14.2. 3D Velocity as obtained from the processing developed with East, North and Up components . . . 198
14.3. Velocity aligned to the sprint heading with Forward, Lateral and Up components . . . 200
14.4. 3-Axis acceleration during 5 strokes . . . 201
14.5. 3-Axis angular velocity during 5 strokes . . . 202
14.6. Roll, Pitch and Yaw angles during 5 strokes . . . 202
14.7. Example of stroke detection using the forward acceleration (5 strokes) 204
14.8. Stroke duration over a 500 m sprint . . . 205
14.9. Stroke rate over a 500 m sprint . . . 205
14.10. Distance covered by each stroke over the 500 m sprint . . . 206
14.11. Stroke distance against stroke duration for each side excluding the first 40 strokes . . . 207
14.12. Stroke distance against stroke duration for each side including the first 40 strokes . . . 207
14.13. Velocity difference between the start and the end of each stroke . . 208
14.14. Detection of the different phases of a stroke . . . 209
14.15. Time duration of the different stroke phases for strokes on side 1 . 210
14.16. Time duration of the different stroke phases for strokes on side 2 . 210
14.17. Time duration of the different stroke phases for all the strokes . . . 211

List of Tables

4.1. Sensors position differences using Ubisense LEC tool 40

4.2. Mean (m) ± standard deviation (m) of the UWB positioning error during two simulated matches with nine tags operating at three different update rates 44

4.3. Distance error of nine tags operating at three different update rates during two simulated matches 50

4.4. Distance estimates and data gaps for different mounting locations . 51

5.1. Speed thresholds in m/s for the different athletes' classes 54

5.2. Speed zones definition based on the thresholds 54

5.3. Relative speed zones definition based on the maximal speed achieved by the athlete 55

8.1. Mean and standard deviation in metres of the horizontal positioning error during a 500 metres sprint. 73

13.1. Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the Leica GS10 receiver 157

13.2. Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the u-blox LEA-6T 157

13.3. Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the u-blox LEA-6T without loop closure 159
13.4. Different types of ephemeris and clock products investigated and their availability 164
13.5. Sprint efforts of the on-water trial session and their associated distance 166
13.6. Drift rates in millimetres per second over different periods obtained with the u-blox LEA-6T with the developed technique 177
13.7. Drift rates in millimetres per second over different periods obtained with the u-blox LEA-6T with RTK processing 177
13.8. Drift rates millimetres per second over different periods obtained with the u-blox LEA-6T without the loop closure constraint 179
13.9. Mean attitude error largely due to misalignment between the units 187
13.10. Standard deviation of the attitude error 187

14.1. Comparison of the different distance calculations with distance in metres and the error percentage 192
14.2. Comparison of the timing error induced by the different distance calculations 192
14.3. Comparison of sprint timing performance using the efforts of the on-water trial session 194
14.4. Comparison of sprint timing performance using the efforts of the on-water trial session after compensation of the start time difference 195
14.5. Distance error induced if the athlete starts from one side of the sprint lane and finishes on the other side 197
14.6. Distance error induced if the athlete deviates by 4 metres between the start and the finish line 198
14.7. Summary of the stroke detection for 2 athletes using different units over multiple days 204

List of Abbreviations

AOA	Angle Of Arrival
ECEF	Earth-Centred, Earth-Fixed
EKF	Extended Kalman Filter
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
IGS	International GNSS Service
IMU	Inertial Measurement Unit
INS	Inertial Navigation Solution
KF	Kalman Filter
MEMS	Microelectromechanical systems
PDF	Probability Density Function
PLL	Phase Lock Loop
PRN	Pseudo Random Noise
RTK	Real-Time-Kinematic
SLAM	Simultaneous Localisation And Mapping
SNR	Signal-to-Noise Ratio
TDD	Temporal Double Differencing
TDOA	Time Difference Of Arrival
TEC	Total Electron Content
TOA	Time Of Arrival
UWB	Ultra-wideband

Part I.

Introduction

0.1. Background

Sports science is a continuously expanding field studying how the human body works during exercise, from a cellular to a whole body perspective. In the past, sport scientists had a particular focus on areas such as biomechanics, physiology and neuroscience. Biomechanics is of significant interest, aiming to find optimal movements leading to specific strength and conditioning training programs. Sport scientists use physiology to assess the energetic cost of a specific movement or pattern. It is also used to assess how metabolism changes with training. Neuroscience also holds great interest, analysing areas ranging from the muscular recruitment capabilities of an athlete to the more complex behavioural neuroscience that is behind strategic decisions of an athlete during a game.

So far, the common point of all the disciplines mentioned is that they all are directly linked to the athlete's body and thus have a clear, physical, link to the athlete's performance. There is, however, growing interest in sports science towards engineering and technology. This ranges from optimising the equipment used in a sport to the development of technology and techniques to measure and quantify the athletic performance. These two aspects are closely linked, as the benefits of a piece of equipment can only be assessed if you have a suitable technique or technology to measure its impact.

Technology has already made a significant impact on many sports over the years. First by providing ways of measuring physiological parameters of an athlete during an effort with the wide use of heart rate monitoring [8, 49, 66], force-plates [83, 92, 148] or strain gauges and power measurement [13, 48] for example. In addition, technology also has a great track record of improving athletic performance by directly improving the equipment used in a specific sport. Some examples are the research on shoes stiffness, whether it is for sprinting [121] or cycling [60]. The evolution of the tennis rackets from wood to composite materials is an apt example, as are the constant aerodynamic developments for speed based sports such as

bobsleigh [30], skiing [12] or cycling [78]. History has also proved that technology improvements can sometimes cross the line of being ethical and acceptable with several controversies occurring. A first example in tennis is the development of the 'spaghetti' stringing system [42, 43] which has been quickly banned by the International Tennis Federation for imparting significantly more spin on the ball than a conventionally strung racket can possibly do. A second example is the more recent controversy around the polyurethane swim suits, which led to many new world record being set when they appeared in 2008 and 2009 before being banned by the International Swimming Federation (FINA) in 2010. A final example is the case of Oscar Pistorius, a double below-the-knee amputate, who competed in the Olympics after several overturned decision based on fairness and the eventual advantages provided by his highly specialized prosthetic limbs [63].

A growing focus in the field is the interaction between the athletes and their environment as a large majority of the sports have a big spatial component. The most obvious ones are all the racing events. Whether it is running, skiing, cycling or on a boat, these all share the same goal: cover a defined distance or course as fast as possible. Although the ultimate goal is to have the fastest finishing time, knowing precisely how the athlete covered the course or distance can provide very useful insights on performance.

From a technology perspective, most of the racing events rely on photocell timing, otherwise known as 'timing gates'. Here a beam is broken by the athlete at the start and at the finish line and perhaps augmented with high-speed cameras. These systems provide a very accurate measurement of the racing time. However, a limitation is that they measure the time difference between two sensors but don't have details of what is happening in between and ultimately, the number of measurements obtained is linear to the number of timing gates used.

Similarly, team sports represent a large group of activities with a major spatial component. Not only would it be useful to measure an individual athletes perfor-

mance, but to examine the relationship between players, which is vitally important for strategic and tactical analysis. The best team is not only the team with the *best individuals*, it critically relies on the interaction between all the players and how they react to the opposition.

From this perspective, there is a growing interest in knowing where an athlete is on a pitch, track or circuit. The awareness of this problem has massively increased recently with the development of cheap tracking devices and associated applications (smartphone, GPS watch, ...). However, elite sports have very different requirements in terms of accuracy, weight and size [9]. It all comes down to being able to precisely locate an athlete in their environment given the specificities of the sport. This is where the use of advanced positioning technologies can play a major role. There are many different positioning technologies available, each of them having a specific set of advantages and limitations. Making an informed choice on the technology used based on the requirements and constraints of the targeted sport is key in obtaining valuable output. In addition, there is a critical need to validate the selected technology in the targeted sport environment. Most positioning technologies have not been developed for any particular sport application and, by definition, elite sport is a very challenging environment involving high dynamics and pushing humans and engineering limits. This stresses the importance of an in-depth review of the selected technology as well as a quality assessment of its performance in the context of the targeted application. Finally, giving a list of x, y positions to athletes, coaches and performance analysts is of limited interest and value. There is a further need to convert the output of the positioning system into measurements and metrics that are meaningful for the coaches and athletes.

0.2. Research Aims and Objectives

This thesis aims to investigate how the latest advances in positioning techniques and technologies can be used for high performance athlete (or equipment) tracking

and propose novel approaches. The work has two main focuses, indoor sport and outdoor sport. A clear separation has been made as both environments rely on different technologies and have varying constraints although the two parts actually share the same goals. Within each part the objectives will be :

1. To undertake an assessment of the state-of-the-art

2. To develop a proposal for the enhancement of the selected positioning system

3. To assess the quality of the developed solution

4. To develop sport specific metrics

5. To undertake a critical review of the work including discussions and conclusions

The research work is supported by two case studies; the indoor sport considered is wheelchair rugby and is thus team focused. The outdoor sport is sprint canoeing and kayaking and is more individual and equipment focused.

Methodology

Objective 1 will be undertaken primarily through a literature review.

Objective 2 will be fulfilled through the findings of the literature review and the development of ideas to improve the system.

Objective 3 will be completed by undertaking practical validation trials.

Objective 4 will be completed through collaborative work with sport specific scientists and coaches.

Objective 5 will be a critical review process.

The end goal is to develop new positioning techniques and deliver impact to high performance sport by providing new information and measurements giving an enhanced insight into an athlete's performance.

0.3. Thesis Outline

This thesis is composed of two independent parts exploring precise location measurements for performance sports indoors and then outdoors. Each part is targeted at a particular application sport and will introduce the requirements and challenges coming with it. This introduction is then followed by the use of state-of-the-art positioning and surveying techniques to obtain and validate positioning data. The last focus of each part is on the development of sport specific metrics and measurements that can be used by coaches and sport performance analysts to improve the performance or their understanding of the performance. The key features of this work are:

Indoor tracking for wheelchair court sports

- (Chapter 2) Literature review of the state-of-the-art indoor positioning techniques that can be applicable.

- (Chapter 3) The research methodology gives an insight into the way trials were conducted to assess the performance of the system.

- (Chapter 4) Quality assessment and analysis of a commercial UWB system for indoor wheelchair court sport athlete tracking.

- (Chapter 5) Development of measurements and metrics specifically designed

for wheelchair court sports using the output of the UWB system.

Outdoor tracking for sprint canoeing and kayaking

- (Chapter 8) Review of existing solutions in the sport as well as a literature review of the state-of-the-art positioning techniques that can be applicable. The findings of the review are summarised in Section 8.3 giving an overview of the selected approach.

- (Chapter 9) In-depth review of the advanced optimisation technique identified.

- (Chapter 12) Contribution to an advanced GPS positioning technique and inertial navigation with the development of a novel tightly-coupled sensor fusion approach. Additionally, a specific focus is given on how the sport specificities can advantageously be used by the developed approach.

- (Chapter 13) Quality assessment of the developed solution both in a controlled and a sport environment.

- (Chapter 14) Development of measurements and metrics specifically designed for sprint canoeing and kayaking.

Part II.

Indoor Athlete Tracking for Wheelchair Court Sports

1. Introduction

This part of the PhD focuses on indoor precise positioning measurement for sports application. There is an interest in the use of technology to enhance the performance of athletes in wheelchair sports whether this is for improving equipment within sports or monitoring athletes performance. However, monitoring athletes during competition can be problematic. There are many challenges both in the collection and analysis of the data collected to fully understand where athletes can improve their performance. To add to the complexity of measuring athletes movements, team sports involve many players competing at one time on a pitch indoors or outdoors.

There are an increasing number of methods for tracking moving objects. Outdoors, Global Navigation Satellite System (GNSS) technology, normally the Global Positioning System (GPS), is often the simplest method to adopt [11] but indoors the situation becomes more complex as visibility of satellites for a GPS solution is not feasible or practical. There are however a number of systems available for indoor tracking [87]. These systems usually rely on different radio technologies such as WiFi [1, 2, 3], RFID [1, 2] or Ultra-Wide Band [2, 50, 134]. There are also a number of image based systems relying on infra-red [25] or traditional images [120]. This work focuses on the use of an Ultra-Wide Band (UWB) radio positioning system [122] which has been developed for the manufacturing, warehouse, and other industrial purposes. UWB systems are based on the use of fixed sensors around the region of movement where these sensors track the positions of tags which are fixed to the target object. The use of the UWB technology has been

dictated by the collaborative nature of this project with Loughborough University which acquired a commercial UWB tracking system just before the beginning of this PhD.

For every tracking application there is a requirement for a certain level of accuracy to ensure the location data collected is fit for purpose. As UWB radio positioning has not been designed for monitoring athletes, there is an even greater need to assess the capability of this technology to ensure the data collected and the derived information is valid.

The work presented below has been published in the Journal of Sports Engineering and Technology [102]. Thanks to the cleared permissions of the Journal publisher [110], part of the work presented is taken from the published article.

1.1. Aims

The general aim of this work is the quality assessment of an UWB system to track wheelchair court sports athletes. In more details, the first aim is the investigation of the system setup procedure and how it can affect the positioning quality obtained. The second aim is the validation and quality control of the UWB positioning system in the context of wheelchair court sports. Finally, it aims to convert the output of the system into measurements and metrics meaningful for the sport.

1.2. Objectives

As a result, the following objectives were identified:

1. Investigation of the system set-up procedure including (Section 4.1):

    - Set-up and configuration quality (Section 4.1.1)

    - Sensors spatial configuration (Section 4.1.2)

2. Validation and quality control of the positioning quality (Section 4.2):

    - Stationary positioning (Section 4.2.1)

    - Dynamic positioning (Section 4.2.2)

    - Impact of wheelchair environment (Section 4.2.3)

    - Filtered positioning quality analysis for distance measurement (Section 4.2.4)

    - Tag mounting location (Section 4.2.5)

3. Development of sport specific metrics converting the output of the UWB system into meaningful metrics for sport performance. (Section 5)

2. Indoor tracking technologies

2.1. Background

UWB is a short-range, large bandwidth radio technology. Its signal properties offer a strong multipath resistance and good penetrability in materials, which makes it particularly suitable for indoor environments. Additionally, the use of UWB pulses enables a very good time-domain resolution which allows such radio systems to be used for precise location and tracking [103, 111, 117, 122, 150].

Impulse-radio-based UWB systems are composed of sensors and tags. Sensors are receivers distributed around the area of interest, while tags are fixed on objects to be located and emit UWB pulses. The UWB pulse is received by a set of sensors which is used to compute a location based on Time Of Arrival (TOA), Time Difference Of Arrival (TDOA) or Angle Of Arrival (AOA) techniques [74, 104].

The TOA technique uses the time of flight of the UWB pulse to determine the distance between the tag and the sensor. For each sensor, this leads to a sphere of possible solutions. The position is then estimated by intersecting spheres of several sensors, a technique commonly known as multi-lateration [136]. However, TOA requires a perfect time synchronisation of all sensors and tags alongside a time-stamped UWB pulse so that the time of flight can be determined. These requirements are critical and generally not practical for commercial systems.

The TDOA technique is based on the TOA principle, but instead of computing the time of flight of an UWB pulse directly, it computes the difference in TOA between several sensors receiving the same UWB pulse. As this technique is time based, all sensors must be time synchronised. However, a benefit over TOA is that tags do not need to be synchronised.

If the orientation of the sensor is known, it is also possible to use the AOA technique to estimate the tag position. This is achieved for each sensor by determining the direction of arrival of the UWB pulse. Position can then be estimated by intersecting AOA of several sensors, a technique commonly known as multi-angulation [136]. All of these techniques are frequently combined together using non-linear regression or Kalman filtering to optimise the positioning quality [94].

2.2. Ubisense Real-Time Location System

This work focuses on the use of a commercial impulse-radio-based UWB system from Ubisense [134] (UWB system). The use of this particular system has been dictated by the collaborative nature of this project with Loughborough University which acquired the Ubisense real-time location system (RTLS) just before the beginning of this PhD.

The RTLS system is one implementation of UWB tracking among others. It operates on a $6 - 8$ Ghz bandwidth and additionally uses a 2.4 Ghz channel for telemetry purposes. The system is composed of sensors (Fig. 2.1a) and tags (Fig. 2.1b). Sensors are stationary and suitably distributed around the playing area to locate the positions of the tags worn by the athletes. Typically, 4 or 6 sensors would be used to surround an indoor wheelchair rugby court (28 m x 15 m).

All sensors are linked by an Ethernet cable to a master sensor for time synchro-



(a) Sensor (20cm x 14cm x 9.5cm) (b) Tag (20cm x 14cm x 9.5cm)

Figure 2.1.: Ubisense real-time location system

nisation. The tag location information is computed and displayed on a computer connected to the system. According to the RTLS specifications, an accuracy of $15 - 30$ cm with an update rate over 10 Hz can be expected. The operating range in open conditions is advertised to be up to 160 metres.

The system is transportable, so it can be easily deployed during competition or training at multiple venues. For this study, the system was installed and calibrated following Ubisenses recommended procedure. The first stage of the calibration is to determine sensor positions by measuring the distance between each sensor and two reference points. This is easily performed with a laser distance measurer. Each distance measured is given as an input to the Ubisense software which determines sensors position. The next step is to perform the cable offset correction to take into account the different length of wires used to allow for time synchronisation. The last step is to determine the orientation of each sensor. Several methods are proposed by Ubisense. The method used in this study is a dual calibration which identifies both cable offset correction and sensors orientation at the same time using a surveyed point, approximately in the middle of the area of interest, so its coordinates are known. A tag is left stationary at this point, while the system calibrates each sensor as the tag position is known [85]. Once this has been done, the system is set up and ready to use. A complete installation and calibration takes approximately 1 hour.

The UWB system only outputs tag positions, so other information such as speed

and distance travelled are derived from these coordinates. In addition, as the vertical component is irrelevant for wheelchair court sports, the work presented always consider the 2-dimensional horizontal positioning quality. In addition, a number of metrics have been developed to convert the position information into measures directly relating to the athlete's performance and are presented in Section 5. Very little was known about the performance and quality of the results achieved. A scientific investigation was proposed to critically assess the system in an ecologically environment.

3. Research Methodology

3.1. Trials and location facility

The aim of the trials discussed is the critical assessment of the Ubisense RTLS system. Trials were undertaken in a large sports hall with a viewing gallery for setting up the surveying equipment (Section 3.2) providing a good view of the playing area (Figure 3.1).

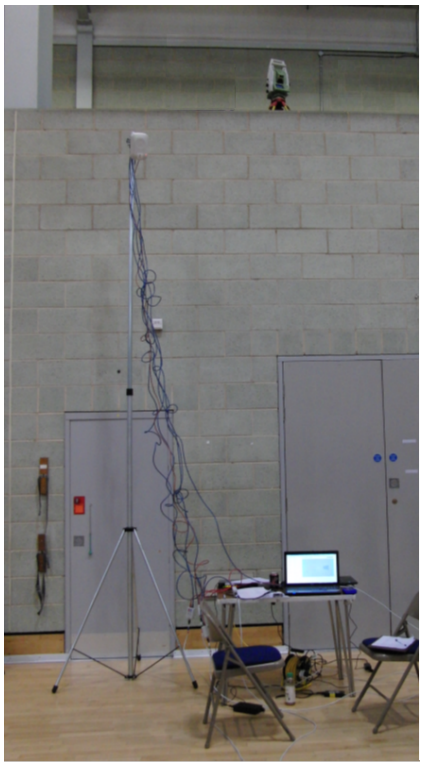


Figure 3.1.: Viewing gallery showing the Leica TS-30 surveying equipment and one UWB sensor on an elevated stand

3.2. Surveying Equipment

The Leica TS-30 is a robotic total station that allows classic surveying tasks (angle and distance measurement) and also tracking of a moving prism shown in Figure 3.3.

According to the Leica specifications [71], the tracking mode gives a positioning accuracy of 3 mm + 4 parts per million (ppm). With a 40 metres maximum range for the trials presented, the ppm part can be neglected. Timing specifications give a maximum measurement rate of 5Hz for the tracking mode.

When computing distance travelled from measured positions, the measurement frequency is a limitation. If the length of a trajectory is computed directly from the sum of distances between successive Leica TS-30 positions, it will be underestimated as straight lines will join the points. In order to correct this underestimation and knowing that Leica TS-30 measurements are precise in position, an interpolation is applied to more closely follow the track (Figure 3.2). The interpolation used is a cubic interpolation, so the interpolated trajectory effectively goes through all Leica TS-30 measured points as they are known to be precise in position.

As a conclusion, we can say that the Leica TS-30 can also be used as a gold standard for the distance assessment. However, this is true as long as the Leica TS-30 gives enough measurements for the interpolation to fill gaps. All Leica TS-30 traces used in this work have been checked to ensure no unacceptable gaps ($>$ 0.5 s) were present.

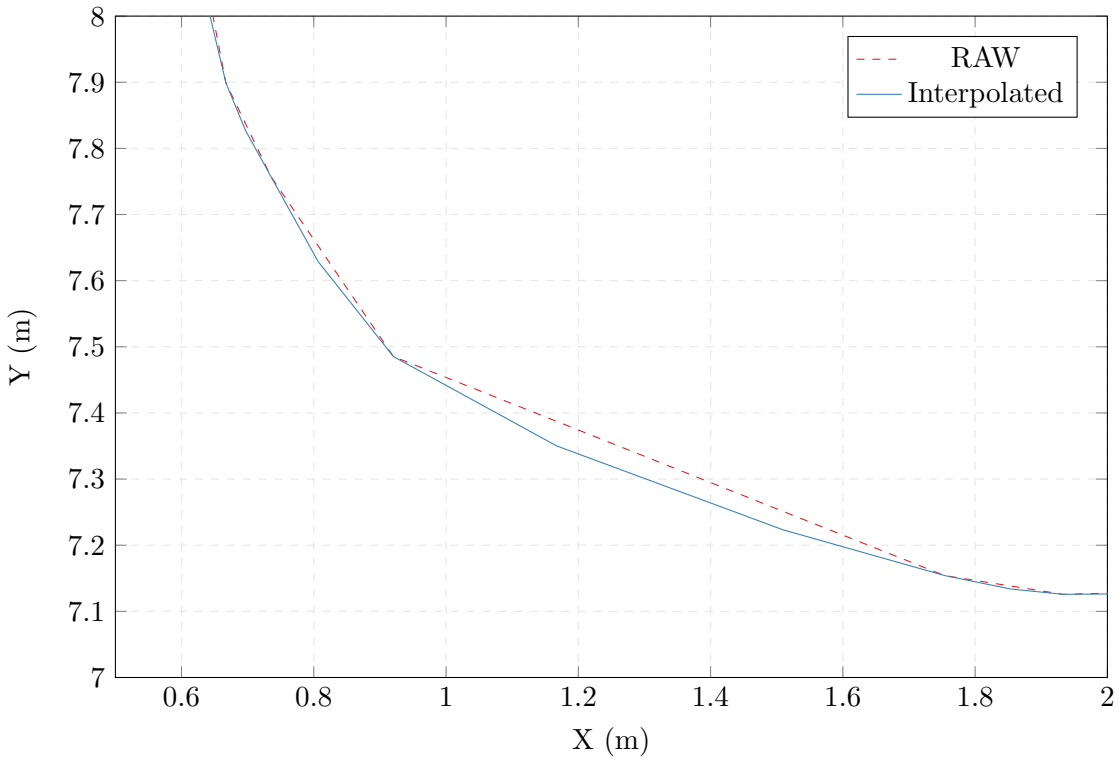


Figure 3.2.: Interpolation benefits to more closely follow the track

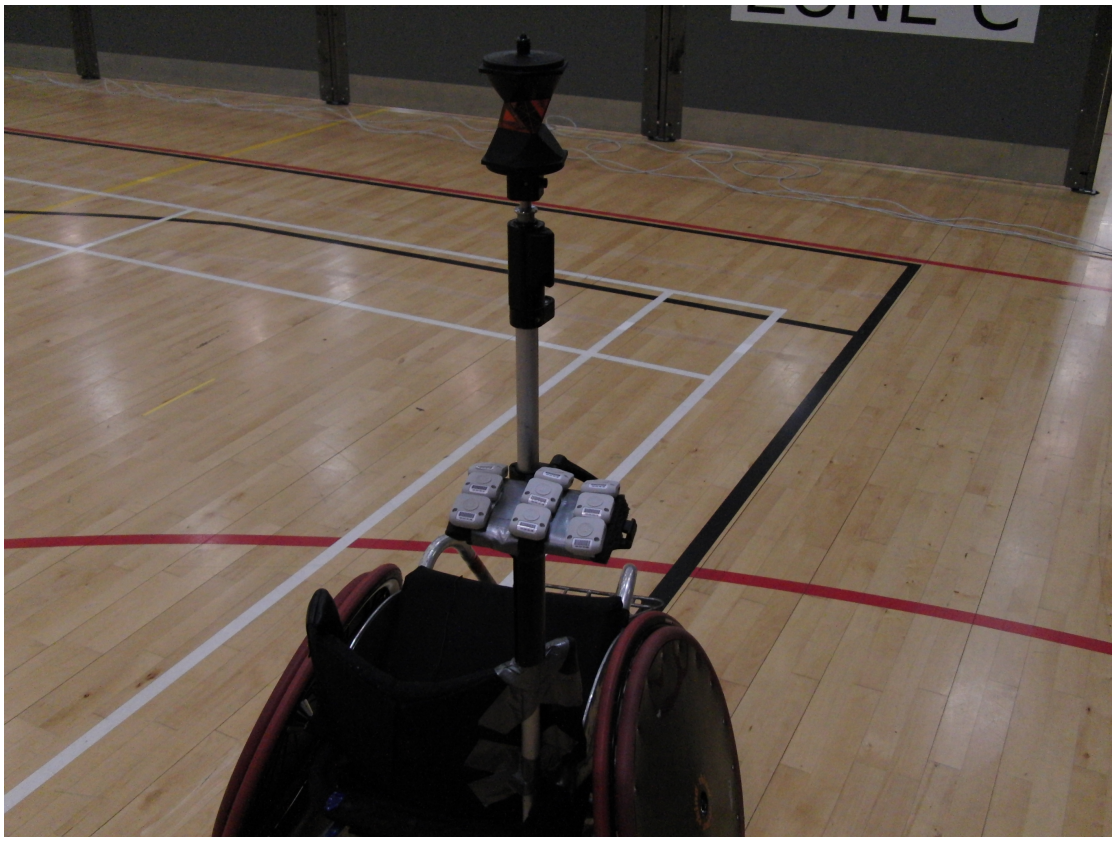


Figure 3.3.: Tags and prism mounted on a wheelchair

3.3. Data collection

As the aim was to assess the UWB positioning quality using the Leica TS-30 as a gold standard, it was necessary to mount tags and prism as close as possible. Another critical requirement with the Leica TS-30 is to maintain a line of sight between the total station and the prism for the tracking. So mounted on a wheelchair was a pole with the prism on top and a plate attached to accommodate the tags, see Figure 3.3. With this set-up, the actual horizontal position of the prism and tags is almost the same, the small offset being negligible compared with the expected system accuracy ($15 - 30$ cm). This set-up was very convenient as the wheelchair can be pushed normally keeping the line of sight between the Leica TS-30 and the prism. Another advantage of this set-up is that it was possible to reproduce wheelchair sports movements in an ecologically valid environment.

3.4. Data processing and smoothing

The processing workflow can be seen in Figure 3.4 and time synchronisation is required after positions are determined as both systems are using their own internal clock for time stamping. The time synchronisation is achieved based on Fast-Fourier Transform (FFT). Once the time synchronisation is complete, the common time window is determined to make sure both systems are running during the analysis period. From here, several movement parameters and quality indicators are computed.

The first quality measure is obtained by computing the horizontal position error of each UWB measurement compared to the interpolated Leica trace considered as the gold standard. The second quality indicator is based on the distance covered as it is one of the key metrics used to monitor sports performance. This was computed by summing up the distance between consecutive points for both Ubisense tags and

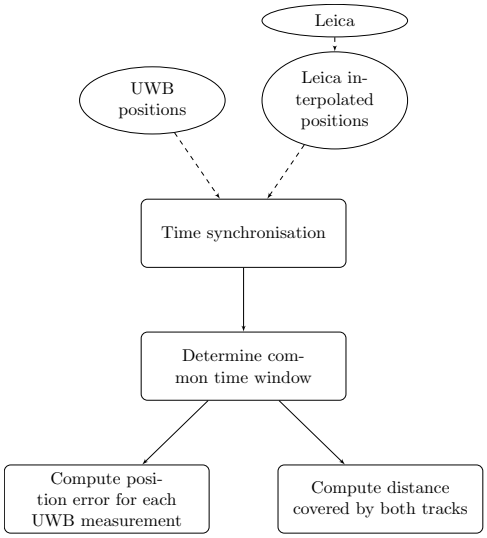


Figure 3.4.: Workflow of the processing used in this study

the Leica TS-30 interpolated data. However, as the UWB positions are subject to random noise, a filter was applied to mitigate this effect. The filtering used was a three-pass sliding average with a window size proportional to the tag update rate.

4. Trials, results and analysis

4.1. System set-up

4.1.1. Set-up and configuration quality

The first objective was an assessment of the set-up procedure recommended by Ubisense, which is a quick, low-cost and practical approach to the system installation. As mentioned in Section 2.2, the main step during this installation is to input sensors locations into the Ubisense software. Sensors locations were obtained by laser distance measurements relatively to two reference points, one at either end of the playing area.

In order to assess the precision of the laser measurement technique, sensors and reference points have been surveyed with the Leica TS-30. Comparing laser distances with Leica TS-30 equivalent distances gives a root-mean-square error of 4 cm which is a typical result according to the measurement technology used.

The distance measurements were provided to the Ubisense Location Engine Configuration (LEC) tool which computed estimates for sensors locations. The comparison between these estimates and the ground truth provided by the Leica TS-30 has been made for one particular set-up of five sensors (Table 4.1). This shows an expected result from the Ubisense set-up procedure. There is some evidence

Table 4.1.: Sensors position differences using Ubisense LEC tool

	X difference (m) Along court	Y difference (m) Across court	Z difference (m) Vertical up	Total difference (m)
Sensor 1	-0.224	-0.031	0.094	0.245
Sensor 2	-0.128	-0.058	0.001	0.141
Sensor 3	-0.069	-0.134	0.017	0.152
Sensor 4	-0.102	+0.050	0.037	0.119
Sensor 5	-0.165	-0.252	0.002	0.301

of small systematic bias (e.g. all negative X differences) and random measurement errors (e.g. the variation in the negative X values) which are of a typical magnitude for the system. More expensive, time-consuming and complex set-up procedures are possible [85, 94] but would not be as easy and quick to deploy for a mobile system used in a sports environment.

4.1.2. Sensors spatial configuration

One advantage of a flexible UWB system is that the sensors distribution can be adjusted to optimise the coverage of the area being tracked. Trials were undertaken in order to find the optimum sensors spatial distribution to cover an indoor wheelchair rugby court.

In order to assess the impact of sensor locations on the output of the system, the wheelchair has been used with both tags and prism mounted. The trajectory pattern used to maximise coverage of the playing area is shown in Figure 4.1. Positioning quality has been evaluated with a statistical analysis on the horizontal positioning error of the UWB system obtained as described in Section 3.4.

In order to optimise the court coverage, different configurations have been investigated using respectively four, five and six sensors. Using four sensors, one in each corner of the court, resulted in a mean error of 0.40 m ($\sigma = \pm 0.28$ m). Similar results were obtained when adding one sensor on one of the middle side with a

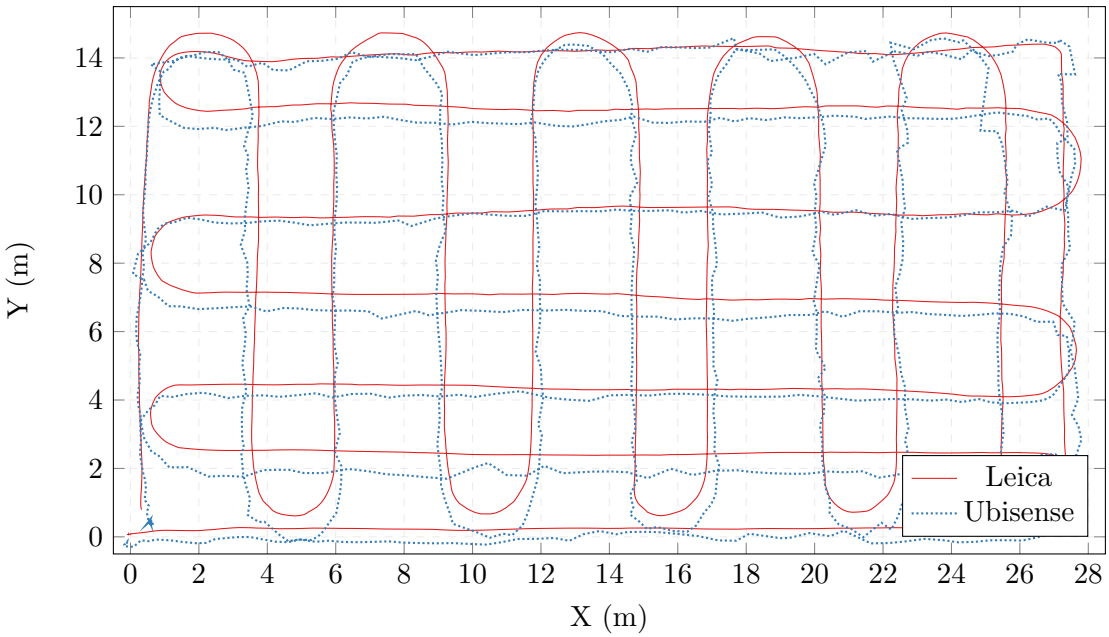


Figure 4.1.: Trajectory pattern used for sensors spatial configuration analysis

mean error of 0.39 m ($\sigma = \pm 0.29$ m). Finally, slightly better results (Ubisense track closer to the TS-30 track) were obtained using six sensors, one in each corner at 4 m height and two on the middle sides at 2 m providing a mean error of 0.35 m ($\sigma = \pm 0.23$ m). This is the optimum set-up tried and is the one used for all the trials presented next. A spatial analysis has been performed to identify possible areas with bad coverage. However, it appeared that the noise and random errors of the system did not allow for the identification of consistently weak areas.

4.2. Positioning quality analysis

4.2.1. Stationary positioning

In order to assess the stationary positioning quality, tags were left stationary for 5 minutes in known court locations. As an example, Figure 4.2 shows the output of one tag left stationary on a corner of the playing area and is a typical pattern of measurements from various positions around the court. The plot shows

that positions out of the UWB system are separated into two distinct clusters. This clustering can be explained by the noise due to sensors sets switching as described in Banerjee et al. [10] who also propose a particle filter to mitigate these effects.

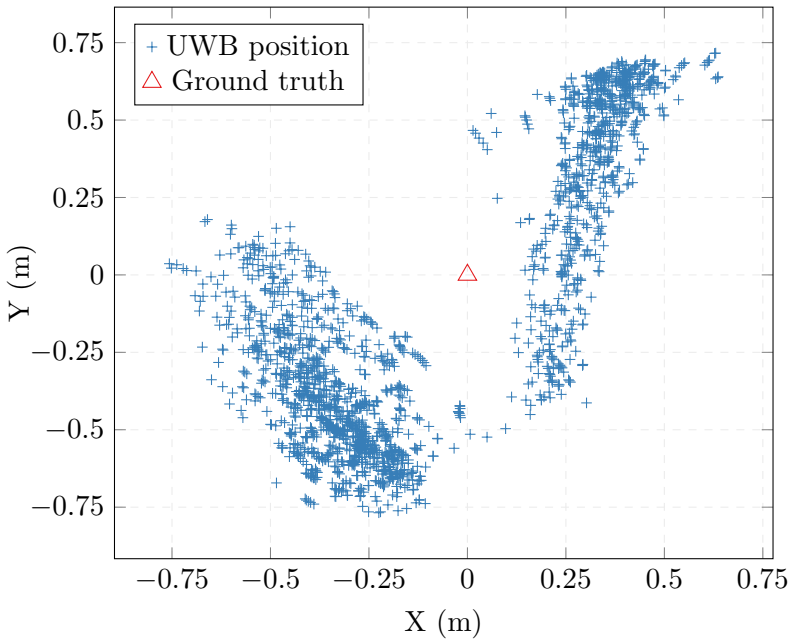


Figure 4.2.: Stationary tag positioning

4.2.2. Dynamic positioning

This section presents an important aspect of the work with the dynamic positioning quality assessment using a robotic total station with tracking capabilities. Previous studies were limited on dynamic assessment by performing only basic linear drills [112], differential comparison (tags comparison instead of a gold-standard comparison) [50] or by asking participants to follow a predefined path marked on the ground [112]. These have shortcomings for the validation of an UWB system in a sports performance context where athletes perform multi-directional movements at varying intensities. A more recent study by Leser et al. [72] addressed most of these shortcomings using a trundle wheel to obtain a distance reference. While this allows more freedom for the players, it still has some limitations in the dynamics that can be tracked and the measure may not reflect exactly the actual distance

covered by the athlete. Finally, as mentioned by Leser et al. [72], the trundle wheel provides a way to assess distance estimation but does not provide any information about the positioning quality. The protocol using a robotic total station addresses these shortcomings as detailed in the results presented below.

The first assessment evaluated the horizontal positioning error of each position output by the UWB system. In order to collect data relevant to indoor wheelchair court sports, trials were conducted during a simulated wheelchair rugby match with one participant playing four quarters of 8 minutes. The participant was asked to simulate a match play including turns, pivots, back and forth movements with rapid changes of speed. An example of a trajectory during a quarter is visible in Figure 4.3.

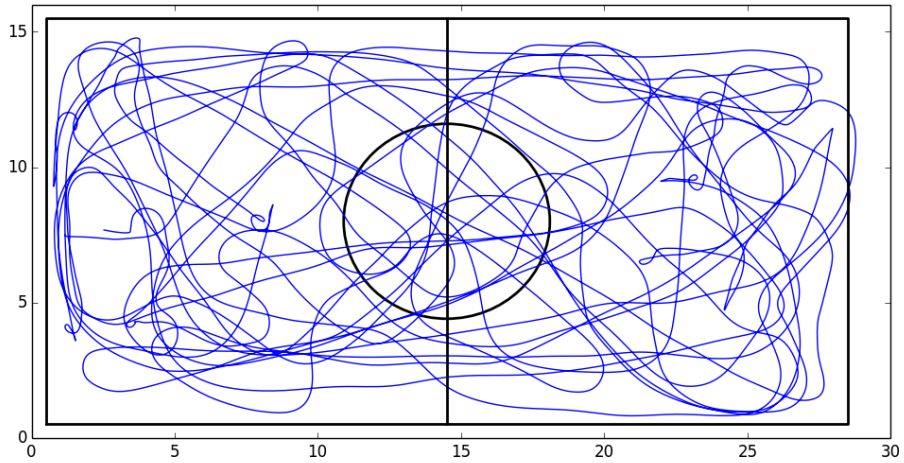


Figure 4.3.: Example trajectory during a quarter of a simulated match obtained with the Leica TS-30

Two matches of this format have been conducted. Table 4.2 presents the statistical analysis of the error for both two matches.

Results illustrated the accuracy of the system with a horizontal positioning mean error of 0.37 m and a standard deviation of ±0.24 m. Detailed results are presented to highlight the consistency of the system. Note that numbers obtained are similar to those obtained in the sensor spatial configuration trial (Section 4.1.2)

Table 4.2.: Mean (m) ± standard deviation (m) of the UWB positioning error during two simulated matches with nine tags operating at three different update rates

	Match1				Match2			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Tag 1 (16Hz)	0.33 ± 0.21	0.31 ± 0.21	0.36 ± 0.26	0.34 ± 0.22	0.39 ± 0.24	0.34 ± 0.22	0.41 ± 0.25	0.39 ± 0.27
Tag 2 (16Hz)	0.36 ± 0.23	0.31 ± 0.22	0.36 ± 0.25	0.35 ± 0.23	0.40 ± 0.22	0.34 ± 0.23	0.40 ± 0.24	0.39 ± 0.26
Tag 3 (16Hz)	0.36 ± 0.22	0.34 ± 0.21	0.39 ± 0.24	0.37 ± 0.22	0.40 ± 0.26	0.37 ± 0.21	0.41 ± 0.25	0.41 ± 0.25
Tag 4 (8Hz)	0.34 ± 0.21	0.33 ± 0.21	0.36 ± 0.25	0.36 ± 0.22	0.39 ± 0.23	0.33 ± 0.23	0.40 ± 0.24	0.39 ± 0.27
Tag 5 (8Hz)	0.36 ± 0.22	0.34 ± 0.21	0.38 ± 0.26	0.37 ± 0.22	0.41 ± 0.24	0.37 ± 0.24	0.41 ± 0.24	0.42 ± 0.25
Tag 6 (8Hz)	0.35 ± 0.21	0.34 ± 0.26	0.37 ± 0.25	0.36 ± 0.23	0.40 ± 0.24	0.37 ± 0.22	0.41 ± 0.25	0.41 ± 0.28
Tag 7 (4Hz)	0.35 ± 0.27	0.32 ± 0.22	0.38 ± 0.25	0.36 ± 0.24	0.39 ± 0.26	0.33 ± 0.23	0.40 ± 0.25	0.40 ± 0.26
Tag 8 (4Hz)	0.35 ± 0.23	0.32 ± 0.21	0.39 ± 0.25	0.36 ± 0.24	0.39 ± 0.26	0.34 ± 0.22	0.41 ± 0.26	0.40 ± 0.26
Tag 9 (4Hz)	0.36 ± 0.23	0.34 ± 0.22	0.41 ± 0.28	0.36 ± 0.24	0.38 ± 0.30	0.38 ± 0.25	0.42 ± 0.26	0.42 ± 0.26
Mean	(16 Hz) 0.37 ± 0.23 m		(8 Hz) 0.37 ± 0.24 m		(4 Hz) 0.37 ± 0.25 m		(Total) 0.37 ± 0.24 m	

where a mean error of 0.35 m and a standard deviation of ±0.23 m were found. Additionally, the tag update rate does not seem to affect the positioning quality according to the values grouped by update rate.

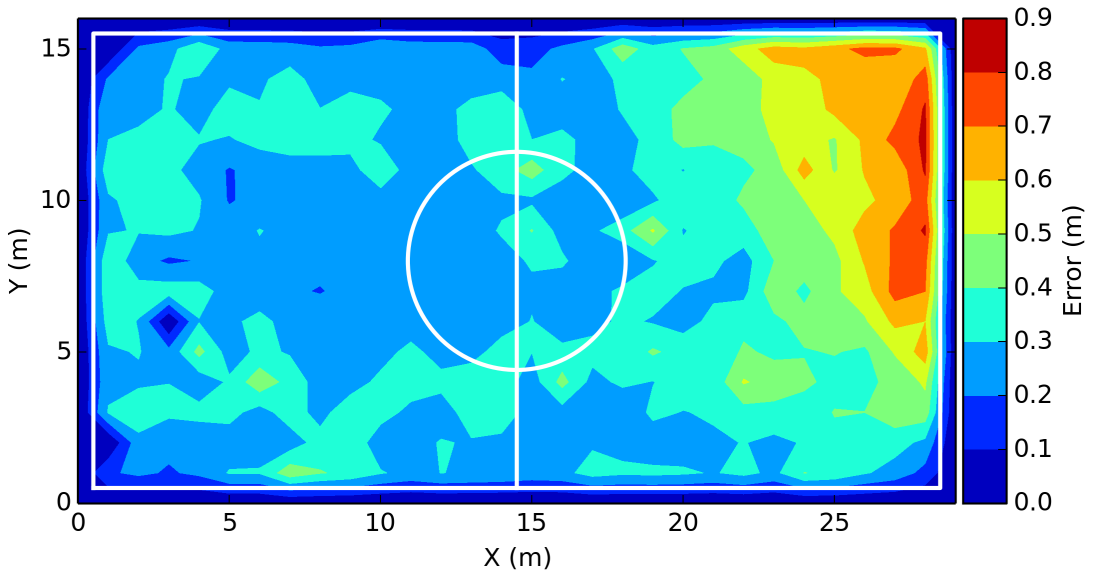


Figure 4.4.: Spatial distribution of positioning error for a particular configuration and setup of sensors

Figure 4.4 shows an example of the spatial distribution of the positioning error. This plot shows errors during one simulated match (4 x 8 minutes) concatenating data from 9 different tags. Each setup of the system will lead to a different error distribution. As can be seen on Figure 4.4, the positioning quality is lower on the top right corner of the court, which might be expected from an error on the

sensors position during the calibration of the system (Section 4.1.1).

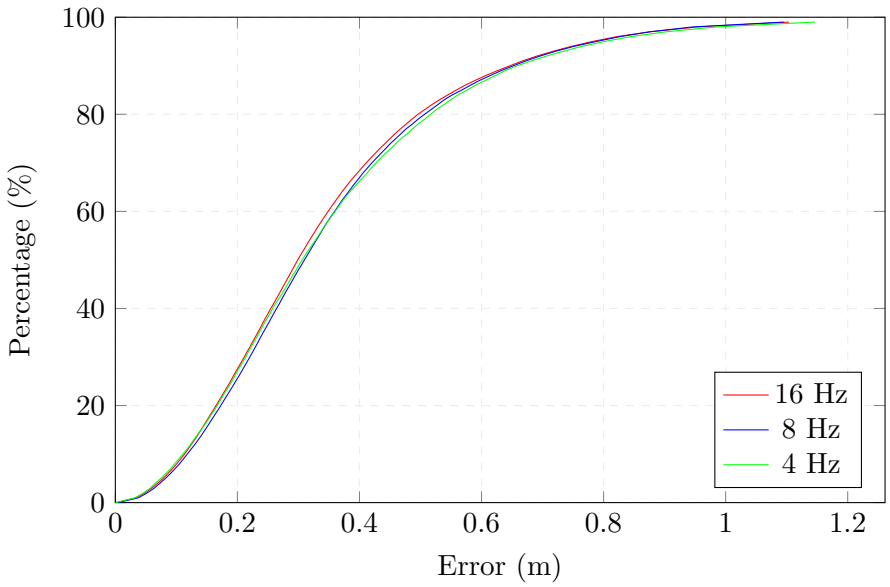


Figure 4.5.: Cumulative distribution function of the positioning error

Finally, the cumulative distribution function (CDF) of the positioning error has been computed and a typical example is shown in Figure 4.5. The CDF has been represented for the three different tag update rates used (4, 8 and 16 Hz) using the horizontal error on one simulated match (four quarters of 8 minutes). The closeness of curves in Figure 4.5 show that the tag update rate does not have a significant impact on the error distribution which confirms that the tag update rate does not affect the positioning quality as already seen in Table 4.2. The 90th-percentile positioning error is 0.63 m regardless of the tag rate. Similar results were obtained with an equivalent set-up procedure in Muthukrishnan and Hazas [94] where it is also compared to more complex system set-ups. The nominal update rates (4, 8 and 16 Hz) were checked against the time-stamped measurement records and agreed within less than 1 Hz. The findings of this section validate the capacity of the UWB system to track highly dynamic movements.

4.2.3. Impact of wheelchair environment

The results presented in previous sections were obtained with only one wheelchair moving on the court which is not representative of a wheelchair rugby environment. In order to address this, a match (four quarters of 8 minutes) has been simulated with another wheelchair interacting on the court. To simulate a game, this involved close engagement between the wheelchairs and more distant separation between them. The positioning quality was assessed using the statistical analysis on the horizontal positioning error described in Section 3.4. These trials were done using the set-up with five sensors described in Section 4.1.2 which reported a mean horizontal positioning error of 0.39 m ($\sigma = \pm 0.29$ m).

During the first two quarters, one wheelchair was tracked by both the UWB system and the Leica TS-30, while the second wheelchair remained untracked. The first and second quarters reported a mean error of 0.38 m (± 0.30 m) and 0.39 m (± 0.35 m), respectively. Such results demonstrate that perturbations caused by a second wheelchair on court do not influence the positioning quality. To determine whether the tags were working entirely independently, the second wheelchair was also tracked by the UWB system during the third and fourth quarters. Subsequently, error values remained similar, with a mean error of 0.38 m (± 0.26 m) and 0.36 m (± 0.27 m). As a result, tracking several independent objects in close and distant relationships did not affect the positioning quality.

Finally, some real match data have been investigated to evaluate the impact of having eight athletes and one referee on the court. Due to the need of a constant line of sight between the surveying equipment and the object tracked, no reference data were available. However, the data have been investigated regarding their availability (presence of data gaps) and visual correctness of the track.

As the real match data were obtained using a tag mounted on the foot strap, the presence of data gaps was compared to the results obtained in Section 4.2.5

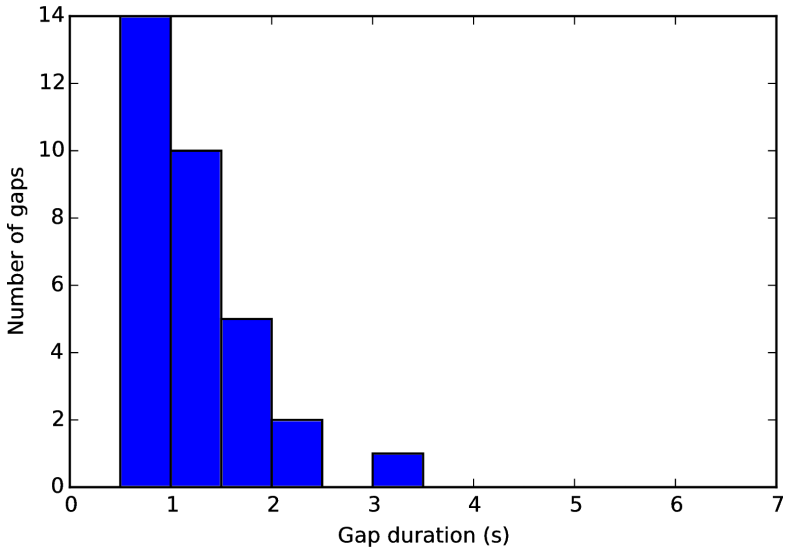


Figure 4.6.: Distribution of the data gaps with one wheelchair on the court

using a similar tag mounting location. A gap is defined as being a data outage for more than 0.5 seconds. The results from Section 4.2.5 are representative of a situation where only one wheelchair is on the court. Over the 488.1 seconds of the trial, 32 gaps were detected, giving an average occurrence of one gap every 15.25 seconds of an average duration of 1.20 seconds. The gaps distribution is visible in Figure 4.6.

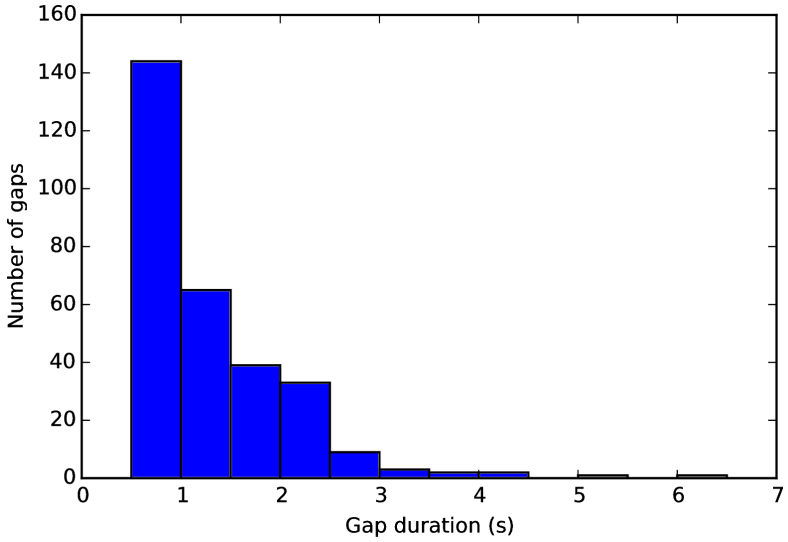


Figure 4.7.: Distribution of the data gaps during a real match with 8 athletes + 1 referee on the court

Moving to the real match data, the whole match was considered with the four quarters of 8 minutes including any stop in the game (ball out of play, fouls etc.)

representing a total dataset of approximately 70 minutes. Over the 4196.8 seconds of the dataset, 299 gaps were detected, giving an average occurrence of one gap every 14.0 seconds of an average duration of 1.28 seconds. Similarly, the gaps distribution is visible in Figure 4.7.

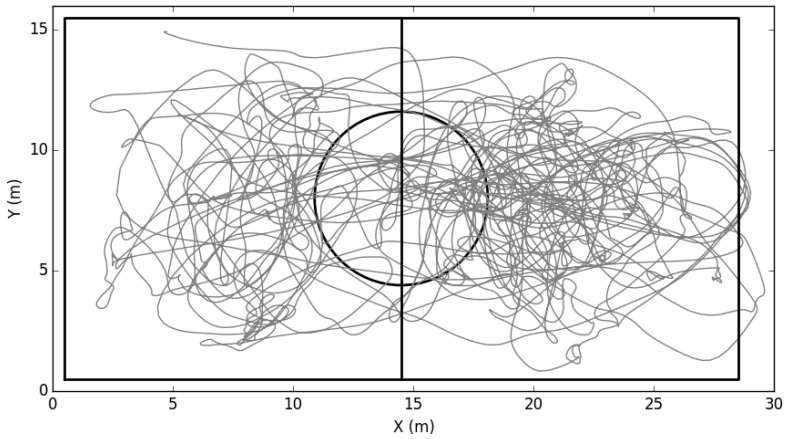


Figure 4.8.: Example of an athletes trajectory during one quarter of a real match

The second aspect of the data that has been investigated is the visual correctness of the track. This visual investigation also included a real-time replay of the track to detect any unnatural or irregular movement of the athlete. After investigation, no major anomalies or irregularities could be detected. An example of the track produced for a quarter during a real match is visible in Figure 4.8.

The findings of this section indicate that having a second wheelchair moving and tracked on the court does not affect the positioning quality. While it was not possible to use the surveying equipment during a real match due to the need for a constant line of sight, some real match data have been investigated. Two aspects were considered, the availability and regularity or smoothness of the track. The availability study was based on a direct comparison to the results obtained in Section 4.2.5. With only one athlete on the court, on average, a data gap occurred every 15.2 s for an average outage of 1.20 s. During a real match, with eight athletes and one referee on the court, the occurrence rate slightly increased to one gap every 14.0 s for an average outage of 1.28 s. The comparison of the respective gaps distribution (Figures 4.6 and 4.7) shows that larger gaps (> 2.5 s) are more likely to appear during a real match situation. However, this difference is not

significant since 94% of the gaps are below 2.5 s in the real match dataset compared to 97% in the single wheelchair situation. Additionally, the visual correctness of the track has been checked using a real-time replay and did not show any evidence of major anomalies or irregularities. As a conclusion, going from a single athlete on the court to a real match situation with eight athletes and one referee only slightly degrades the performance of the tracking system. However, according to the results of the analysis, the degradation is not significant and does not affect the suitability of the system to be used for indoor wheelchair court sports.

4.2.4. Filtered positioning quality analysis for distance measurement

This section focuses on the estimation of the distance covered by an athlete, an important metric in analysing athletes performance. A study by Leser et al. [72] assessed that the same UWB system can provide a distance estimate with an error of $3.45 \pm 1.99\%$ of the course length in a basketball context. These results were obtained using a combination of Kalman filter and low-pass filter. An alternative approach is proposed in this work.

In order to compute more accurately the distance covered, UWB measurements were processed using the 3-pass sliding average smoothing described in Section 3.4. Then distances are computed for both UWB smoothed measurements and interpolated Leica tracks and compared. Detailed results are presented in Table 4.3 which summarises the errors obtained for nine tags, operating at three different update rates, during the two simulated matches.

The mean error on the distance travelled for all tags of the UWB system is 0.45% of course length. Each quarter trajectory was approximately 1000 metres in length; therefore, a 0.45% error is equivalent to 4.5 metres. Additionally, results showed that the higher the update rate, the better the distance estimate. This can be

Table 4.3.: Distance error of nine tags operating at three different update rates during two simulated matches

	Match1				Match2			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Tag 1 (16Hz)	-0.05%	-0.40%	0.00%	0.04%	-0.55%	0.22%	-0.07%	-0.36%
Tag 2 (16Hz)	-0.62%	-1.05%	-0.25%	-0.56%	-0.90%	-0.35%	-0.14%	-0.36%
Tag 3 (16Hz)	-0.44%	-0.76%	-0.06%	-0.43%	-0.72%	0.13%	-0.08%	-0.02%
Tag 4 (8Hz)	-0.63%	-1.09%	-0.38%	-0.50%	-0.80%	-0.56%	-0.46%	-1.04%
Tag 5 (8Hz)	-0.40%	-0.97%	0.01%	-0.35%	-0.72%	0.00%	-0.04%	-0.17%
Tag 6 (8Hz)	0.27%	0.20%	0.49%	0.33%	-0.12%	0.78%	0.42%	0.63%
Tag 7 (4Hz)	-0.92%	-1.68%	-0.80%	-0.82%	-1.22%	-1.07%	-0.46%	-1.19%
Tag 8 (4Hz)	-0.05%	-0.56%	0.07%	0.31%	-0.54%	0.00%	0.22%	-0.12%
Tag 9 (4Hz)	-0.21%	-0.72%	-0.09%	-0.27%	-0.51%	0.49%	-0.01%	-0.11%
Mean	(16 Hz) 0.36%		(8 Hz) 0.47%		(4 Hz) 0.52%		(Total) 0.45%	

explained as the trajectory will be composed of more points which give a better recording of the dynamics with the same positioning quality (Section 4.2.2).

As a summary, using the proposed 3-pass sliding average smoothing approach, distance can be known with a mean error below 0.5% of course length. This is a significant improvement over the results presented by Leser et al. [72]. The difference is mainly explained by the different filtering applied. It may also be partly explained by the protocol used in Leser et al. [72] as the trundle wheel does not follow exactly the same path as the tag. Also, the tags update rate in Leser et al. [72] (4 Hz) may contribute to a reduced quality of distance measurement. The present work found that higher tag update rates ($\geqslant$ 8 Hz) are more suitable for distance estimation as outlier effects can be more easily mitigated by the processing and are also giving the finest recording of the dynamics.

4.2.5. Tag mounting location

Section 4.2.4 presented the results for distance estimation considering different update rates with tags attached as shown in Figure 3.3. In order to find mounting

Table 4.4.: Distance estimates and data gaps for different mounting locations

	Distance estimate	Error (%)	Data gaps
Prism pole	751.8 m	-0.13%	1
Vest	755.4 m	0.35%	6
Frame	750.1 m	-0.36%	31
Foot strap	745.9 m	-0.91%	32
Camber bar	736.9 m	-1.85%	34

locations more appropriate to wheelchair court sports, a specific trial was conducted with tags located in different places on the wheelchair or worn by the athlete. Below are the detailed tag mounting locations considered:

- Vest: Tag positioned between the scapula using a GPS vest worn by the participant.

- Frame: Tag attached to the wheelchair frame located at the front of the chair.

- Foot strap: Tag positioned onto the foot strap of the wheelchair.

- Camber bar: Tag secured to the camber bar of the chair located beneath the seat.

Additionally, one tag was left attached to the prism pole to allow for a direct comparison with the results presented in Section 4.2.4. A quarter of 8 minutes has been simulated and distance estimates have been computed for each tag. The Leica TS-30 has measured a distance of 752.81 m for this quarter. The results are summarised in Table 4.4 which also includes the number of data gaps (no measurement for more than 0.5 s) for each tag.

The tag attached to the prism pole provided the best distance estimate with results in agreement with those found in Section 4.2.4. Regarding mounting locations relevant to wheelchair sports, the vest appeared to show the smallest error. It

also appeared that lower mounting places are more subject to data gaps which could affect the distance estimate. Nevertheless, the magnitude of error was still minimal regardless of location, and lower places may offer the most practically relevant tag locations.

5. Sport specific metrics

To furthermore support wheelchair court sports, sport specific metrics have been developed. Whilst the distance covered estimation was the primary focus of the work, the use of an indoor tracking system permits the development of many other metrics. This chapter presents the different metrics that have been developed as part of this work. All the metrics have been integrated in a series of software to deliver a direct impact for elite sports (Section 5.4). The software were developed to specifically assist sport scientists and coaches to use the output of the UWB system. It has already delivered a great impact for the sport with example applications in [101, 105, 106, 107].

5.1. Speed and Speed Zones

As previously mentioned, the UWB system only delivers timestamped positions as an output. As a consequence, the speed measurement used in this section is obtained by differentiating position. Because the positions used have been previously smoothed, the velocity obtained is actually an average velocity over the smoothing window period. This has the advantage of mitigating the effects of measurement noise but does degrade the granularity of the data. Because the metrics developed in this section are all summarising the speed information, this does not represent an issue. However, if instantaneous velocity was required for a metric, a specific quality assessment should be conducted.

The initial speed metrics developed were as basic as the average speed and the maximal speed achieved by an athlete during a match or a training session. However, the data obtained with the UWB system has the potential to be used for much more refined speed metrics. In collaboration with Loughborough University, it has been decided to develop and investigate a set of speed zones covering the different intensities of an athlete's effort during a match [86]. The different speed zones are defined by a set of 4 thresholds which are also dependant on the athlete's class. The International Wheelchair Rugby Federation has established a functional classification system to accommodate players with different physical function based on their ability to perform a variety of movement. The different athlete's classes are 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5. The higher the class is, the more ability the player has. Table 5.1 summarises the different speed thresholds used for each athlete class.

Table 5.1.: Speed thresholds in m/s for the different athletes' classes

Theshold	Athlete's class			
	0.0 - 0.5	1.0 - 1.5	2.0 - 2.5	3.0 - 3.5
Z1	0.28	0.28	0.28	0.28
Z2	0.60	0.69	0.73	0.76
Z3	1.50	1.72	1.84	1.91
Z4	2.39	2.75	2.94	3.06
Z5	2.84	3.27	3.49	3.63

Based on these thresholds, the different speed zones are defined in Table 5.2.

Table 5.2.: Speed zones definition based on the thresholds

	Velocity
Static	$< Z1$
Low Speed	$Z1 - Z2$
Low-Mod Speed	$Z2 - Z3$
Moderate Speed	$Z3 - Z4$
High Speed	$Z4 - Z5$
Maximal Speed	$> Z5$

With the different speed zones being defined, it is possible to compute the time spent in each of them, both absolutely, with the time spent in seconds and relatively, with the percentage of the session or match spent in each zone. Similarly,

the distance covered in each speed zone is computed both absolutely, in metres, and relatively as a percentage of the total distance. Visual examples of the output are visible in Section 5.4.

In addition to the class defined speed zones, relative speed zones have been developed based on a percentage of the maximal speed achieved by the athlete. The default relative speed zones are defined in Table 5.3.

Table 5.3.: Relative speed zones definition based on the maximal speed achieved by the athlete

Zone 1	$< 20\%$ Max speed
Zone 2	$20 - 50\%$ Max speed
Zone 3	$50 - 75\%$ Max speed
Zone 4	$> 75\%$ Max speed

The relative speed zones have the advantage of being scaled to each athlete rather than their classification and as a result provide different insights. It can be viewed as the classification based speed zones gives indications on how well an athlete performs in his class whilst the relative speed zones gives indications on how well an athlete performed compared to his standards.

5.2. Accelerations and Sprints

Another area investigated is the detection of accelerations and sprints based on the UWB data. Accelerations have been divided into two classes, normal and extreme accelerations. A normal acceleration is defined as an increase in velocity of more than 2 m/s in less than 10 seconds. Similarly, an extreme acceleration is defined as an increase in velocity of more than 3 m/s in less than 10 seconds. These values have been empirically chosen based on the collaborative work with Loughborough University. One of the software, the analysis viewer (Section 5.4), offers the possibility for the user to manually adjust the different speed thresholds as well as the time period.

In addition to accelerations, sprints have also been investigated. A sprint occurs when the athlete's speed gets above the threshold value Z4 (Table 5.1) and finishes when the speed drops below the same threshold. For each sprint, the duration, distance, average speed and maximal speed are extracted as well as the time in the match it has occurred at. In addition, global statistics of the sprints over a full match are computed such as their average distance and duration.

In line with the velocity information extracted from the UWB system, the acceleration is obtained by double differentiation of the positions. This process increases the noise in the acceleration obtained compared to the positioning noise. As accelerations and sprints metrics described above are all summarising the acceleration over fairly long time windows, it is not critical. However, if instantaneous acceleration was of interest, a specific quality assessment should be conducted. From a technology perspective, the use of accelerometers is much more suited for this kind of applications.

Examples of the output for both accelerations and sprints are visible in Section 5.4.

5.3. Court occupation

The court occupation metrics developed in this section aims to provide ways of measuring where on the court an athlete has been spending its time during a training session or more importantly during a match. It is indeed crucible that a player respects its role in the team and covers the area of the court that he has been assigned.

As these metrics are a direct representation of spatial information they are best represented with a map form factor. The first one developed is a heatmap of the court occupation by a player. An example heatmap of a player during a match

can be seen on Figure 5.1 and 5.2. Heatmaps provide a visually intuitive way of viewing the court occupation. However, it does not allow for a precise and quantitative analysis as results are 'blurred' in a colour scale. To overcome this limitation, a second output has been developed with a map of the court divided in strategic zones representing the time spent (Figure 5.3) and the distance covered (Figure 5.4) in each zone. Both absolute and relative figures are given to allow for precise analysis of an event but also comparison between different events.

5.4. Analysis software

In order to deliver a direct impact to elite sports and to support ongoing research at Loughborough University, a range of software have been developed. The first goal was the completely automated generation of reports for an athlete or a team after a training session or a match summarising the results of the different metrics. This report can be exported and printed to make it easy accessible to anybody. An example output is visible on Figure 5.1. In addition to reporting, the developed software can also replay the session or match displaying the position on court as well as the velocity trace.

A second software has been developed to generate match reports for an athlete. An example output of the match report software is shown on Figure 5.2.

A third software has been developed to offer a more advanced analysis interface to be used by sport scientists at Loughborough University. This interface allows the manual adjustment of different thresholds and zones to explore and create new metrics meaningful for the sport. Figure 5.3 and 5.4 presents different outputs from the analysis software.

On a more technical aspect, other pieces of software were also developed to facilitate the use of the system in a sport environment. The first focus was on

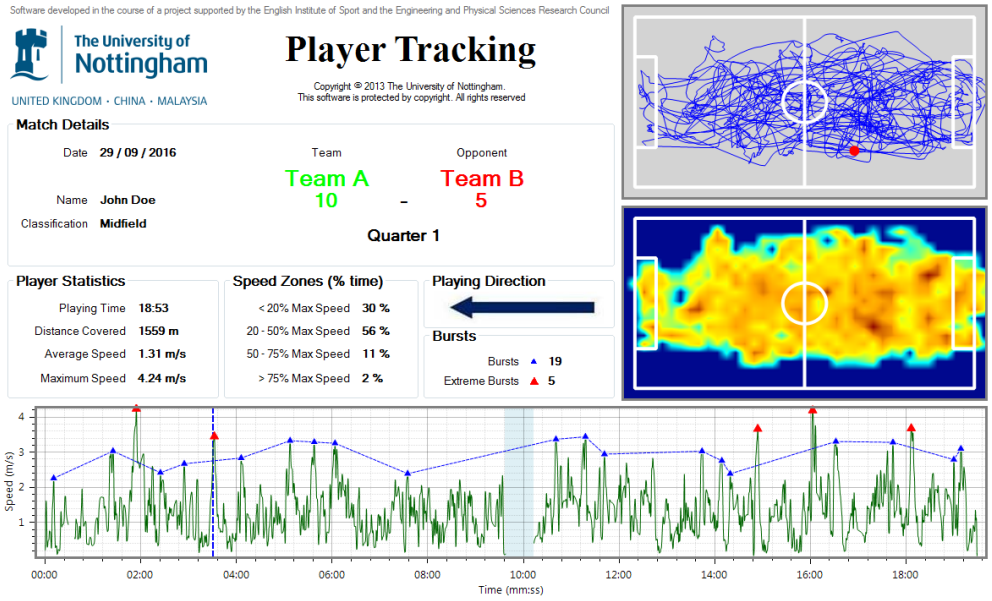
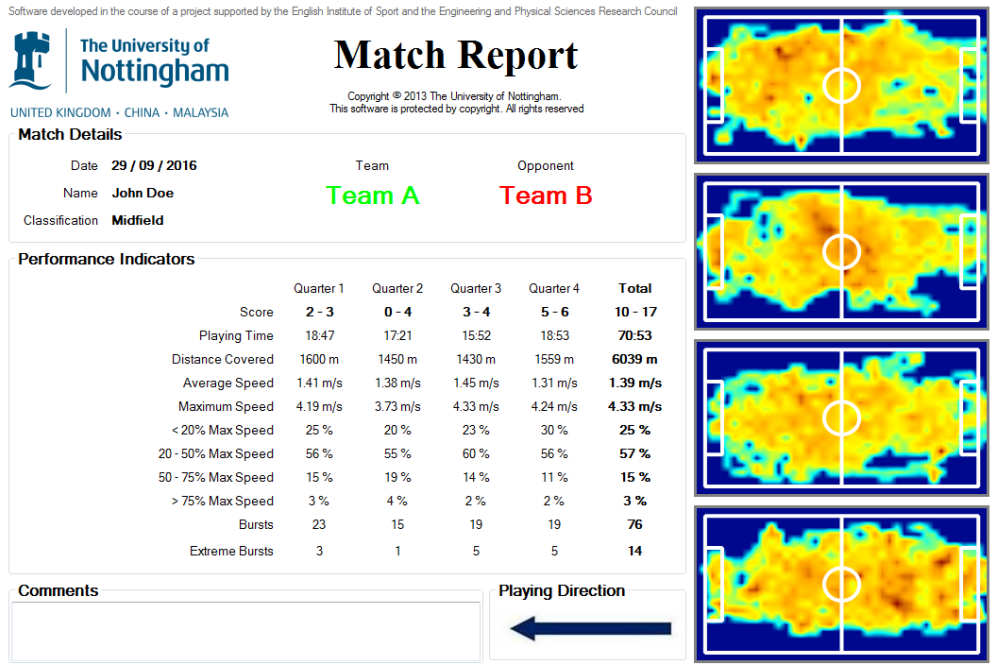


Figure 5.1.: Example output of the developed software offering a report of activity for athletes and coaches



	Quarter 1	Quarter 2	Quarter 3	Quarter 4	Total
Score	2 - 3	0 - 4	3 - 4	5 - 6	10 - 17
Playing Time	18:47	17:21	15:52	18:53	70:53
Distance Covered	1600 m	1450 m	1430 m	1559 m	6039 m
Average Speed	1.41 m/s	1.38 m/s	1.45 m/s	1.31 m/s	1.39 m/s
Maximum Speed	4.19 m/s	3.73 m/s	4.33 m/s	4.24 m/s	4.33 m/s
< 20% Max Speed	25 %	20 %	23 %	30 %	25 %
20 - 50% Max Speed	56 %	55 %	60 %	56 %	57 %
50 - 75% Max Speed	15 %	19 %	14 %	11 %	15 %
> 75% Max Speed	3 %	4 %	2 %	2 %	3 %
Bursts	23	15	19	19	76
Extreme Bursts	3	1	5	5	14

Figure 5.2.: Example output of the developed match report software offering a full match report for an athlete

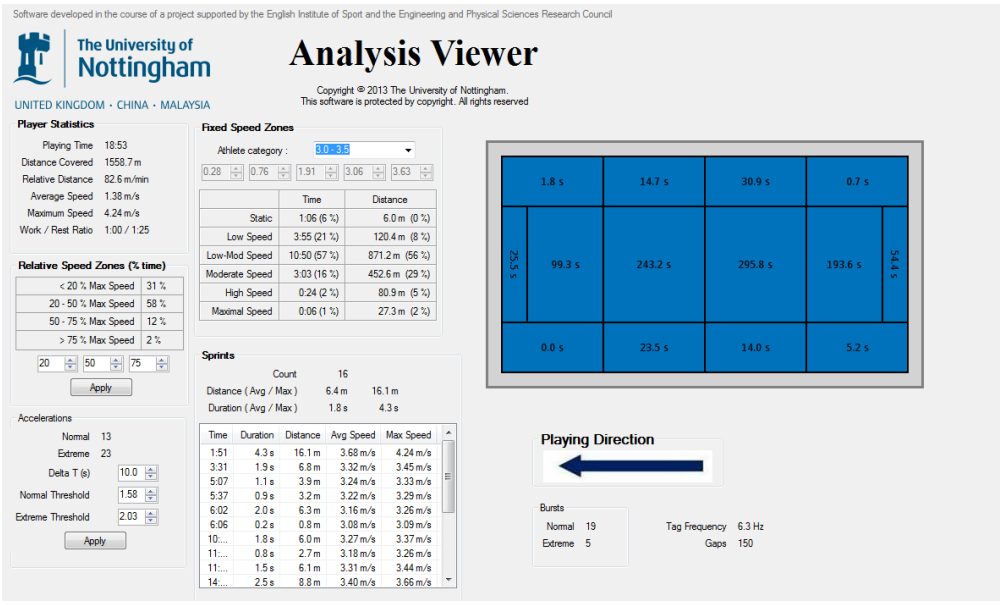


Figure 5.3.: Example output of the developed software for enhanced analysis displaying the time spent in different areas of the court

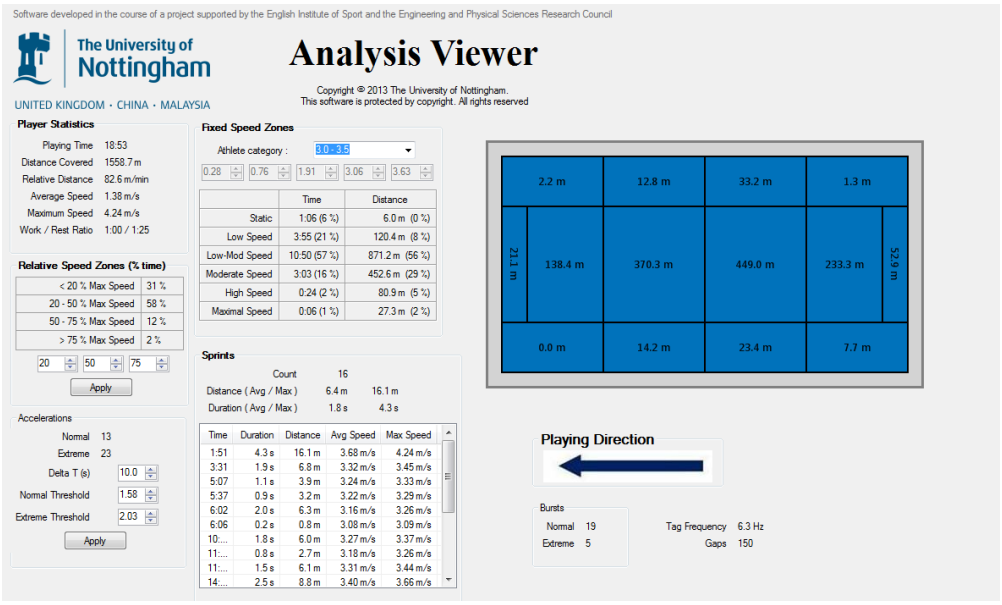


Figure 5.4.: Example output of the developed software for enhanced analysis displaying the distance covered in different areas of the court

supporting the set-up and calibration process while the second focus was to make players substitutions easier. While both tasks can be performed using the proprietary software accompanying the Ubisense system, it is not designed for the sport environment and thus these simple tasks proved to be very time consuming.

6. Conclusion

The use of a commercially available Ultra-Wide Band (UWB) positioning system to track wheelchair court sports players has been investigated. The literature review shown that an increasing number of studies use UWB systems for sport tracking applications. It also highlighted the need for a quality assessment to validate the capabilities of such systems. Existing studies in the literature did not evaluate the performance of UWB systems in a sport context against a gold-standard reference or only considered pre-defined patterns (Objective 1). The possibilities of the development of a novel positioning algorithm were limited as the commercial UWB system did not allow for internal modifications. However, a specific filtering has been proposed to significantly improve the computation of the distance covered by an athlete. (Objective 2). Through the use of state-of-the-art surveying techniques, a quality assessment of the UWB system has been performed in an ecologically valid environment with a particular emphasis on dynamic tracking. The challenges induced by the wheelchair environment have also been considered. It has been demonstrated that with a quick and easy deployment procedure, dynamic tracking can be achieved with a mean horizontal positioning error of 0.37 m ($\sigma = \pm 0.24$ m). Additionally, distance covered can be determined with an error below 0.5% of course length with adequate data processing. Tag update rates did not have a significant impact on the positioning quality; however, higher rates ($\geqslant 8$ Hz) provided a greater number of points to more closely record the high dynamic movements. It was also found that having many wheelchairs on the court did not have a significant effect on the positioning. Finally, several tag mounting places have been tried with the smallest error obtained for the tag

worn by the athlete in a GPS vest (Objective 3). The second contribution is the development of sport specific metrics converting the system's output into meaningful metrics for performance analysis (Objective 4). This work has supported a number of publications in the sport science community. It has also supported the use of the UWB system during international wheelchair rugby competitions as well as at Team GB training camps. Finally, a critical review of the research has been conducted including the applicability of the UWB system to other sports (Objective 5).

6.1. Contribution to knowledge

Contribution to knowledge has been made with the quality assessment of an UWB system for wheelchair court sports athlete tracking and published in [102]. Whilst the quality assessment was targeting a specific application, a number of validation trials are also valid in other contexts. In addition to the quality assessment, tools and metrics have been developed to convert the system's output into sport specific metrics and measurements. A number of studies [86, 101, 105, 106, 107] have been completed, all based on the UWB system assessed and the tools developed in this work. In addition to the clear contribution to the sport science community, the presented work also has a contribution in the positioning and sensing communities with references to this work in Alarifi et al. [5] and Eryildirim et al. [33].

6.2. Impact for elite sports

Besides the contribution to knowledge, impact for elite sports has been delivered with the development of specific tools to generate match reports and activity summaries of an athlete. The system and its tools have already been used during a number of international wheelchair rugby competitions as well as Team GB

training camps.

6.3. Discussion : Other Sports Applicability

Whilst the work presented had a strong emphasis on wheelchair rugby due to the collaborative work with Loughborough University and their interest in this specific sport, it can be extended to a range of other indoor court sports. First of all, as the impact of wheelchair environment has been assessed (Section 4.2.3), the use of the system can be extended to wheelchair basketball and wheelchair tennis. In addition, it can also be used to track fully-abled athletes in sports such as handball, netball or basketball. One limitation of the work presented is that it did not investigate the positioning quality on the vertical component as it is irrelevant to wheelchair court sports. As a result, if it was to be used, a specific investigation and quality assessment would be required.

Part III.

Outdoor Athlete Tracking for Sprint Canoeing & Kayaking

7. Introduction

This part of the PhD focuses on outdoor precise positioning and attitude measurement for sport tracking applications. To support and apply this research to an ecologically valid environment, it has been agreed with the English Institute of Sport to work with GB Canoeing for the sprint canoeing and kayaking disciplines. GB Canoeing has a primary interest in position tracking so that sprint times can be derived. In addition, the organisation has a growing interest for technologies and techniques that can be used to provide new insights on the sprinting dynamics for performance analysis.

7.1. Aims

The preliminary focus of this work is the identification of the sport requirements as well as a review of the existing commercial solutions and their limitations. Based on the findings, state-of-the-art positioning technologies and techniques that have the potential to meet the sport requirements and overcome the limitations of existing solutions will be identified. By proposing new developments around the identified techniques, this work aims to develop an innovative processing technique. In addition, it aims to propose specific developments to take advantage of the specificities of the sprint canoeing and kayaking application to further improve the quality and robustness of the solution obtained. The next aim is the validation and quality control of the developed technique. Finally, it aims to convert the

processing output into measurements and metrics meaningful for the sport.

7.2. Objectives

In order to fulfil the different aims described above, the following objectives were identified:

1. Identification of the sport requirements. (Section 7.3)

2. Review of the existing commercial solutions and of the current state-of-the-art positioning techniques. (Chapter 8)

3. Based on the state-of-the-art and the sport requirements, identify potentially suitable positioning technologies and techniques as well as an optimisation strategy to be used. (Section 8.3)

4. Propose new developments around an advanced GPS positioning technique, inertial navigation and the fusion of IMU and GPS measurements as well as a comprehensive review of it (Chapter 12)

5. Development of specific constraints taking advantage of the target application specificities. (Section 12.3)

6. Validation and quality control of the developed solution in a controlled environment. (Section 13.1)

7. Validation and quality control of the developed solution in the targeted application against a gold-standard. (Section 13.2)

8. Development of sport specific metrics converting the output of the developed

processing into meaningful measures for sport performance. (Chapter 14)

These objectives link to the overall PhD objectives as follow :

PhD Objective 1 : Section objective 1 - 2

PhD Objective 2 : Section objective 3 - 4 - 5

PhD Objective 3 : Section objective 6 - 7

PhD Objective 4 : Section objective 8

PhD Objective 5 : Section objective 5 - 6 - 7 - 8

7.3. Sport requirements, constraints and goals

Below is a list of GB Canoeing requirements relevant to the position tracking part :

- Typical boat velocity : $4 - 6$ m/s
- Typical sprint duration :
  - 200m : 35”
  - 500m : 1’36”
  - 1000m : 3’23”
- Timing precision : $< 1/10$ s at the very least

- Equivalent positioning precision : $< 40 - 60$ cm

Additional constraints :

- Low-cost
- Lightweight
- Real-time
- Easy to setup

Regarding the more general look for technologies and techniques that can be used to provide new insights on the sprinting dynamics for performance analysis, a similar list of requirements cannot be precisely defined. One of the main reason is the novelty of the subject. The existing literature typically uses video or instrumented paddles to obtain insights on the sprinting dynamics. A strong contender to complement these systems is the use of inertial tracking which can provide position, velocity and orientation of the boat at high-rate (typically 200 Hz). No evidences on the use of such an inertial tracking system have been found during the literature review. There are however studies using raw inertial measurements and especially the accelerometers output for stroke rate extraction and forward acceleration studies [28, 89, 137, 138]. GB Canoeing has a growing interest in having precise orientation measurements of the boat during strokes alongside the boat position and more importantly its velocity. Understanding the correlation between these different variables can give crucible information on key factors leading to performance.

8. Review of existing systems and technologies available in the sport

The first part of this chapter is a review of the existing commercial systems for sprint canoeing and kayaking tracking and their limitations. The second part is a review of the state-of-the-art positioning technologies and techniques that can be applicable to the sport. Finally, based on both investigations, the last part will briefly introduce the choices made in this research to develop a suitable solution.

8.1. Existing commercial systems

8.1.1. Catapult - MinimaxX

The Catapult MinimaxX is a widespread GPS tracking device used in many sports. It was historically one of the first GPS tracking device conceived for sport applications. Although the exact components are not specified, the system uses a low-cost GPS receiver and an IMU. The validity and reliability of this device for sprint kayaking applications has been evaluated in Janssen et al. [59]. The im-

portant finding of that study is that the device under-reported the kayak velocity by 0.14-0.19 m/s and acceleration by 1.37 m/s^2. The magnitude of these numbers cannot be neglected when applied to elite sport.

8.1.2. Igtimi - YatchBot Tracker

The second device is the YatchBot Tracker made by Igtimi, a New-Zealand company specialised in sailing boat tracking. The actual components inside the YatchBot Tracker device are unknown. From the positioning quality obtained, one could guess that a good GPS receiver is used in standalone C/A code[1] mode. The company's website claims a 3-metres positioning accuracy which is compatible with this assumption. The YatchBot Tracker additionally embeds an IMU. Based on a preliminary analysis of the noise and bias stability, it is likely to be a low-cost MEMS IMU sensor in the $\sim$£100 price range.

Regarding the output, the YatchBot Tracker provides a timestamped position in geographic coordinates. It also outputs the acceleration as measured by the IMU and an attitude estimate in Euler angles timestamped by the IMU internal clock.

8.1.3. McLaren Applied Technologies - Athena

The third device available, called Athena, has been developed for UK Sports and GB Canoeing by McLaren Applied Technologies. It is not strictly speaking a commercially available device as it has been specifically designed for the needs of GB Canoeing and is only accessible by them. However, as it offers an all-in-one operational box, like a commercial solution, it was included in this category.

[1]Standalone C/A code is the basic GPS positioning technique typically used in smartphones or satnavs

Similarly to the Igtimi device, there is very little that is known about what is inside the box. It uses a low-cost u-blox 6 GPS receiver configured in the solution mode[2] and does not have raw measurements capabilities. This implies that the GPS receiver directly outputs the position (and eventually Doppler-based velocity) without any means of post-processing the measurements. The exact model of the receiver in the u-blox range is not known. The Athena device also embeds an IMU which once again can not be precisely identified. However, analysis of the measurements indicates it is an ultra-low-cost sensor, similar to what can be found in smartphones, in the £10 price range.

Regarding the output, similarly to the Igtimi device, it provides a position in geographic coordinates and a filtered velocity estimate. It also outputs the acceleration as measured by the IMU and an estimate of the attitude in Euler angles timestamped by the IMU internal clock.

8.1.4. Vaaka - Kayak Cadence Sensor

The last device available is slightly different as it is a cadence sensor. It is attached at the centre of the paddle and outputs the stroke rate using the embedded IMU and a specifically developed algorithm. The Vaaka cadence sensor needs to be used with a compatible ANT+ or Bluetooth device as it does not store or display any information itself. If the connected device has a GPS, stroke distance is also computed in addition to the stroke rate. A review of the device is made in Croft et al. [28], however, there is no critical assessment of the quality and reliability of the measured stroke rate. It is actually not known what is considered as a stroke by the sensor, when it starts and when it is completed. In addition, the distance per stroke measurement, if available, is limited to the paired device GPS accuracy, usually a smartphone or a GPS watch, which is well below the stated requirements of the sport.

[2]In solution mode, a GPS receiver only outputs the estimated position but does not output the measurements that were used to obtain the estimate.

8.1.5. Limitations of the existing systems

As described, the existing systems would have limitations if they were to be applied to elite sport for in-depth performance analysis. GPS based tracking devices are all using standalone C/A code positioning which has a very limited precision (5 - 10 metres). In order to improve the solution, heavy filtering is applied, which has non-negligible consequences on the data obtained. This is particularly critical during the start phase where the filtering can completely modify the shape of the boat acceleration profile.

As part of the validation trials conducted in this work (Section 13.2), the Athena and Igtimi devices were used alongside a gold-standard reference positioning system. For the details and the methodology of the trials please refer to Section 13.2. In order to identify the limitation of these devices, their performance has been compared against the reference system during a 500 metres sprint. Figure 8.1 presents the horizontal positioning error in metres during the sprint.

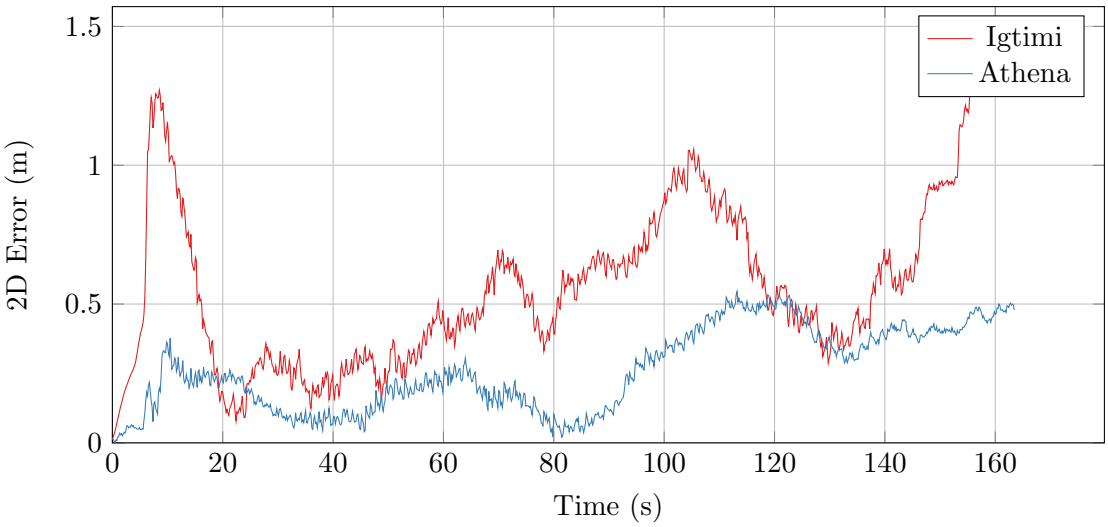


Figure 8.1.: Horizontal positioning error in metres during a 500 metres sprint

As expected, there is a limited horizontal positioning accuracy obtained by the standalone C/A code positioning used by the Athena and Igtimi devices. Coarse code solutions can only provide relative positioning with an error < 1 metre by

utilising built-in motion filters. Even with the filtering, solutions are still subject to random variations of a significant magnitude (> 50 cm). Interestingly, it can be seen on both Athena and Igtimi solutions that the error is rapidly growing during the first 10 seconds of the sprint before reducing again. This is a typical effect of the receivers' built-in motion filters which are trying to detect stationary periods to fix the receiver's position to improve the solution. A side effect of these filters is to introduce a lag when the motion changes. A summary of the results is presented on Table 8.1.

Table 8.1.: Mean and standard deviation in metres of the horizontal positioning error during a 500 metres sprint.

Device	Mean error (m)	Std (m)
Athena	0.26	0.15
Igtimi	0.58	0.32

A second aspect of the investigation is a review of the velocity measured by the device. Only the Athena device has been considered as the Igtimi one does not produce a velocity estimate. As a result, velocity would have to be obtained by differentiating the positions which would lead to a fairly noisy velocity estimate, especially knowing the positioning accuracy shown on Figure 8.1. For clarity, Figure 8.2 focuses on the beginning of the 500 metres sprint.

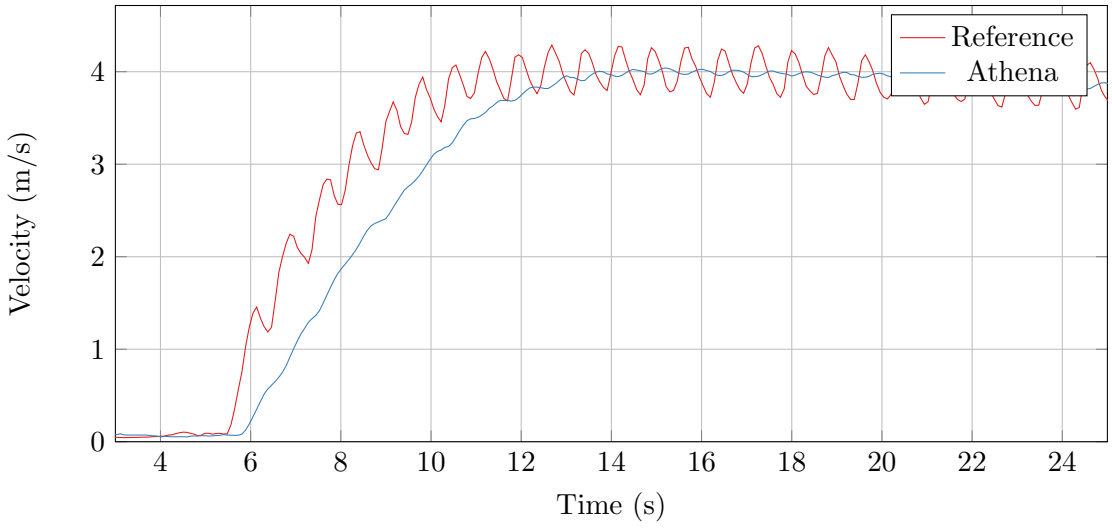


Figure 8.2.: Boat velocity during the first part of a 500 metres sprint

Once again, the Athena estimation is affected by a lag which is likely to be a side effect of the built-in motion filter. Additionally, it is clear that there is a different level of precision between the Athena and the reference. With the Athena, strokes can be distinguished but the change in velocity due to them is not valid and the output looks much more like an averaged velocity rather than an instantaneous velocity. This also stresses the need for quality assessment before using any of these device. The velocity output has the correct global shape and it would be easy to trust it if it was not plotted against a reference. This can have a dramatic impact for the sport if the data was used for in-depth performance analysis. One example would be the measurement of the change in velocity of each stroke which would be analysing and drawing conclusions on incorrect data.

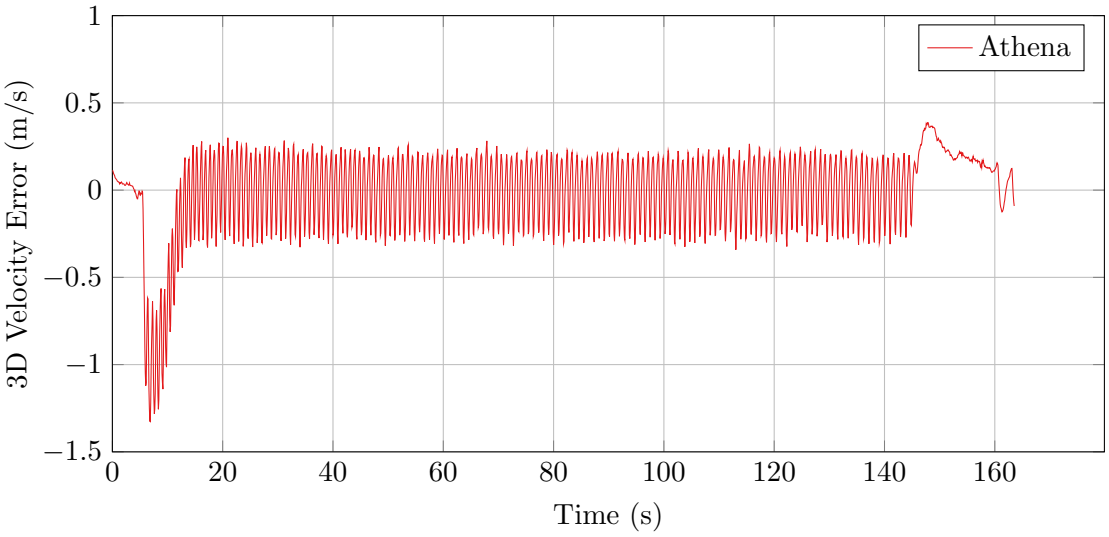


Figure 8.3.: Boat velocity error reported by the Athena device during a 500 metres sprint

The difference in level of precision and the initial lag are also visible on Figure 8.3 showing the 3D velocity error. The Athena reported velocity with a mean error $\mu = 0.18$ m/s and a standard deviation $\sigma = 0.16$ m/s. Interestingly, the numbers are comparable to the ones obtained in the review of the Catapult MinimaxX by Janssen et al. [59], which is thought to be relying on the same GPS technology, where it was found to under-report the kayak velocity by 0.14-0.19 m/s.

With these results, it is clear that existing devices have limitations. The results

also support the need for a quality assessment of the existing solutions and for research in the GPS area as stated by the community below:

> " *One major issue for the application of this information to coaches is that there is not enough evidence that commercially available products have an acceptable level of validity for measuring intra-stroke velocity, acceleration, or paddle force measures of elite paddlers.* " [89]

> " *More developments and research are needed in the area of GPS-based technology and instrumented paddle systems, as these devices will be useful for coaches and athletes in the future if reliability and validity measures can be improved.* " [89]

In fact, most of the research on sprint canoeing and kayaking do not use these commercial systems due to their limitations. As strokes are the critical component of a sprint, the majority of the studies are based on their detection and the extraction of related information. Using video is the first common alternative to detect strokes as in Mcdonnell et al. [88, 89]. However, it requires a time consuming post-processing task as strokes are manually detected frame by frame. In addition, the cited studies used videos with a very limited acquisition rate (24 Hz and 60 Hz respectively) which directly impacts the precision of the stroke detection that can be achieved. The second common alternative for stroke detection is the use of accelerometers as in Vadai et al [138] and Robinson et al. [108].

Whilst strokes can be detected using video or accelerometers, these systems on their own do not provide any positioning or velocity information. Because scientists want to know the velocity or distance covered by each stroke, they had to use the very little information that was available which left them with the following two solutions:

- Use the GPS standalone C/A code positioning which has a very limited

precision to derive velocity and distance from it [58]

- Use the total sprint time to derive an average boat velocity that is applied to all the strokes [88]

It is obvious that both solutions have a very limited precision or validity although they were the best available to the researchers. This work aims to investigate what positioning technologies and techniques could be used to overcome the limitations of existing solutions while integrating the requirements of the sport. The next section is a literature review of outdoor tracking technologies that can potentially be used for sprint canoeing and kayaking performance monitoring.

8.2. Outdoor tracking technologies

Based on the requirements of sprint canoeing and kayaking, a solution with a minimal infrastructure is key as athletes travel around the world for training camps and competitions. In addition, the area to cover can be as large as 1 km long. Having to transport and setup an infrastructure each time would be impractical if not infeasible. This rules out the use of systems such as the UWB system used in the first part of this thesis for example. It would also rule out any video based system as they require extensive calibration if measurements are to be made on them.

A strong contestant when it comes to outdoor positioning is the use of the Global Positioning System (GPS). The GPS technology has been around for almost 40 years now and has evolved in many different techniques and equipments each of them satisfying different accuracy and cost requirements. It has been seen in the previous section that all the existing commercial solutions are using standalone C/A code GPS positioning. This is the standard and basic GPS positioning technique. Whilst it has a lot of advantages in terms of robustness, size, power con-

sumption and availability, it has a limited accuracy (5 - 10 metres). This accuracy is perfectly acceptable when used for personal navigation or as an automotive navigation system but does not meet the requirements of elite sprint canoeing and kayaking. The absolute accuracy can be improved by using differential GPS (DGPS) where a static base station is used to obtain corrections which typically brings the accuracy down to 0.7 - 3 metres. However, the measurement noise would still be of significant magnitude (> 50 cm)

More advanced GPS positioning techniques exist, most of them use the carrier phase part of the GPS signal. Once again, different techniques with different equipments are available. The most common of these is the real-time kinematic (RTK) positioning technique which requires a static base station to be running at the same time as the rover attached to the athlete or object to track. The base station is used to send corrections to the rover in order to mitigate the effects of atmospheric propagation and other common GPS errors by measurement differentiation. A detailed review of the common GPS error sources is given in Chapter 10. This advanced GPS positioning technique can achieve performances in the range of a few centimetres and is commonly used in surveying applications. In terms of drawbacks, the technique requires the installation of a base station affecting the low-cost and easy to setup objectives.

The need for a base station can be overcome by using one of the commercial correction service recently created. It can take the form of a virtual base station like the Trimble VRS service [132] or the form of corrections sent by satellite over the L-band like the Trimble RTX service [131]. Both techniques would provide sufficient positioning accuracy for the sprint canoeing and kayaking application. One drawback is the need to pay for a subscription which increases the cost of the system.

Another alternative GPS positioning technique is know as Precise Point Positioning (PPP) which relies on precise satellite positions and clocks and a dual-

frequency GNSS receiver. PPP can provide position at a centimetre to decimetre level. One significant drawback, however, for any real-time application is the particularly long initialisation time and the need for integer ambiguity resolution which are both particularly critical with any sport application.

Finally, a completely different GPS processing approach has been identified with time-differenced carrier-phase processing allowing precise differential carrier phase positioning without the need for a base station or external corrections. This technique can be applied using only one, single frequency, low-cost GPS receiver. It also removes the need for ambiguity resolution and initialisation periods. This is achieved by differencing measurements of a single receiver at different epochs instead of differencing measurements of two receivers at the same epoch. While this technique has a lot of clear advantages, it also has some limitations, the main one being that it only achieves relative precise positioning. Additionally, because it does not use external correction, the positioning quality degrades over time.

From a more general perspective, a further drawback of any GPS-only solution is the limited measurement rate of GPS receivers. The common measurement rate is 1 Hz and can go up to 20 Hz with expensive specialised receivers. This is a limitation for any dynamic application such as sprint canoeing and kayaking where a stroke can be as short as 0.5 seconds and you want to obtain intra-stroke insights. A common way of addressing this issue is to combine GPS with an IMU to obtain an integrated solution known as inertial navigation solution.

Combining an IMU with a GPS receiver is a very popular choice as both systems have complimentary error characteristics. GPS provides a consistently accurate positioning quality over time as long as there is a line-of-sight to at least 4 GPS satellites. On the other hand, inertial navigation is an autonomous system that does not use any external signal. It can provide high accuracy short-term position, velocity and attitude estimates at a high data rate. However, inertial navigation accuracy degrades quickly over time mainly due to the inertial sensor's imperfec-

tions. Combining GPS and IMU can offer a greater reliability and accuracy by taking advantages of both systems. The GPS position is used to constrain the drift of the inertial system whilst the inertial system can be used to aid the detection of GPS outliers or to cover gaps due to a GPS outage.

In order to combine GPS and IMU measurements, a sensor fusion technique is required. The objective of sensor fusion is to find the optimal solution based on the measurements of multiple sensors. Because most sensors have non-linearities in their model, the underlying task is the optimisation of a large non-linear problem which is usually achieved, in the navigation community, by using Kalman filtering. A Kalman Filter provides an optimal solution if the problem solved is linear. As most navigation problems have non-linearities, a non-linear version, Extended Kalman Filters (EKF), has been developed. EKF are actually linearising the problem around the current estimate point. To allow high update rate and offer real time performance, EKF are aggregating all the past states and estimate only the current state. Because they are aggregating all previous states, EKF are unable to achieve a complete re-linearisation. As a result, augmented-state EKFs and fixed-lag smoothers have been developed to address this issue by storing some of the past states in the filter. Whilst the linearisation problem is solved, the computational cost of the filter is significantly increased as the state matrix size is extended.

Recent advances in mathematical and computer science domains have led to new techniques to solve large non-linear problems. The robotics community has largely applied them to solve the SLAM (Simultaneous Localisation And Mapping) problem by formulating SLAM as a maximum-a-posteriori estimation problem. In addition, the formalism of factor graphs [68, 76] is often used to describe the inter-dependence of the variables and as a result the underlying structure of the problem. The first advantage of using a factor graph to hold the structure of the problem is that it provides an insightful visualisation of it. The second advantage is its generality as it can be used to represent all kind of problems, from sim-

ple ones to more complex inference problems with multiple heterogeneous factors and arbitrary interconnections. As a direct consequence, it enables a class of well known generic solvers to be used to infer the problem. The final advantage is that it clearly exhibit the structure of the problem and in the case of a navigation problem, its sparsity.

In order to optimise the problem encoded by the factor graph, non-linear least squares are commonly used. An introduction to this technique is given in Section 9.3. A key part of it, is the computation of the Jacobian matrix which is required as part of the solving process. With large problems and matrices, obtaining the Jacobian is very computationally intensive and quickly becomes a limitation. Taking advantage of the sparsity of the navigation problem is the key insight of this recent technique. In more detail, the sparsity of the Jacobian matrix is dictated by the topology of the underlying factor graph. This sparsity can then advantageously be used by specialised fast linear solvers [57, 64, 69] relying on the latest advances in the Cholesky factorisation and its implementation [24]. As a result, these specialised solvers can solve problems with a sparse structure holding tens of thousands variables in a few seconds.

8.3. Overview of the selected approach

It has been shown that existing commercial devices have limitations, especially if applied to elite sport environments for in-depth performance analysis. All GPS tracking devices currently available are based on standalone C/A Code positioning which does not meet the sport requirements in terms of positioning accuracy.

Looking at the positioning requirements, it is clear that a code based solution will not reach the level of precision needed. Moving to a carrier-phase based solution enables such accuracy requirements to be achieved. However, such solutions usually require a base station to difference observations. This is not optimal as it

adds both to the complexity of setup and also violates the low-cost constraint. An alternative GPS processing technique has been identified by time differencing the carrier phase measurements of the receiver which removes the need for a base station [15, 128, 129]. The main limitation being that it only achieves relative precise positioning (Section 10.2). However, in the context of this PhD, the absolute accuracy rarely matters. Instead, in sports the main interest is to know precisely the evolution of an object or an athlete relative to its environment.

Because GPS receivers have a limited update rate ($1 - 10$ Hz), they are sometimes combined with an IMU. Not only this overcomes the limited GPS update rate, it also has the potential to deliver a fully integrated inertial navigation solution. Typically, position estimates as well as 3-axis velocity and orientation can be computed at high rate (200Hz) providing a lot of insights from the sport application perspective.

However, combining GPS and IMU measurements to obtain an inertial navigation solution is not straightforward. The task belongs to the sensor fusion domain and specific filters have to be used. While EKF have shown good performances, a novel approach is presented in Chapter 12 resulting in higher accuracy as it enables full re-linearisation of the problem. In the proposed approach, factor graphs are used to represent the sensor fusion problem. This representation gives an insightful visualisation of the problem and easily accommodates complex relationships between variables. It also provides a convenient way to develop new factors which has advantageously been used in this work.

The potential of the developed technique for its sport tracking application also needs to be investigated. The quality achieved by the developed solution is presented in Chapter 13 whilst Chapter 14 presents sport specific insights that can be obtained for performance analysis.

9. Graph-based problem optimisation

This chapter aim is to provide an overview of the graph-based problem optimisation technique which is gaining popularity, especially within the robotics community to solve the SLAM problem. SLAM is a challenging task, generally relying on the fusion of several sensors, most of them having non-linear models, over a long time span and ideally keeping all the past states in the optimisation process to identify common features and loop closures.

The technique can easily be split into two distinct parts. The first one is to model the problem to be optimised as a factor graph while the second part is to optimise the generated graph using non-linear least squares to find an optimal solution. This chapter will follow a similar separation after detailing the underlying probabilistic formulation in the first section. The second section will present how the probabilistic formulation can be represented as a graph. Finally, the last section will develop how the graph can be optimised to find an optimal solution using non-linear least squares.

9.1. Probabilistic formulation

Let's denote $\mathbf{x} = \{x_1, \ldots, x_n\}$ a set of continuous random variables and $p(\mathbf{x})$ its joint probability density function.

Using the chain rule of conditional probability, it is possible to write :

$$p(\mathbf{x}) = \prod_{i=1}^{n} p(x_i|x_1, \ldots, x_{i-1}) \tag{9.1}$$

Now, supposing that the random variables $\mathbf{x} = \{x_1, \ldots, x_n\}$ form a Markov chain in that order, it is possible to simplify the factorisation of the joint probability density function as :

$$p(\mathbf{x}) = \prod_{i=1}^{n} p(x_i|x_{i-1}) \tag{9.2}$$

The Markov chain property can also be interpreted as the variable x_i encodes all the relevant information up to index i.

In many cases, the variables are not directly observable but have a conditional dependency on one or several measurements. A measurement can affect one variable or a subset of them. Let's denote $\mathbf{z}$ a set of measurements. The joint probability density function given the set of measurements $\mathbf{z}$ then becomes :

$$p(\mathbf{x}|\mathbf{z}) = \prod_{k=1}^{K} p(\mathbf{x}|z_k) \tag{9.3}$$

However, it is rare for a measurement to have a conditional dependency on all the variables which means that for a specific measurement z_k, only a subset $\mathbf{x}_k$ of all the variables $\mathbf{x}$ are involved and Eq 9.3 can be written as:

$$p(\mathbf{x}|\mathbf{z}) = \prod_{k=1}^{K} p(\mathbf{x}_k|\mathbf{x}_{k_p}, z_k) \tag{9.4}$$

where, according to the Markov chain rule, $\mathbf{x}_{k_p}$ encodes all the previous states information considering the states $\mathbf{x}_k$ affected by the measurement z_k. From a

graph perspective, as presented on Figure 9.1, $\mathbf{x}_{k_p}$ is the set of variables that are direct neighbours to $\mathbf{x}_k$.

Let's now consider $\mathbf{x}$, a n-dimensional random variable affected by white Gaussian noise with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. By definition, the probability density function of $\mathbf{x}$ is:

$$p(\mathbf{x}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{\frac{n}{2}} |\boldsymbol{\Sigma}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right) \tag{9.5}$$

For each measurement z_k of $\mathbf{z}$ subject to noise w_k, it is possible to define a measurement model f_k such as :

$$\mathbf{x}_k = f_k(\mathbf{x}_{k_p}, z_k) + w_k \tag{9.6}$$

As it can be seen, the measurement model takes the previous states and the measurement itself as an input to predict the updated states as an output. Under the assumption that the noise w_k of the measurement model is a zero-mean Gaussian variable of covariance Σ, it is possible to write :

$$\phi_k = p(\mathbf{x}_k|\mathbf{x}_{k_p}, z_k) \propto \exp\left(-\frac{1}{2}(\mathbf{x}_k - f_k(\mathbf{x}_{k_p}, z_k))^\top \Sigma^{-1}(\mathbf{x}_k - f_k(\mathbf{x}_{k_p}, z_k))\right) \tag{9.7}$$

Which allows the definition of an error or cost function $\mathbf{e}_k$ as :

$$\mathbf{e}_k = \mathbf{e}_k(\mathbf{x}_k, \mathbf{x}_{k_p}, z_k) = \mathbf{x}_k - f_k(\mathbf{x}_{k_p}, z_k) \tag{9.8}$$

By substituting Eq 9.8 back in Eq 9.7, it becomes :

$$\phi_k \propto \exp\left(-\frac{1}{2}\, \mathbf{e}_k^\top\, \Omega_k\, \mathbf{e}_k\right) \tag{9.9}$$

where $\Omega_k = \Sigma_k^{-1}$, known as the information matrix, is the inverse of the measurement covariance matrix. And finally, Eq 9.4 can be written as :

$$p(\mathbf{x}|\mathbf{z}) = \prod_{k=1}^{K} \phi_k \propto \prod_{k=1}^{K} \exp\left(-\frac{1}{2}\, \mathbf{e}_k^\top\, \Omega_k\, \mathbf{e}_k\right) \tag{9.10}$$

As it can be seen, Eq 9.10 is a factorisation of the joint probability density function of $\mathbf{x}$ given a set of measurements $\mathbf{z}$.

The ultimate goal of the optimisation is to find the optimal set of parameters $\mathbf{x}^*$ that best explains the set of measurements $\mathbf{z}$. This is achieved by maximising the probability density function of $\mathbf{x}$ which is equivalent to minimising its *negative log-likelihood*, also named as *cost*:

$$\mathbf{F}(\mathbf{x}) = -\log p(\mathbf{x}|\mathbf{z}) \propto \sum_{k=1}^{K} \mathbf{F}_k = \sum_{k=1}^{K} \mathbf{e}_k^\top \, \Omega_k \, \mathbf{e}_k \tag{9.11}$$

which is done by solving the equation :

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \sum_{k=1}^{K} \mathbf{e}_k^\top \, \Omega_k \, \mathbf{e}_k \tag{9.12}$$

The equation 9.12 is the common form of a least squares estimation problem. The general solving of least squares estimation problems is detailed in Section 9.3.

As it has been seen in this section, each measurement z of $\mathbf{z}$ can also be viewed as an individual factor ϕ_k in the factorisation of the global probability density function (Eq 9.10)

9.2. Graph representation

A convenient way to represent the factorisation of Eq 9.10 is to use a graphical model known as factor graph. According to its definition [68, 76], a factor graph is a bipartite graph representing the factorisation of a function. A factor graph can be defined as $G = (\mathcal{X}, \mathcal{F}, \mathcal{E})$ where:

- $\mathcal{X}$ is the set of parameters vertices. Parameters are the variables of the problem to be estimated, sometimes also referred as states.

- $\mathcal{F}$ is the set of factor vertices. A factor encodes a constraint on or between a subset of parameters.

- $\mathcal{E}$ is the set of edges connecting each factor to the parameters it involves. Edges exist between factors and parameters only, and only if, the constraint encoded by the factor involves this parameter.

Thus, a factor graph describes the relations between state variables through factors encoding a measurement or constraint. One advantage of a factor graph is that it explicitly describes the factorisation of a global function and clearly exhibit the relations between the variables, also referred as the structure of the problem.

Each factor represents a local cost or error function. The set of all the factors of a graph represent a global cost function to be minimised. The goal of the optimisation is to find the optimal set of parameters that best explains the measurements and minimise the global cost function.

The initial theory of factor graph representation is described in Kschischang [68]. A good course introduction on factor graphs and how they can encode a probabilistic problem is given in Loeliger [76]. The link between factor graphs, non-linear least squares and Kalman filters is detailed in another publication of Loeliger [75]. Finally, the application of factor-graph based optimisation to navigation and especially to the SLAM problem has been presented more recently by Kummerle et al. [69] and Indelman et al. [55]. An example of a factor graph representing a navigation problem is shown on Figure 9.1 based on the work of Indelman et al. [55].

The represented problem uses multi-frequency measurements from low frequency GPS (red) and stereo vision (blue) measurements to high-frequency IMU measurements (black). It also includes some calibration parameters (green).

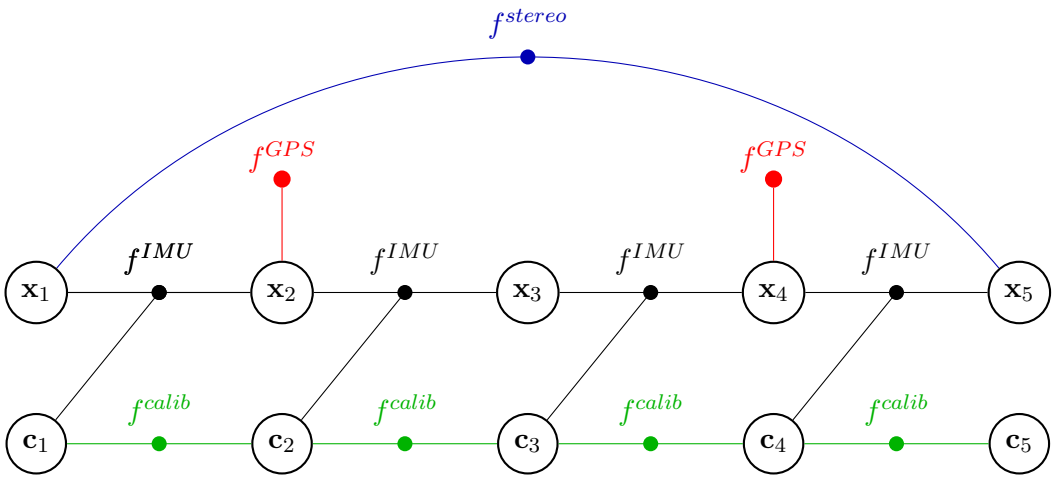


Figure 9.1.: Factor graph representation of a navigation problem

As part of this work, a novel approach on the use of GPS measurements and their factor graph representation has been developed in Section 12.1. Similarly, a novel formulation of the IMU factor has been proposed in Section 12.2.2 as well as a novel tightly-coupled formulation of the GPS and IMU sensor fusion in Section 12.2.1. Finally, new factors have been developed to take advantage of the sprint canoeing and kayaking application specificities in Section 12.3.

Now that the factor graph representation of a probabilistic problem has been defined, the next step is the use of non-linear least square to find an optimal solution to the problem encoded by the graph.

9.3. Non-linear Least Squares

As detailed in Section 9.2, it is possible to optimise a problem represented as a factor graph by solving a least square estimation problem. Indeed, it has been seen that the graph can be expressed as a finite sum of cost functions to be minimised.

Generally speaking, least squares problems are interested in finding the minimum

of a function $\mathbf{F} : \mathbb{R}^n \to \mathbb{R}$ such as :

$$\arg\min_{\mathbf{x}} \|\mathbf{F}(\mathbf{x})\|^2 \iff \arg\min_{\mathbf{x}} \frac{1}{2} \sum_{i=1}^{m} \|f_i(\mathbf{x})\|^2 \tag{9.13}$$

where $\mathbf{x} \in \mathbb{R}^n$ is a n-dimensional vector of variables and $\mathbf{F}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_m(\mathbf{x})]^\top$ is a m-dimensional function of $\mathbf{x}$ with $n < m$. $\mathbf{F}(\mathbf{x})$ represents a global cost function to be minimised composed of multiple cost functions f_i. The target of the optimisation is to minimise the global cost function $\mathbf{F}(\mathbf{x})$ which is also equivalent to finding the optimal set of parameters $\mathbf{x}^*$ that minimises $\mathbf{F}(\mathbf{x})$.

It can be seen that this formulation is directly applicable to the factors defined in Eq 9.11 and resulting in:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \frac{1}{2} \sum_{i=1}^{m} \|f_i(\mathbf{x})\|^2 = \arg\min_{\mathbf{x}} \frac{1}{2} \sum_{k=1}^{m} \mathbf{e}_k^\top \, \Omega_k \, \mathbf{e}_k \tag{9.14}$$

The solving of a non-linear least square problem typically involves an iterative process where parameters updates are computed at each iteration with the general aim of reducing the global cost using the following pattern:

1. Approximate $\mathbf{F}(\mathbf{x})$ around the current parameters estimate $\tilde{\mathbf{x}}$

2. Solve the resulting system to obtain a parameter update $\boldsymbol{\Delta}\mathbf{x}$

3. Update the parameter estimate with the computed step : $\tilde{\mathbf{x}} = \tilde{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}$

4. Iterate until a convergence criteria is met

A comprehensive review of the different existing methods to solve a non-linear least square problem is beyond the scope of this thesis. A brief introduction to the

Levenberg-Marquandt method, the one used in this work, is presented in the next section. The brief introduction is based on specialised literature [41, 77, 84].

9.3.1. Levenberg-Marquardt

Levenberg-Marquardt is one of the standard technique for solving non-linear least-squares problems and the one used in this work. From a technical perspective, it can be seen as a combination of the gradient descent and Gauss-Newton methods.

A gradient descent method reduces the sum of the squared errors by taking a step in the steepest-descent direction to update the parameters. A Gauss-Newton method reduces the sum of the squared errors by assuming that the cost function is locally quadratic at the current parameters estimate and find its minimum to obtain a parameters update.

The Levenberg-Marquardt method behave more like a gradient descent method when the parameters are far from they optimal value and more like a Gauss-Newton method when the parameters are closed to their optimal value. This behaviour change is controlled through a damping-parameter λ.

Let's consider our cost function:

$$\mathbf{F}(\mathbf{x}) = \frac{1}{2}\, \mathbf{e}(\mathbf{x})^\top \mathbf{\Omega}\, \mathbf{e}(\mathbf{x}) \tag{9.15}$$

Its 2nd-order Taylor expansion at the current parameters estimate $\tilde{\mathbf{x}}$ can be written as:

$$\mathbf{F}(\tilde{\mathbf{x}} + \Delta\mathbf{x}) = \mathbf{F}(\tilde{\mathbf{x}}) + \nabla\mathbf{F}\Delta\mathbf{x} + \frac{1}{2}\Delta\mathbf{x}^\top \mathbf{H}_\mathbf{F} \Delta\mathbf{x} \tag{9.16}$$

where $\nabla\mathbf{F}$ is the gradient vector composed of the first derivatives of F and $\mathbf{H}_\mathbf{F}$ is the Hessian matrix composed of its second derivatives.

Based on the literature [84, 119], a step with the gradient descent method would then be defined as :

$$\Delta \mathbf{x}_{GD} = -\nabla \mathbf{F}^{\top} \tag{9.17}$$

and a Gauss-Newton step as:

$$\Delta \mathbf{x}_{GN} = \mathbf{H}^{-1} \nabla \mathbf{F}^{\top} \tag{9.18}$$

The Levenberg-Marquardt step is defined as a combination of the two:

$$\Delta \mathbf{x}_{LM} = -\left(\mathbf{H} + \lambda \operatorname{diag}(\mathbf{H})\right)^{-1} \nabla \mathbf{F}^{\top} \tag{9.19}$$

The value of the cost is monitored in each iteration to update the values of λ. If a step is successful and reduces the cost, the value of λ is decreased for the next iteration, pushing the algorithm behaviour towards the Gauss-Newton method. On the other hand, if the step is not successful, resulting in an increased cost, the value of λ is increased for the next iteration, pushing the algorithm behaviour towards a gradient descent method.

A comprehensive review of the Levenberg-Marquart algorithm as well as the other existing methods to solve non-linear least-squares is available in Madsen et al. [84].

9.3.2. Local Parametrisations

The non-linear least squares optimisation methods such as Levenberg-Marquardt are assuming that the space of parameters is Euclidean. However, this is not always the case, particularly in the context of a navigation problem. The main example is when orientations are optimised as part of the problem. Typically, these

parameters are elements of the non-Euclidean rotation group $SO(2)$ or $SO(3)$. A common approach to incorporate such parameters in the problem is to express the parameters update $\Delta\mathbf{x}$ in a space different of the parameters themselves.

As an example, two main choices are available to represent a rotation:

- Use a minimal set of parameters (e.g. euler angles) to parametrise a non-Euclidean space resulting in singularities and non-linearities.

- Use an over-parametrised representation of the rotation (e.g. quaternions or rotation matrices).

It is generally preferable to use an over-parametrised representation to avoid singularities. Unfortunately, directly applying the update $\tilde{\mathbf{x}} = \tilde{\mathbf{x}} + \boldsymbol{\Delta}\mathbf{x}$ on the over-parametrised states is likely to violate the constraints of the over-parametrisation (unit quaternion or orthogonal matrices). A commonly used alternative is to compute perturbations and updates $\Delta\mathbf{x}$ using a minimal representation such as Euler angles whilst keeping the parameters $\mathbf{x}$ over-parametrised. Because perturbations and updates are usually very small at each step, they are far from the singularities and are thus safe to use. The parameters themselves are still benefiting from the robustness of the over-parametrisation. One critical requirement to apply this technique is the definition of a non-linear $\boxplus$ operator to apply the minimal representation update $\Delta\mathbf{x}$ to the over-parametrised parameters.

$$\tilde{\mathbf{x}} = \tilde{\mathbf{x}} \boxplus \boldsymbol{\Delta}\mathbf{x} \tag{9.20}$$

In this work, the orientations have been parametrised as quaternions and a local parametrisation has been used so that perturbations and updates are represented as Euler angles. The parametrisation as well as the definition of the non-linear $\boxplus$ operator used is the classical formulation as described in Agarwal et al. [4].

9.3.3. Loss Functions

For least-squares problems that are likely to contain outliers, the use of loss functions can greatly improve the quality of the minimisation. This is sometimes referred to as robust least squares. The objective of a loss function is to reduce the cost of large residuals and their influence on the final solution. A standard non-linear least square problem is easily transformed into a robust least squares problem with:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \frac{1}{2} \sum_{i=1}^{m} \rho_i \left(\| f_i (\mathbf{x}) \|^2 \right) \tag{9.21}$$

where ρ_i is the loss function.

A sane loss function should produce an output close to the squared cost for small residuals and have the following properties [4]:

$$\begin{aligned} \rho(0) &= 0 \\ \rho'(0) &= 1 \\ \rho'(s) &< 1 \text{ In the outlier region} \\ \rho''(s) &< 0 \text{ In the outlier region} \end{aligned}$$

In addition, the loss function should be scaled based on the specificities of the cost function f_i involved. This is basically defining which area or values are considered correct and which one are outliers and should be robustified.

Different loss functions exists, the one used in this work (Section 12.1.2) is a Cauchy loss function. The synthesis written by Triggs et al. [130] contains a more extensive review of the different loss functions and their use.

10. GPS Positioning

When it comes to outdoor positioning, the Global Positioning System (GPS), and more generally, Global Navigation Satellite Systems (GNSS) technologies are a reference. GPS provides instantaneous position and velocity estimates using measurements coming from a constellation of satellites that offers a global coverage of the globe. Whilst historically the technology is linked to military applications, it is now largely used beyond this scope. Outside of the military usage, the first logical field of application relates to all the surveying and mapping activities which are not only used by surveyors but also by geophysicists, geologists or archaeologists for example. The second large field of application is navigation which is available from maritime to space applications. Hence, GPS navigation has applications with automotive, boats, planes and more recently even spacecraft. It is also used by pedestrians, cyclists, hikers or runners either to obtain guidance or to monitor an activity.

This section firstly introduces the essential principles of GPS (Section 10.1). The second part (Section 10.2) will detail a particular GPS processing technique, namely time-differenced carrier-phase positioning, which has been identified and found particularly suitable for outdoor sport tracking.

10.1. GPS Fundamentals

The NAVigation System with Timing And Ranging (NAVSTAR) Global Positioning System (GPS) is an American satellite positioning system. It is a passive, 24 hour, all-weather system which provides instantaneous location and time information when direct line of sight to at least four GPS satellites is available. Making a full description of the system is a difficult and large task which is beyond the scope of this thesis and has already been accomplished excellently in many books [34, 52, 65]. This section, mainly based on the content of the books cited, aims to introduce the fundamentals while keeping it restrained to the notions that are directly involved in the GPS processing technique used in this work. For a more complete and truly comprehensive review of GPS, the reader is invited to refer to the books cited above for example.

10.1.1. System description

As of the time of this manuscript, 71 GPS satellites had been launched since 1978 with 31 currently operational. The GPS constellation, often referred to as the space segment, is managed by the U.S. Air Force which targets to maintain a constellation of at least 24 operational GPS satellites, 95% of the time. To comply with this commitment, 31 GPS satellites were maintained operational for the past few years. The constellation itself is composed of satellites from different generations with each generation benefiting from the advances in knowledge and technology to improve the quality of service, availability and integrity. All the GPS satellites are spread between six medium-Earth orbits with an orbital inclination of 55°, at an altitude of about 20200 kilometres and with an orbital period of roughly 12 hours. Historically, GPS satellites broadcast two L-band navigation signals at 1575.42 Mhz (L1) and 1227.60 Mhz (L2).

The second part of the system is known as the user segment. Due to the mili-

tary history of GPS, the user segment is divided into a civil and military usage. Civil users can access the standard positioning service (SPS) through the use of coarse acquisition (C/A) code combined with a navigation message which are both broadcasted on the L1 frequency. As all the satellites are broadcasting on the same carrier, the C/A code is modulated by binary phase shift keying onto the L1 carrier with a Pseudo-Random Noise, referred to as PRN code. Each satellite has a unique PRN code which allows the separation of the signals and the identification of the satellite involved after demodulation. This technique, known as spread-spectrum, improves the resistance to natural interferences and voluntary jamming. The modulated C/A code has a wavelength of approximately 300 metres.

In addition, military users have access to the Precise Positioning Service (PPS) through the use of a Precise P-code on both L1 and L2 frequency. In order to restrict its access to authorised users only, the P-code is encrypted by a Y-code and has a wavelength of 30 metres.

Besides the C/A-code and the P-code, a navigation message is continuously transmitted, containing the satellite ephemeris, satellite clock and ionospheric modelling coefficients which are necessary to compute a position from the measurements as well as satellite health information to monitor the system integrity.

More recently, following the GPS modernisation campaign, an additional civil signal (L2C) has been added on the L2 frequency for non safety critical applications. It appeared with the launch of the first IIR-M satellite in 2005. More lately, a new signal on the L5 band (1176.45 Mhz) has been introduced with the first Block IIF satellite launch in 2010. The L5 signal is the third civil GPS signal and it aims at supporting safety-of-life applications for aviation as well as an improved availability and accuracy. Finally, a fourth civil signal will appear with the first launches of Block III satellites.

GPS also relies on a network of worldwide ground stations, also known as the ground segment. These stations monitor and control the satellites as well as determining their orbit and uploading the navigation messages.

As a final note, the American NAVSTAR GPS system is the most famous and most used global navigation satellite system, however, it is not the only one available. There is currently another fully functional global navigation satellite system, maintained by Russia and named GLONASS. China is expending its regional navigation satellite system, BeiDou, which is hoped to become the global Compass navigation system by 2020. The European Union is also in the process of deploying its own system, known as Galileo, which is not expected to be fully operational before 2020.

10.1.2. GPS Observables

After the rapid overview of the system, this section describes the three fundamental types of GPS observations namely, pseudorange, carrier phase and Doppler observations. It will additionally cover the signal strength observable which does not provide any location information but contains a useful indicator of the measurement validity. These observations are all accessible with low-cost, single-frequency (L1) receivers. Slightly more expensive, there also exist civil dual frequency receivers, in which case, carrier phase, Doppler and signal strength values are also available for the L2 frequency.

Pseudorange Observation

The pseudorange is directly derived from the C/A code by matching a receiver's internally generated code with the received code from a GPS satellite. It is a direct time measurement of the signal propagation from a satellite to the receiver's

antenna which can be converted into a distance measurement after multiplication by the speed of light. Let's denote ρ the true range between a satellite s and a receiver r. The true range ρ is defined as :

$$\rho = \sqrt{(x_s - x_r)^2 + (y_s - y_r)^2 + (z_s - z_r)^2} \tag{10.1}$$

where x, y and z are the satellite and receiver Earth-Centered Earth-Fixed (ECEF) Cartesian coordinates.

The so-called pseudorange measured by the receiver is affected by errors of different sources. Section 10.1.3 details the origin as well as the impact of the different error sources in more details, however, a common expression of the measured pseudorange p in terms of these error sources is

$$p = \rho + c\ (\delta^S - \delta^R) + E + T - I + m + \varepsilon \tag{10.2}$$

where:

ρ	is the true receiver to satellite range (m)
c	is the speed of light (m/s)
δ^S	is the satellite clock error (s)
δ^R	is the receiver clock error (s)
E	is the ephemeris error (m)
T	is the tropospheric delay (m)
I	is the ionospheric delay (m)
m	is the multipath error (m)
ε	is the measurement noise (m)

Because receivers don't have a highly precise clock, there will always be an error, δ^R, between GPS time and the receiver's estimated GPS time. As a one millisecond difference equals to a range error of about 300 km, the receiver clock error cannot be neglected and has to be estimated as part of the positioning process. This is also why the measurement is called pseudorange since it is a range measurement

up to a time uncertainty.

Carrier Phase Observation

The carrier phase was not initially designed for positioning purposes. Its was only used to "carry" and transmit the C/A code message. However, different techniques have emerged to make use of this carrier signal for positioning purposes.

According to the reference literature by Hoffman et al. [52] and Traugott [129], an ideal phase measurement of the satellite S at time t by the receiver R with a first lock obtained at t_0 can be represented with the following mathematical model :

$$\varphi^S(t) = f^S t - \frac{\rho(t)}{c} f^S - \varphi_0^S(t)\varphi^R(t) \quad = f^R t - \varphi_0^R(t) \tag{10.3}$$

Where φ^S is the phase of the received carrier with a frequency f^S and φ_R is the phase of the receiver-generated carrier with a frequency f_R. Additionally, φ_0^S and φ_0^R are the initial phase biases caused respectively by the satellite and receiver clock offset compared to GPS time.

$$\varphi_0^S(t) = -\delta^S(t) f^S + \varphi^S(t_0)\varphi_0^R(t) \quad = -\delta^R(t) f^R + \varphi^R(t_0) \tag{10.4}$$

The phase observable is actually a receiver measurement on the beat frequency between the received carrier from the satellite and the receiver-generated reference. Hence, the beat phase ϕ can be written as :

$$\phi^{RS}(t) = \varphi^S(t) - \varphi^R(t) \tag{10.5}$$

By substituting (10.4) in (10.3), it becomes :

$$\phi^{RS}(t) = (f^S - f^R)t - \frac{\rho(t)}{c} f^S + \delta^S(t) f^S - \delta^R(t) f^R + \varphi^R(t_0) - \varphi^S(t_0) \tag{10.6}$$

According to Hoffman et al. [52], the deviation of the frequencies f^S and f^R com-

pared to the nominal frequency f (generally L_1 or L_2) can be neglected. Similarly, according to Traugott [129], the satellite clock offset can be neglected for clarity, although it will be added to the error term later. Hence, the beat phase can be simplified as :

$$\phi^{RS}(t) = -f\frac{\rho(t)}{c} - f\delta^R(t) + \underbrace{\varphi^R(t_0) - \varphi^S(t_0)}_{N} \tag{10.7}$$

The term N, commonly named ambiguity, is not a function of time and so is constant. A geometrical interpretation of (10.7) can be obtained by multiplying it by the wavelength λ and thus obtaining the phase pseudo-range:

$$\Phi(t) = \rho(t) + c\,\delta^R(t) + \lambda N \tag{10.8}$$

So far, just ideal carrier phase measurement models have been studied, however, the measured phase observable will be affected by the typical GPS errors :

$$\chi = c\,\delta^S + E + T - I + m + \varepsilon \tag{10.9}$$

Where

	c	is the speed of light (m/s)
	δ^S	is the satellite clock error (s)
	E	is the ephemeris error (m)
	T	is the tropospheric delay (m)
	I	is the ionospheric delay (m)
	m	is the multipath error (m)
	ε	is the measurement noise (m)

In the end, the phase observable measurement can be expressed as :

$$\Phi(t) = \rho(t) + c\delta^R(t) + \lambda N + \chi \tag{10.10}$$

Doppler Observation

The Doppler observation is a range rate type measurement available on both L1 and L2 frequencies based on the frequency change of the carrier. This frequency shift is caused by the relative motion of the receiver compared to the satellite. Hence, the Doppler frequency can be written as :

$$f_D = -\frac{\dot{\rho}}{\lambda} \tag{10.11}$$

where $\dot{\rho}$ is the time derivative of the receiver-satellite range and λ is the carrier wavelength. It has been seen that both the pseudorange observation and the carrier phase observation were affected by the receiver clock error δ^R. As the Doppler observation is a range rate type measurement, it is not affected by the receiver clock error itself but by its drift. The receiver clock drift $\dot{\delta^R}$ originates from the offset between the internal receiver oscillator and the nominal GPS frequency. According to Kaplan & Hegarty [65], the impact of the receiver clock drift on the Doppler measurement can be approximated by

$$\Delta_D \approx -f\ \dot{\delta^R} \tag{10.12}$$

and thus, the Doppler measurement becomes:

$$f_D = -\frac{\dot{\rho}}{\lambda} - f\ \dot{\delta^R} \tag{10.13}$$

This equation still expresses the frequency shift due to the relative receiver-satellite motion. However, it can be converted into a range rate by multiplying it by the carrier wavelength λ, giving :

$$D = \dot{\rho} + c\,\dot{\delta^R} \tag{10.14}$$

Similarly to the way Doppler measurement is affected by the receiver clock drift rather than the receiver clock error, it is, apart from the measurement noise, affected by the time derivatives of the usual error sources. Following the notation

used in Section 10.1.2, the Doppler observation error expression is :

$$\dot{\chi} = c\,\dot{\delta^S} + \dot{E} + \dot{T} - \dot{I} + \dot{m} + \varepsilon \tag{10.15}$$

The Doppler observation is obtained directly from the Phase Lock Loop (PLL) of the GPS receiver. Different techniques are used to make the observation, usually with an underlying averaging, which can introduce a delay in the observation [20]. To overcome this potential issue, different techniques have been developed to extract the Doppler measurement by post-processing carrier observations [19, 124].

Signal-to-noise Ratio

The signal-to-noise ratio (SNR) is another GPS observable which provides information about the strength of the signal received. From a receiver perspective, these measurements are initially required for the tracking lock detection [45]. However, it has also been found to be a potentially useful estimate of the measurement noise and indirectly its quality [26, 113]. The SNR can then be used to pre-process and weight the measurements in the positioning algorithm.

The different GPS receiver manufacturers have adopted several ways to report the SNR. The most widely used is the carrier-to-noise power density ratio C/N_0 and defined by:

$$C/N_0 = \frac{P_C}{N_0}\ \text{Hz} \tag{10.16}$$

which can also be expressed in dBHz as :

$$C/N_{0\ \text{dBHz}} = 10\log_{10}\frac{\frac{P_C}{N_0}\ \text{Hz}}{1\ \text{Hz}} \tag{10.17}$$

where P_C is the received signal power in [W] and N_0 is the noise power spectral density in [W/Hz]. In general and to allow for pre-processing and observation weighting, the signal strength is reported by the GPS receiver for every satellite

received and for every band used.

The carrier-to-noise power density ratio C/N_0 is an indicator of how well the signal is tracked by the GPS receiver tracking loops which reflects on the GPS observation error [70].

10.1.3. Errors and Correction Models

As it has been described in the different GPS observables, GPS measurements are affected by typical errors. These errors have several origins such as clock errors, satellite ephemerides, atmospheric refractions caused by the ionosphere and the troposphere, multipath and also the inevitable measurement noise. Some of these errors can be mitigated through the use of correction models while some others can benefit from the use of external correction data which may sometime only be available for post-processing. The following sections detail the different error sources and how they have been modelled and mitigated in this work.

Satellite orbit

As GPS satellites are placed on different orbits, it is possible to predict their position. These predictions are made available in the satellite ephemeris as a set of coefficients describing the orbit of each satellite. Several types of ephemerides exists with different purposes. The broadcast ephemeris is essential as it is the ephemeris transmitted in the navigation message and used by the GPS receivers to compute a position. It is available in real-time with coefficients for a daily interval and have an accuracy of approximately one metre.

Besides, precise orbits can be downloaded on the world wide web from agencies such as the International GPS Service (IGS). All the satellite ephemerides pro-

duced by the IGS have a sample interval of 15 minutes. The accuracy of the orbit is dependant on the product used. The Ultra-Rapid orbits are available in real-time with the first half based on observations and the second half predicted with an accuracy of approximately 5 cm. The rapid and final products are both based on observations and available with a latency of 17 to 41 hours for the rapid product and 12 to 18 days for the final product. For both the orbit accuracy is of approximately 2.5 cm.

As the satellite position, based on the ephemeris, is considered fixed when solving for the receiver's position, an error in the satellite's position is directly propagated into an error in the receiver's position. This error can be largely removed using common advanced GPS processing techniques by differencing observations for short baselines [52].

Satellite clock

Alongside the satellite ephemeris, the satellite clock error is also transmitted in the navigation message. Although, GPS satellites use high precision atomic clocks, there is still a non-negligible time drift from the nominal GPS time. This time drift is monitored by the ground segment for each satellite and a correction based on three polynomial coefficients is included in the navigation message.

Alternatively, the IGS proposes several precise clock products. From the ultra-rapid product available in real-time providing a satellite clock error with an accuracy of 3 ns to the final product available within 12 to 18 days latency with an accuracy of 75 ps.

As two observations from a same satellite made at the same time are affected by the exact same satellite clock error, its effect can largely be removed by differencing observations similarly to the satellite orbit error.

Ionosphere

The ionosphere is a layer of the earth's atmosphere which extends from around 80 to 1000 kilometres above the earth's surface. It is a dispersive medium which means that the propagation velocity of a radio signal through it, including a GPS signal, is dependent on the signal frequency [52]. This change in velocity generates a change in the time-of-flight of the signal which is commonly called the ionospheric delay, regardless of it being an actual delay or an advance. Additionally, when a modulated signal travels through a dispersive medium, the carrier and the modulating signals will be affected differently.

From a GPS perspective, this has two implications; on a particular band, L1 for example, the code observable and the carrier phase observable will be subject to a different ionospheric delay. As a general rule, the code observable will experience a delay whilst the carrier phase observable will experience an advance [93]. Secondly, signals on the L1 band will have different propagation velocities than signals on the L2 band and thus a different delay.

The ionosphere contains a high concentration of ions and free electrons in constantly changing proportions. More precisely, the electron density is linked to the number and the dispersion of free electrons which are released when gas molecules are ionized by the sun's radiations. A common expression of this electron density is the *total electron content* (TEC) which is the total number of electrons integrated between two points, along a tube with a one metre squared cross section. The TEC is varying temporally over a number of cycles such as the 11 year sun spot cycle, a seasonal cycle and the diurnal cycle [65].

The ionospheric delay on a GPS signal is proportional to the TEC along the signal path. Many models are available to attempt to mitigate the effect of the ionosphere such as the Klobuchar model which is one of the most commonly used [67]. The Klobuchar model is a function of time and latitude. Its coefficients are

transmitted as part of the navigation messages so that standalone GPS receivers can mitigate the ionospheric effect. However, the Klobuchar model only accounts for approximately 50% of the ionosphere effects at mid latitudes and residual errors will remain [34].

Moreover, as the ionospheric delay is highly dependant on the satellite elevation angle, the Klobuchar model needs to be used alongside a mapping function to refine the Klobuchar estimated delay based on the satellite elevation.

On another perspective, as the ionosphere is a dispersive medium, a dual frequency GPS receiver can mitigate most of the ionospheric effect by combining the observations on the two different frequencies. As an example, taking pseudorange observations on both L1 and L2 allows the following equation to be written [15]:

$$I_{L1} = \left(\frac{L_2^2}{L_2^2 - L_1^2} \right) (\rho_{L1} - \rho_{L2}) \tag{10.18}$$

where:

	I_{L1}	is the pseudorange ionospheric delay on L1 (m)
	$L1, L2$	are the frequencies of the signals (Hz)
	ρ_{L1}, ρ_{L2}	are the pseudorange measurements on each frequency (m)

In this work, the ionospheric delay has been mitigated using the Klobuchar model alongside its associated mapping function.

Troposphere

The troposphere is a layer of the earth's atmosphere which extends from the ground to an altitude varying between 7 km to 20 km. The tropospheric delay is commonly divided into a wet component, mainly composed of water vapour and a dry component, mainly formed of the dry gases N_2 and O_2. Unlike the ionosphere, the troposphere is a non-dispersive medium which results in equal delay errors for L1 and L2 carrier and code signals. As such, dual frequency GPS measurements

can not be used to mitigate its effect.

The dry component, accounting for approximately 90% of the total tropospheric delay, is relatively well modelled [52]. The dry component model is a function of temperature, pressure and humidity along the propagation path of the GPS signal. As such measurements are not easily accessible, it is common to use a standard atmospheric model to estimate them. A standard atmospheric model is usually function of the latitude, height and the Julian day. Additionally, if the ground surface's temperature and pressure are known, the residual error can be down to within 2 to 5% of the total dry component delay [44]. Similarly to the ionosphere, the correction model is generally associated to a mapping function to also take into account the satellite elevation.

Unlike the dry component, the wet component delay, mainly caused by water vapour, is difficult to model due to its temporal and spatial variation. However, differential or interferometric GPS processing techniques can be used to mitigate its effect especially for short baselines.

In this work, the tropospheric delay has been mitigated using the Saastamoinen model [109] whose performance against other models is discussed in Mendes & Langley [90] combined with the Neill mapping function [95].

Multipath

Multipath is the term given to the GPS satellite signals which are reaching the receiver's GPS antenna via an indirect path. Multipath is caused by the environment surrounding the receiver's antenna such as building where the GPS signal can be reflected. It can also be generated by the ground, for example on water, where the signal bounces before reaching the receiver's antenna. Multipath results in a signal arriving at the antenna multiple times and delayed in time.

In general, it is complicated to model multipath due to the virtual infinity of different environmental reflection scenarios. However, there exists a few options which can be used in particular scenarios. For a static GPS receiver, the multipath effect can be reduced thanks to the repeatability of the satellite orbit and the unchanged environment. On the down side, this requires a long occupancy time and can obviously not be applied to kinematic scenarios.

Regarding the code observable, GPS receivers manufacturers have also implemented specific real-time signal processing techniques on the receiver's correlator to mitigate the effect of multipath [143].

Besides, in the presence of multipath, if the direct ray is also tracked by the receiver's PLL, the carrier phase observable error is bounded to a quarter cycle [100].

Finally, the use of an appropriate GPS antenna with regard to its surrounding environment is essential. There exists antennas with ground-planes or choke rings which aim to physically block some of the multipath rays by preventing them to reach the antenna itself.

Under certain conditions, the Right Hand Circular Polarisation (RHCP) property of the GPS signal, can be used to detect and reject multipath. In more details, the GPS signal is emitted by satellites using a RHCP and keep this polarisation if it reaches the receiver's antenna through a direct path. However, in presence of multipath, the received signal might have been reflected as a Left Hand Circular Polarisation (LHCP) signal, depending on the angle of incidence and the material at the reflection point. This property can advantageously be used as GPS antennas are designed to receive a RHCP signal and thus have a much higher gain to receive RHCP signals than LHCP. As a side note, a RHCP signal can turn into a LHCP after a first reflection and back to a RHCP signal after a second reflection. However, each time the signal is reflected, it becomes weaker which implies a

reduced SNR which can also be used to reject multipath (Section 10.1.2).

Receiver clock

Unlike GPS satellites, receivers use relatively inexpensive quartz crystal oscillators to keep the time. As discussed in Section 10.1.3, time synchronisation is critical for GPS. As a consequence, the receiver clock error, i.e. the difference between the receiver time and the nominal GPS time, is estimated as an unknown variable in the positioning algorithm. This additional unknown variable explains why signals from at least 4 satellites are required to compute a position. In an ideal situation where a GPS receiver would be perfectly synchronised on GPS time, signals from only 3 satellites would be required to obtain a position. However, this is not a realistic solution due to the size and cost constraint it would imply on the receiver design.

Receiver measurement noise

Like any measurement device, the observations made by a GPS receiver are affected by a measurement noise. This measurement noise is the result of the signal passing through the different parts of the receiver as well as the antenna and the cable connecting the two. The measurement noise is mainly a white thermal noise whose magnitude is proportional to the quality of the components used. The receiver measurement noise can be estimated by conducting a zero-baseline test [6, 15].

10.2. Time-Differenced Carrier Phase Positioning

Time-differenced carrier-phase positioning is an alternative processing technique allowing precise differential carrier phase positioning without a base station. This technique can be applied using only one, single frequency, low-cost GPS receiver. It also removes the need for ambiguity resolution and initialisation periods. This is achieved by differencing measurements of a single receiver at different epochs instead of differencing measurements of two receivers at the same epoch. Whilst this technique has a lot of clear advantages, it also has some limitations, the main one being that it only achieves relative precise positioning. Additionally, the time differentiation limits the time span to several minutes to maintain precise positioning.

A typical scenario of time-differenced carrier-phase positioning would be to choose a starting time t_0. The position at t_0 would usually be determined by traditional coarse code positioning. Let's denote Δ the bias between the determined position and the true location at t_0. The magnitude of Δ is typically a couple of metres with coarse code positioning. For all the following epochs t, the position will be determined relatively to the starting location by time-differenced carrier-phase. This means that the vector between the position at t_0 and t will be determined precisely using carrier phase. However, the general accuracy will not be improved as the initial bias Δ remains. That is why this processing technique has been described by Traugott [129] as achieving only relative precision as opposed to absolute accuracy.

10.2.1. Time-Differences of Phase Measurements

While Section 10.1.2 describes the phase observable measured by the receiver, it is not the measurement directly used in this technique. As extensively described in the literature, one of the critical part of carrier phase based GPS processing is

to solve for the ambiguities N. If the tracking is continuous without loss of lock to a given satellite, the ambiguity for that satellite remains constant.

Let's assume that the receiver keeps a constant lock with a given satellite PRN between t_1 and t_2. From (10.8) we get :

$$\Phi(t_1) = \rho(t_1) + c\delta^R(t_1) + \lambda N + \chi(t_1) \tag{10.19}$$

$$\Phi(t_2) = \rho(t_2) + c\delta^R(t_2) + \lambda N + \chi(t_2) \tag{10.20}$$

By time differencing these two observations, we obtain :

$$\Phi(t_2) - \Phi(t_1) = \rho(t_2) - \rho(t_1) + c(\delta^R(t_2) - \delta^R(t_1)) + \lambda(N - N) + \chi(t_2) - \chi(t_1) \tag{10.21}$$

The main interest of this differentiation is the ambiguity cancellation. Additionally, there is a "pseudo-cancellation" of the error term χ. In the case of classical RTK GPS processing, the error cancellation is based on the fact that the rover and the base station are close enough to each other to be subject to the same errors. Here, the same principal applies but from a time perspective. In Equation (10.21), there is an error cancellation if the error at t_2 is similar to the error at t_1. However, this assumption is not entirely valid and logically degrades with the time span between t_1 and t_2. To take into account the remaining error sources, let's define:

$$\dot{\chi} = \chi(t_2) - \chi(t_1) \tag{10.22}$$

Using only the strict definition of the error sources χ (Eq 10.9), it is possible to write :

$$\begin{aligned} \chi(t_2) - \chi(t_1) = c\left(\delta^S_{t_2} - \delta^S_{t_1}\right) + (E_{t_2} - E_{t_1}) + (T_{t_2} - T_{t_1}) - (I_{t_2} - I_{t_1}) \\ + (m_{t_2} - m_{t_1}) + (\varepsilon_{t_2} - \varepsilon_{t_1}) \end{aligned} \tag{10.23}$$

Which then lead to :

$$\begin{aligned}\dot{\chi} = &\left[c\,\frac{\partial \delta^S}{\partial t} + \frac{\partial E}{\partial t} + \frac{\partial T}{\partial t} - \frac{\partial I}{\partial t} + \frac{\partial m}{\partial t}\right](t_2 - t_1) \\ &+ \left[\frac{\partial T}{\partial \psi} - \frac{\partial I}{\partial \psi}\right](\psi(t_2) - \psi(t_1)) \\ &+ \varepsilon_{t_2} - \varepsilon_{t_1}\end{aligned} \tag{10.24}$$

Which highlights the fact that the ephemeris error (satellite orbit and satellite clock drift) and the multipath error only have a temporal dependency. On the contrary, tropospheric and ionospheric errors both have a temporal dependency, due to the time variations in the atmosphere, and a spatial dependency which model the geometrical change between two epochs. Details of each error term in the context of time-differenced carrier phase measurements are given in Section 10.2.4.

10.2.2. Positioning algorithm

While Section 10.2.1 has detailed the construction of the time-differenced observable, the current section will present the algorithm used to extract positioning information from the time-differenced carrier phase observable.

Based on the nomenclature used by Blake [15], the following terms are defined :

Ψ = ECEF XYZ position of the satellite

ψ = ECEF XYZ position of the receiver

ρ = Range between satellite and receiver

e = line of sight unit vector

By definition, the range can be expressed as :

$$\rho(t) = e(t) \cdot (\Psi(t) - \psi(t)) \tag{10.25}$$

Let's denote t_1 a base epoch in GPS time and t_2 another epoch ($t_2 > t_1$) with a continuous tracking of satellites between the two epochs to ensure the cancellation of ambiguities. Adding a time difference, leads to :

$$\begin{aligned} \rho(t_2) - \rho(t_1) &= e(t_2) \cdot (\Psi(t_2) - \psi(t_2)) - e(t_1) \cdot (\Psi(t_1) - \psi(t_1)) \\ &= e(t_2) \cdot \Psi(t_2) - e(t_1) \cdot \Psi(t_1) - e(t_2) \cdot \psi(t_2) + e(t_1) \cdot \psi(t_1) \end{aligned} \tag{10.26}$$

Let's define the term $\Delta\Psi(t_1, t_2)$ representing the change in range due to satellite motion between the two epochs as :

$$\Delta\Psi(t_1, t_2) = e(t_2) \cdot \Psi(t_2) - e(t_1) \cdot \Psi(t_1) \tag{10.27}$$

Additionally, let's represent the receiver position at t_2 as being its position at t_1 plus an increment :

$$\psi(t_2) = \psi(t_1) + \delta\psi \tag{10.28}$$

Substituting (10.27) and (10.28) back in (10.26) :

$$\rho(t_2) - \rho(t_1) = \Delta\Psi(t_1, t_2) - e(t_2) \cdot (\psi(t_1) + \delta\psi) + e(t_1) \cdot \psi(t_1) \tag{10.29}$$

Finally, let's define $\Delta geo(t_1, t_2)$ representing the change in satellite to receiver geometry between the two epochs as :

$$\Delta geo(t_1, t_2) = e(t_2) \cdot \psi(t_1) - e(t_1) \cdot \psi(t_1) \tag{10.30}$$

Equation (10.29) can now be simplified as :

$$\rho(t_2) - \rho(t_1) = \Delta\Psi(t_1, t_2) - \Delta geo(t_1, t_2) - e(t_2) \cdot \delta\psi \tag{10.31}$$

Substituting (10.31) in (10.21) makes it clear how the time-differenced observables

can be used to compute the motion of the receiver between two epochs :

$$\begin{aligned}\Phi(t_2) - \Phi(t_1) =& \Delta\Psi(t_1, t_2) - \Delta geo(t_1, t_2) - e(t_2) \cdot \delta\psi + c(\delta^R(t_2) \\ & - \delta^R(t_1)) + \Delta\chi(t_1, t_2)\end{aligned} \tag{10.32}$$

For clarity, let's rewrite the receiver clock offset difference $c(\delta^R(t_2) - \delta^R(t_1))$ as a clock drift $c\dot{\delta}_\tau$.

$$\Phi(t_2) - \Phi(t_1) = \Delta\Psi(t_1, t_2) - \Delta geo(t_1, t_2) \quad -e(t_2) \cdot \delta\psi + c\dot{\delta}_\tau + \Delta\chi(t_1, t_2) \tag{10.33}$$

The unknown variables of (10.32) are :

- $\delta\psi$: Receiver position increment
- $c\dot{\delta}_\tau$: Effect of the receiver clock drift on the range

Which can be represented as the vector of parameters $\mathbf{x} = (\delta\psi_x, \delta\psi_y, \delta\psi_z, c\dot{\delta}_\tau)$.

As there are 4 unknown variables, time-differenced observables to at least 4 satellites are required for the problem to be optimised. If $n > 4$ satellites are tracked, a system of n equations can be formed and solved by least squares.

10.2.3. Trajectory reconstruction

As detailed by J. Traugott [129] there are two strategies to reconstruct the trajectory from the time-differenced carrier-phase observables, namely the over-all strategy and the accumulation strategy.

- Over-all strategy :

    1. A base epoch t_b is fixed with its position determined by coarse code

positioning

2. Positions for the following epochs t are obtained by time-differencing between t and t_b.

3. Hence, the change in position since the first epoch is computed.

- Accumulation strategy :

1. A base epoch t_{i_0} is fixed with its position determined by coarse code positioning

2. Positions for the following epochs t_i are obtained by time-differencing between t_i and t_{i-1}.

3. Hence, the change in position is computed incrementally with the last epoch t_i being then used as the base epoch when the next measurement t_{i+1} comes in.

Both strategies are having advantages and drawbacks which are detailed in the work of J. Traugott [129]. As a brief summary, the over-all strategy requires a continuous tracking starting from the base epoch whereas the accumulation strategy just requires a continuous tracking from the previous epoch. Besides, the accumulation strategy is accumulating measurements affected by random errors which can lead to unbound, random-walk effects [129].

10.2.4. Errors after time-differentiation

The common errors affecting GPS measurements and how they are traditionally mitigated have been covered in Section 10.1.3. When using the time-differenced

carrier phase observable, the budget of each error source is different due to the differentiation. As a general rule, the errors detailed in Section 10.1.3 represent a range error whilst after time-differentiation, they represent a range-rate error which corresponds to the change in the error between the two observations differenced. This section details the error generated by each individual error source in the context of time-differenced carrier phase positioning and is mostly based on the seminal work of J. Traugott [129].

Satellite orbit

An error in the satellite position can be divided into an along-track error, a cross-track error and a radial error. Based on geometry considerations, the along-track error and the change of radial error are the main contributions to the satellite position error budget in the context of time-differenced carrier phase processing [129, 145].

Satellite clock

After time-differentiation, the offset of the satellite clock compared to the nominal GPS time is not directly affecting the observable. However, the satellite clock drift between the two differenced observations will remain and contribute to the error budget. It is hard to entirely separate the satellite orbit errors from the satellite clock errors as they are both very tightly linked. As an example, an error in the satellite clock will directly lead to an error in the satellite position.

Ionosphere

The ionospheric error, after time-differentiation, can be divided into three components. The change in value of the TEC, the change in value of the mapping function and finally the effect of ionospheric scintillation. The TEC value change is due to the temporal variation of the ionospheric activity between the two observations. The change in the mapping function value is due to the different satellite elevations between the differenced observations. And finally, the ionospheric scintillation is more of a random process which is really hard to model and currently subject to a lot of research [7, 27, 53, 123]. Its effect on time-differenced carrier phase positioning is limited (< 1 mm/s) [129] and thus is not mitigated in this work.

Troposphere

After time-differentiation, the tropospheric error can be divided into two components. The first one is the temporal variation of the zenith path delay which is due to the evolution of the weather and secondly the change of the mapping function value which is, similarly to the ionosphere, due to the satellite elevation change. The temporal variation of the zenith path delay is usually not exceeding 0.2 mm/s [129]. Besides, the change in satellite elevation can cause significant range rate errors. Satellites above 10° elevation cause a range rate error of approximately 13 mm/s whilst it can go up to 80 mm/s for satellites with an elevation below 3° [129]. These elevation dependant variations are the reason for the use of mapping functions, such as the Neill mapping function used in this work, to accurately model them.

Multipath

Whilst multipath generates a range error with traditional GPS positioning techniques, it is its change in time that affects time-differenced carrier phase positioning. The resulting error originates from the change in the distance travelled by the reflected signal. This change in distance travelled is affected by both satellite and receiver motion. The error due to satellite motion typically results in range rate errors of the cm/s to the dm/s magnitude. This limited impact is a consequence of the high altitude of GPS satellites, hence the change in geometry due to the satellite motion from a receiver perspective is limited. Besides, a movement of the receiver does significantly affect the geometry in relation to its surrounding environment. Due to the limited bandwidth of the GPS receiver's phase lock-loop (PLL), instantaneous multipath error will be filtered and the time-differenced observable will be affected by an averaged multipath error. In case of continuous tracking, i.e. no loss of lock, this averaged error usually have a mm/s magnitude range [129].

Receiver clock

As described in Equation 10.33, the receiver clock drift is defined as one of the unknown variables to be estimated during the positioning process and is modelled as a random process.

Receiver measurement noise

Under the same white-noise assumption as Section 10.1.3, after time-differentiation between subsequent epochs, the receiver noise is increased by a factor of $\sqrt{2}$ based on the uncertainty and error propagation rules.

11. Strapdown Inertial Navigation

Inertial navigation is a form of navigation belonging to the Dead Reckoning (DR) family. An inertial navigation system uses an Inertial Measurement Unit (IMU) which is generally composed of 3-axis accelerometers, 3-axis gyroscopes and sometimes 3-axis magnetometers. Through the use of the IMU measurements, an inertial navigation system calculates via dead reckoning the position, velocity and attitude of an object or person relative to a known starting position, velocity and attitude. This chapter will introduce the fundamental concepts of inertial navigation known as strapdown inertial navigation.

11.1. Introduction

Inertial navigation is a dead reckoning navigation system which uses an IMU to measure the accelerations and turn rates of an object in three orthogonal directions. Originally, inertial navigation systems were based on gyro-stabilised platforms with a platform placed at the centre of three gimbals which let the platform keep the same orientation regardless of the rotations it is experiencing. The platform's attitude (roll, pitch and yaw) is measured directly using the bearing of the gimbals. Besides, acceleration is measured through the accelerometers attached to the platform. As the platform is gyrostabilised, the accelerometers are always reporting measurements along the same direction. However, such systems have limitations in terms of size and cost. Especially, they require a lot of expensive

high-precision mechanical parts which, as they are subject to movement, can wear out or jam. More recently, with the advent of Micro-Electro-Mechanical Systems (MEMS) sensors, IMU have greatly reduced in size, weight and cost whilst improving in sensing quality with increased range, accuracy and stability. This evolution has lead to the creation of strapdown systems which removed the need for a gyrostabilised platform. With strapdown systems, the IMU is directly attached to the object or person to be tracked which removes the need for extended calibration required with gyrostabilised platforms. MEMS technology also has the benefits of having less moving parts making the sensor more robust. However, these changes were only possible with the increase in sensing range of the MEMS gyroscopes since they need to be able to keep up with the object or person changes in attitude. The rest of this chapter details the principles of strapdown inertial navigation from the mechanisation of raw measurements made along the sensors axes to the production of position, velocity and attitude estimates.

11.2. Coordinate frames

Before tackling the core algorithm and the main equations behind inertial navigation, it is essential to define the different coordinate frames involved and their relationship.

11.2.1. Inertial Frame

The inertial frame (i-frame) has its origin at the Earth's centre of mass and is non-rotating with respect to distant galaxies. It is an ideal coordinate frame where accelerometers and gyroscopes fixed to it would both have zero outputs. The coordinate system is defined with the z-axis parallel to the spin axis of the Earth, the x-axis pointing towards the average vernal equinox and the y-axis completing

a right-handed orthogonal coordinate system [116].

11.2.2. Earth Frame

The Earth frame (e-frame) is also geocentric with axes fixed with respect to the Earth hence rotating with respect to the inertial frame. It is defined with the x-axis pointing towards the intersection of the Greenwich meridian and the Equator, the z-axis is parallel to the mean Earth spin axis and the y-axis is completing a right-handed orthogonal coordinate system.

The Earth's rotation rate about the z-axis which is also the rotation rate of the e-frame with respect to the i-frame projected in the e-frame is defined by :

$$\omega_{ie} = 2\pi \frac{1 + 365.25}{365.25 \times 24 \times 60 \times 60} \approx 7.2921158 \times 10^{-5}\ rad/s \tag{11.1}$$

Both cartesian and geodetic coordinates can be used to express a position in the e-frame. The cartesian Earth-Centered Earth-Fixed (ECEF) coordinate system defines a position as a (x,y,z) vector. This cartesian position can also be expressed from a geodetic position defined by its latitude (φ), longitude (λ) and height (h) using the following equation :

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} (R_N + h)\cos\varphi\cos\lambda \\ (R_N + h)\cos\varphi\sin\lambda \\ (R_N(1 - e^2) + h)\sin\varphi \end{bmatrix} \tag{11.2}$$

where R_N is the Earth radius of curvature in the prime vertical and e is the first eccentricity of the ellipsoid used.

The use of the Earth mean spin axis is a simplification as the spin axis is slowly varying, however, this approximation is considered to have negligible effects on the navigation error [35].

Finally, from an inertial perspective, if an IMU is fixed to the e-frame, gyroscopes sense the Earth rotation and accelerometers sense the associated Coriolis acceleration. Both effects can be compensated as described in Section 11.5.

11.2.3. Navigation Frame

The navigation frame (n-frame) is a local geodetic frame whose origin can be any point on the Earth's surface. It is defined with the x-axis pointing to the ellipsoidal north, the y-axis pointing to the ellipsoidal east and the z-axis pointing down the ellipsoidal normal. Due to its axis definition, it is commonly associated to the North-East-Down (NED) coordinate system.

The n-frame is traditionally used in navigation as it refers to notions which are appealing to the user with the familiar axes North, East and Down. Its origin is generally taken at the current position of the navigation system which, in this case, also coincides with the origin of the body frame (Section 11.2.4).

The rotation rate of the n-frame with respect to the e-frame is usually referred to as the transport rate ω_{en}. The transport rate is generated by the movement of the point origin of the n-frame with respect to the e-frame due to the displacement of the navigation system.

11.2.4. Body Frame

The body frame (b-frame) is, in an inertial navigation context, defined by the instrument axes and the origin point of the IMU. The IMU is generally rigidly attached to the object or person to track. The body frame is also sometimes referred to as the sensor frame. IMU measurements are directly made in the body frame.

11.3. Rotations

Rotations are a key component of inertial navigation, first of all as they are used to represent the attitude of the object or person tracked. Secondly, because they are required to transform from one coordinate frame to another. Different representations are possible for rotations, each of them having advantages and limitations. This section aims to provide a quick overview of different representations. An in-depth study of the different attitude and rotation representations is presented in the work of Diebel [31].

11.3.1. Direction Cosine Matrix

Let's consider the body frame B and the navigation frame N. Both coordinate frames are defined by an origin point and more importantly, regarding rotations, a set of axis respectively $[b_1\ b_2\ b_3]$ and $[n_1\ n_2\ n_3]$. Let's define the angle α_{ij} as the angle between b_i and n_j. The direction cosine matrix (DCM) representing the rotation between the two frames is defined as:

$$[C_b^n] = \begin{bmatrix} \cos(\alpha_{11}) & \cos(\alpha_{12}) & \cos(\alpha_{13}) \\ \cos(\alpha_{21}) & \cos(\alpha_{22}) & \cos(\alpha_{23}) \\ \cos(\alpha_{31}) & \cos(\alpha_{32}) & \cos(\alpha_{33}) \end{bmatrix} \tag{11.3}$$

Direction cosine matrices are convenient to represent rotations as they don't suffer from singularities as Euler angles may suffer. On the down sides, they are not the most compact solution as each DCM is composed of 9 entries which is quite costly.

11.3.2. Euler angles

A common attitude representation is to use the Euler angles (roll, pitch and yaw). Euler angles can also be used to represent a rotation by performing three successive rotations on three different axes [61, 127]. The Euler angles are a convenient attitude representation since it is a three-dimensional vector representing the attitude on the three axis. Hence, each number has an intuitive meaning in terms of orientation. On the downside, Euler angles are subject to singularities also referred to as the gimbal lock. Singularities occur when the second of the three Euler angles rotation sequence aligns the first and third rotation axes. In which case, the three degree of freedom representation degenerates into a two degree of freedom representation. An example is when the pitch angle is equal to 90 degrees. In this situation, the roll and yaw axis becomes aligned and monitor the same rotation, loosing one degree of freedom. As this problem is critical for strapdown inertial navigation, direction cosine matrix and quaternions are preferred over Euler angles for this application [127].

11.3.3. Quaternions

Quaternions are an alternative way of representing a rotation which solve the singularity problem of the Euler angles by adding an extra degree of freedom. Quaternions are based on the principle that a rotation $\|\mu\|$ about a three-element vector $\mu = [\mu_x \ \mu_y \ \mu_z]^T$ can be used to rotate from one coordinate frame to another [127].

From this rotation vector, a quaternion q, representing the same rotation, can be constructed by:

$$q = [q_w \ q_x \ q_y \ q_z]^T \tag{11.4}$$

$$
\begin{aligned}
q_w &= \cos\left(\frac{\|\mu\|}{2}\right) \\
q_x &= \mu_x \; sin\left(\frac{\|\mu\|}{2}\right) \\
q_y &= \mu_y \; sin\left(\frac{\|\mu\|}{2}\right) \\
q_z &= \mu_z \; sin\left(\frac{\|\mu\|}{2}\right)
\end{aligned}
\tag{11.5}
$$

A quaternion representing a rotation must have a norm of 1, i.e. it has to be a unit quaternion. Re-normalisation of a quaternion might be required during processing when the accumulation of rounding errors breaks the unity norm rule and can be achieved with the following equation:

$$
q_{normalised} = q/\sqrt{q^T q} \tag{11.6}
$$

Quaternions, in addition to be robust to gimbal lock, offers a more compact representation (4 numbers) of a rotation than a DCM (9 numbers). It also makes the interpolation of a rotation between two angles much easier than using other rotation representations. Finally, it can easily be re-normalised when necessary, which is not always the case for direct cosine matrices which might be hard to convert back to proper orthogonal matrices.

Thanks to these different advantages, quaternions are the rotation representation used in this work and the rest of this section will provide additional details around the use of quaternions.

Euler angles to/from quaternion

Let's define the Euler angles (φ, θ, ψ) representing respectively roll, pitch and yaw and a quaternion $q = [q_w \; q_x \; q_y \; q_z]^T$.

Knowing the Euler angles an equivalent quaternion can be created with:

$$\begin{aligned} q_w &= \cos(\varphi/2) \times \cos(\theta/2) \times \cos(\psi/2) + \sin(\varphi/2) \times \sin(\theta/2) \times \sin(\psi/2) \\ q_x &= \sin(\varphi/2) \times \cos(\theta/2) \times \cos(\psi/2) - \cos(\varphi/2) \times \sin(\theta/2) \times \sin(\psi/2) \\ q_y &= \cos(\varphi/2) \times \sin(\theta/2) \times \cos(\psi/2) + \sin(\varphi/2) \times \cos(\theta/2) \times \sin(\psi/2) \\ q_z &= \cos(\varphi/2) \times \cos(\theta/2) \times \sin(\psi/2) - \sin(\varphi/2) \times \sin(\theta/2) \times \cos(\psi/2) \end{aligned} \tag{11.7}$$

On the contrary, given a quaternion the Euler angles can be retrieved by:

$$\begin{aligned} \varphi &= \operatorname{atan}\left(\frac{2\left(q_w \times q_x + q_y \times q_z\right)}{1 - 2\left(q_x^2 + q_y^2\right)}\right) \\ \theta &= \operatorname{asin}\left(2\left(q_w \times q_y - q_x \times q_z\right)\right) \\ \psi &= \operatorname{atan}\left(\frac{2\left(q_w \times q_z + q_x \times q_y\right)}{1 - 2\left(q_y^2 + q_z^2\right)}\right) \end{aligned} \tag{11.8}$$

11.4. Gravity compensation

For any measurement made on the Earth, accelerometers will sense, in addition to the acceleration induced by the movement of the object or person tracked, the gravity acceleration. In order to be able to convert the accelerometers output into a valuable information regarding the object or person movement, a gravity model is required.

By using the navigation frame as the reference frame to mechanise the inertial measurements (Section 11.5), the gravity acceleration vector can easily be expressed [118] as:

$$g^n = \begin{bmatrix} 0 \\ 0 \\ \gamma\left(\lambda, h\right) \end{bmatrix} + \begin{bmatrix} \zeta_0 \\ -\eta_g \\ \delta_g \end{bmatrix} \tag{11.9}$$

where $\zeta_0, -\eta_g$ and δ_g are the 3-axis components representing a gravity anomaly

and $\gamma(\lambda, h)$ is a mapping function defined by:

$$\gamma(\lambda, h) = a_1 \left(1 + a_2 \sin^2 \lambda + a_3 \sin^4 \lambda\right) + \left(a_4 + a_5 \sin^2 \lambda\right) h + a_6 h^2 \qquad (11.10)$$

where the coefficients a_i are associated to the WGS-84 coordinate system [115].

However, as explained in the previous section (Section 11.2) and due to the relativistic theory [23], the acceleration generated by the object or person movement and the gravity acceleration can not be distinguished. This explains why the attitude estimation is critical in strapdown inertial navigation as, if the attitude is known, the different gravity axis components can be removed from the accelerometers measurements.

11.5. Navigation frame initialisation and mechanisation

The navigation frame mechanisation is the process of converting raw IMU measurements into an updated position, velocity and attitude of the navigation system. This mechanisation has been extensively described in many reference books [18, 35, 45, 127].

While different coordinate frames can be used for the IMU mechanisation, the navigation frame is a common choice as the results can easily be interpreted by the user. A complete description and comparison of the mechanisation in the different possible frames alongside their advantages and drawbacks is detailed by Groves [45].

11.5.1. Initialisation and alignment

As strapdown inertial navigation is a dead reckoning navigation system, the IMU mechanisation is, by definition, an iterative process where updates of the position, velocity and attitude of the object or person tracked are computed. Like most iterative processes, there is a need for initialisation. In the context of strapdown inertial navigation, the initialisation can actually divided into two key parts. The first one is the determination of the initial position and velocity whilst the second is the determination of the initial attitude which is generally termed as the alignment procedure.

For an outdoor usage, the initial position and velocity are usually determined using GPS. Indoors, a local position is usually of a bigger interest than the global position. That is to say that indoors, the user is more interested to know where he is in a building than knowing where he is on the globe. Local positions can be obtained by several means and partly refers back to what has been covered in the Part II of this thesis. However, although the local position is of a bigger interest, there is still a need for a global position to estimate the gravity acceleration vector (Section 11.4) or the turn rate of the body frame with respect to the navigation frame (Section 11.5.2) for example.

Coarse alignment

The coarse alignment procedure compares the accelerometers output to the local gravity vector to estimate the initial attitude. This alignment procedure can only be applied during a static initialisation period where the system is only subject to the gravity acceleration so that accelerometers output can directly be matched to the gravity acceleration vector. One limitation is that only the roll and pitch angles can be estimated as the gravity vector is always orthogonal to the yaw axis

and thus does not provide any information that can be used to estimate yaw.

Practically, the accelerometers output is averaged over the static period giving a mean acceleration vector in the body frame $f^b = [f_x \ f_y \ f_z]^T$. Under the static assumption, it is possible to write:

$$g^n = C_b^n f^b \tag{11.11}$$

where g^n is the gravity vector expressed in the navigation frame, f^b is the averaged accelerometers output made in the body frame and C_b^n is the direct cosine matrix representing the rotation from the body frame to the navigation frame. From this equation linking the static accelerometers output to the system attitude, it is possible to estimate roll (φ) and pitch (θ) using:

$$\begin{aligned} \varphi &= \operatorname{atan} 2\left(-f_y, f_x\right) \\ \theta &= \operatorname{atan} 2\left(f_x, \sqrt{f_y^2 + f_z^2}\right) \end{aligned} \tag{11.12}$$

The accuracy of this alignment method is thus directly proportional to the accelerometers accuracy. It is also directly correlated to the quality of the static period.

Finally, once roll and pitch have been estimated, if the navigation system and more precisely the IMU includes 3-axis magnetometers, an estimate of the initial heading can be computed. Magnetometers, in the absence of local magnetic perturbations, sense the Earth's magnetic field (H_0) components on the three axes (H_x, H_y, H_z). Similarly to the Equation 11.11, we have:

$$\begin{bmatrix} H_x \\ H_y \\ H_z \end{bmatrix} = C_n^b \begin{bmatrix} \cos\delta\cos\gamma \\ \cos\delta\sin\gamma \\ \sin\delta \end{bmatrix} H_0 \tag{11.13}$$

where δ and γ are respectively the angle of dip and the angle of declination. As the angle of dip is known from the Equation 11.12, an estimate of the navigation

system's heading can be determined. In the particular case where roll and pitch are equal to zero, the heading can be determined with:

$$\psi = \operatorname{atan} 2\left(-H_y, H_x\right) \tag{11.14}$$

Otherwise, in a more general situation where roll (φ) and pitch (θ) estimates are known with:

$$\psi = \operatorname{atan} 2\left(-H_y \cos\varphi + H_z \sin\varphi, H_x \cos\theta + H_y \sin\varphi \sin\theta + H_z \cos\varphi \sin\theta\right) \tag{11.15}$$

If no magnetometers are present, the heading might be retrieved during the dynamic part of the navigation where the course-over-ground (COG) can be extracted from the GPS measurements and used to provide a heading estimate. This technique's accuracy increase with greater velocity.

The principal strength of the coarse alignment procedure is that it is available and functional even for low-cost sensors.

Fine alignment

The fine alignment procedure is an extension of the coarse alignment that can be used when higher quality gyroscopes are available in the IMU. During a stationary period, it has already been seen that accelerometers are only subject to the Earth's gravity acceleration. Similarly, gyroscopes are only subject to the Earth's rotation whilst the system is left static. Unfortunately, unlike the gravity vector which is of a significant magnitude, the Earth's rotation rate is of a very small magnitude and undistinguishable from the noise of low-cost gyroscopes. However, if the gyroscopes are good enough, the Earth's rotation can be sensed and used to provide a heading estimate.

Let's define the Earth's rotation rate $\omega_{ie}^{n} = \begin{bmatrix} \Omega\cos\varphi & 0 & -\Omega\sin\varphi \end{bmatrix}^{T}$ and the direction cosine matrix of the navigation frame with respect to the body frame C_{n}^{b}. According to Titterton et al. [127], when the system is static, the gyroscope measurements ω are equals to:

$$\omega = \omega_{ie}^{b} = C_{n}^{b}\omega_{ie}^{n} \tag{11.16}$$

An estimate of the system's heading can be determined as all the other variables are known or have already been estimated.

Dynamic alignment

In certain applications or scenarios, having a static initialisation period is not feasible. An alternative is to perform the alignment during the navigation. One clear drawback of this technique is that, for real-time users, until the alignment is completed, the position, velocity and attitude reported by the navigation system will be unusable. Besides, as a post-processing alternative, this solution is totally viable. In many ways, the dynamic alignment is just about applying the navigation frame mechanisation which will be detailed in the next section. It is based on the same principles as the coarse alignment with the gravity vector always being sensed by the 3-axis accelerometers. The heading can be determined from the GPS course-over-ground as already mentioned above.

Coarse and fine alignments have the advantage of being robust techniques to initialise the navigation system prior to the actual navigation. Whilst they are not mandatory, they can usually easily be performed and improve the robustness of the navigation system.

11.5.2. Navigation Frame Mechanisation

The navigation frame mechanisation is the process of transforming raw IMU measurements into updated estimates of the position, velocity and attitude of the system. As strapdown inertial navigation is part of the dead reckoning family, it results in an iterative process where estimates (position, velocity and attitude) are updated each time a new measurement is made by the IMU.

Before processing the instrument measurements, the algorithm applies corrections to both accelerometers and gyroscopes output. The correction terms can either be based on an existing factory or laboratory calibration or either be coming from correction terms that are estimated as part of the optimisation process (Section 12.2.3). It is also not uncommon to have a combination of both correction types.

Let's denote w^b and f^b the gyroscopes and accelerometers measurements of the IMU made in the body frame.

The first step of the mechanisation is to correct the gyroscope measurements for any rotation rate sensed that is not generated by the system's movement. In more details, two effects need to be compensated. The first one is to compensate for the Earth's rotation rate, ω_{ie}^n, and the second one is to compensate for the navigation frame transport rate, ω_{en}^n. After these two compensations, the turn rate of the body frame with respect to the navigation frame referenced in the body frame, ω_{nb}^b, is obtained and directly represents the rotating motion of the object or person tracked.

$$\omega_{nb}^b = \omega^b - C_n^b \left(\omega_{ie}^n + \omega_{en}^n\right) \tag{11.17}$$

Where the different variables are defined as :

$$\begin{aligned}
\omega_{ie} &= 2\pi \frac{1 + 365.25}{365.25 \times 24 \times 60 \times 60} \approx 7.2921158 \times 10^{-5}\ rad/s \\
\omega_{ie}^{n} &= \begin{bmatrix} \omega_{ie}\cos\varphi & 0 & -\omega_{ie}\sin\varphi \end{bmatrix}^{T} \\
\omega_{en}^{n} &= \begin{bmatrix} \frac{v_E}{R_\lambda + h} & \frac{-v_N}{R_\varphi + h} & \frac{-v_E \tan\varphi}{R_\lambda + h} \end{bmatrix}
\end{aligned} \tag{11.18}$$

with the latitude (φ), the longitude (λ) and the altitude (h). In addition, R_φ and R_λ are the latitudinal and longitudinal Earth's radius associated to an ellipsoid, usually WGS-84.

This compensated gyroscope measurement, ω_{nb}^{b}, can now be used to compute an updated attitude. Let's define a quaternion $\mathbf{q}$ representing the navigation system's attitude and ω_{nb}^{b} the compensated gyroscope measurements. Before computing the actual update, let's define:

$$\begin{aligned}
\omega &= \omega_{nb}^{b} \times \mathrm{dt} \\
\omega_0 &= \|\omega\| = \sqrt{\omega_x^2 + \omega_y^2 + \omega_z^2}
\end{aligned} \tag{11.19}$$

where dt is the measurement rate of the gyroscope. Let's additionally define:

$$\Omega = \omega_{nb}^{b} \times \frac{\sin\left(\omega_0/2\right)}{\omega_0} \tag{11.20}$$

From these definitions, a quaternion $\mathbf{dq}$ representing the update in rotation measured by the gyroscope is defined as:

$$\mathbf{dq} = \begin{bmatrix} \cos\left(\frac{\omega_0}{2}\right) & \Omega_x & \Omega_y & \Omega_z \end{bmatrix} \tag{11.21}$$

Finally, the updated attitude is computed with:

$$\mathbf{q}_{\mathrm{t}} = \mathbf{q}_{\mathrm{t}-1} \times \mathbf{dq} \tag{11.22}$$

where $\times$ is the quaternion multiplication.

Once the updated attitude has been computed, it can be used to transform the acceleration measured in the body frame to the equivalent acceleration in the navigation frame. This is done by rotating the measured acceleration f^b with the updated quaternion $\mathbf{q}$ to obtain the acceleration in the navigation frame f^n.

The next step, similarly to what has been done with the gyroscope measurements, is to compensate for any acceleration experienced by the system which has not been generated by its motion. More precisely, the accelerometers measurements have to be compensated for gravity and Coriolis acceleration to obtain compensated accelerometers measurements:

$$\dot{v^n} = f^n - (2\omega_{ie}^n + \omega_{en}^n) \times v^n + g^n \tag{11.23}$$

An updated velocity estimate of the system in the navigation frame is obtained by integrating Equation 11.23 over the accelerometers measurement rate dt. Repeating this integration a second time yields to the updated position estimate.

At this stage all the INS components (position, velocity and attitude) have been updated and the navigation frame mechanisation is finished. This process is repeated each time a new measurements is made by the IMU leading to updated estimates.

12. A Novel Graph-Based GPS and INS integration

12.1. Time-Differenced Carrier Phase Positioning

12.1.1. A Novel Graph-based implementation

This section describes the different stages necessary to model the time-differenced carrier phase observables into a factor graph to be then optimised. The first step is to encode the problem as a factor graph. This can easily be done by having one vertex for each GPS epoch holding the parameters to be estimated (position and receiver clock drift) and edges connecting two epochs with the time-differenced carrier-phase observables. Figure 12.1 shows an example of the novel factor graph formulation of time-differenced carrier phase positioning proposed.

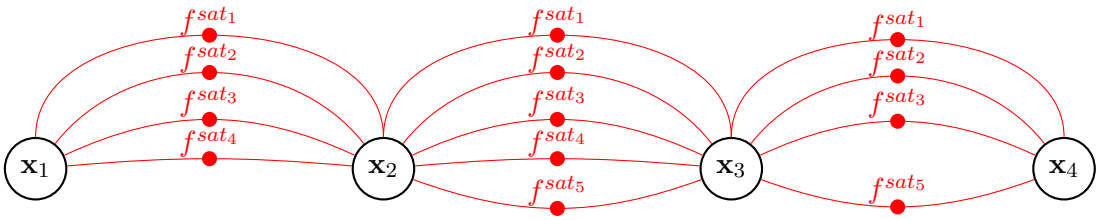


Figure 12.1.: Factor graph representation of time-differenced carrier phase positioning

Each vertex contains the set of parameters to be estimated for each epoch. Vertices

are then connected by a set of edges encoding the time-differenced carrier phase measurement with one edge per satellite. On the situation depicted on Figure 12.1, 4 satellites are received between the first two epochs, then a 5^{th} satellite is received before loosing track of the 4^{th} satellite in the last epoch represented.

Now that the graph representation of the problem has been detailed, the second step is to define the error function encoded by the time-differenced carrier phase positioning factor.

As mentioned above, each vertex holds estimates of the following parameters $\mathbf{x}_t = (\delta\psi_x, \delta\psi_y, \delta\psi_z, c\dot{\delta}_\tau)$ at time t. These parameters were defined in Section 10.2.2.

Starting from (10.33) and ignoring the term modelling the uncompensated errors (χ), the receiver clock drift can be expressed as :

$$c\dot{\delta}_\tau = \Delta\Phi(t_1, t_2) - \Delta\Psi(t_1, t_2) + \Delta geo(t_1, t_2) + e(t_2) \cdot \delta\psi \tag{12.1}$$

Hence, the error function on the receiver clock drift is :

$$Error(c\dot{\delta}_\tau) = [\Delta\Phi(t_1, t_2) - \Delta\Psi(t_1, t_2) + \Delta geo(t_1, t_2) + e(t_2) \cdot \delta\psi] - c\dot{\delta}_\tau \tag{12.2}$$

Similarly, the error function on the receiver position is :

$$Error(-e(t_2)\cdot\delta\psi) = [\Delta\Phi(t_1, t_2) - \Delta\Psi(t_1, t_2) + \Delta geo(t_1, t_2) - c\dot{\delta}_\tau] - e(t_2)\cdot\delta\psi \tag{12.3}$$

The 3 components of the position residual error in ECEF coordinates are then retrieved by multiplying Eq. (12.3) by $-e(t_2)$. The combination of these two error functions defines the complete 4-dimensional residual of the time-differenced carrier phase positioning factor.

12.1.2. Loss Function

As presented in Section 9.3.3, non-linear least squares can be robustified by using loss functions. A loss function scales the residuals to reduce the weight of potential outliers. Whilst a loss function can be applied to every single residual, this has a significant impact on the optimisation time as it adds several mathematical operations to each residual. As a trade off, this work has only applied a loss function to residuals that were more likely to be outliers.

As detailed previously, it is known that GPS measurements made on satellites with a lower elevation are more likely to be affected by outliers. This is mainly due to larger atmospheric propagation effects and an increased likelihood of multipath reflection. As a result, this work has applied a Cauchy loss function for all the measurements made on satellites with an elevation lower than 30°. The value of the elevation threshold has been determined empirically.

The scaling of the loss function has been done based on the GPS receiver's measurement noise. The carrier phase measurement noise is typically of a few millimetres depending on the receiver's quality. In this work, the Cauchy loss function was scaled using $\alpha = 0.005$ metres in accordance to the carrier phase measurement noise.

12.1.3. Weighting Scheme

As described in Section 10.1, GPS signals are affected by several sources of noise and errors which, for most of them, are satellite dependent. Hence, at one epoch, measurements to different satellites are affected by different measurement noise. During the position estimation process, if all the measurements are assumed to have an equal noise and as such an equal weight, there is a risk for a noisy satellite measurement to corrupt the position estimate. Different techniques have been

developed to propose a weighting model of the measurement noise [26]. The majority of the weighting techniques are either using a satellite elevation weighting model or either use the signal-to-noise ratio (Section 10.1.2) output by some GPS receivers. In the context of time-differenced carrier phase positioning, the use of a weighting scheme is also critical. Whilst the work from Traugott [128, 129] does not use a weighting scheme, the work from Blake [15] incorporate a SNR weighting scheme. In the present work, both elevation and SNR based weighting schemes have been investigated. This section will detail the different weighting scheme evaluated, whilst practical results are presented in Section 13.1.

Elevation Weighting Scheme

A common way to weight GPS observations is to use a weighting function dependent of the satellite elevation angle as viewed by the GPS receiver. Different weighting schemes have been developed in the literature, all of them being empirically derived and trying to improve the positioning accuracy. Whilst this may appear as a new concept, it has actually already been introduced in this manuscript with the use of mapping functions for tropospheric and ionospheric delays detailed in Section 10.1.3 and Section 10.2.4. It is then not a big surprise to find that an historical elevation weighting function used to weight GPS observations is the tropospheric mapping function (TMF) itself. Its use is based on the fact that the tropospheric delay error increases with low elevation angles, as does the noise affecting the GPS signal. The TMF as defined by Black and Eisner [14] is :

$$TMF(e) = \frac{1.001}{\sqrt{0.002001 + \sin^2(e)}} \tag{12.4}$$

where e is the satellite elevation angle as seen by the GPS receiver in radians. The weight to be used in the weighted least-squares is the inverse of the TMF so that a measurement made from a satellite at the zenith has a maximal weight of 1 with the weight decreasing towards 0 with decreasing elevations.

Based on the historical TMF weighting function, an alternative weighting model has been developed as part of the Rtklib GPS processing library [126]. The Rtklib elevation based weighting scheme is defined as :

$$q(e) = 0.003 + \frac{0.003}{\sin(e)} \tag{12.5}$$

where the weight to be used in the weighted-least squares is $w = q(e)/q_0$ with q_0 being the maximal weight for a satellite at the zenith (i.e. $e = \pi/2$).

Another elevation based weighting scheme investigated is the Bernese one defined as:

$$w = 1/\sqrt{(}\sin(e)) \tag{12.6}$$

Figure 12.2 illustrates the different weights obtained by the TMF, Rtklib and Bernese weighting schemes depending on the satellite elevation angle.

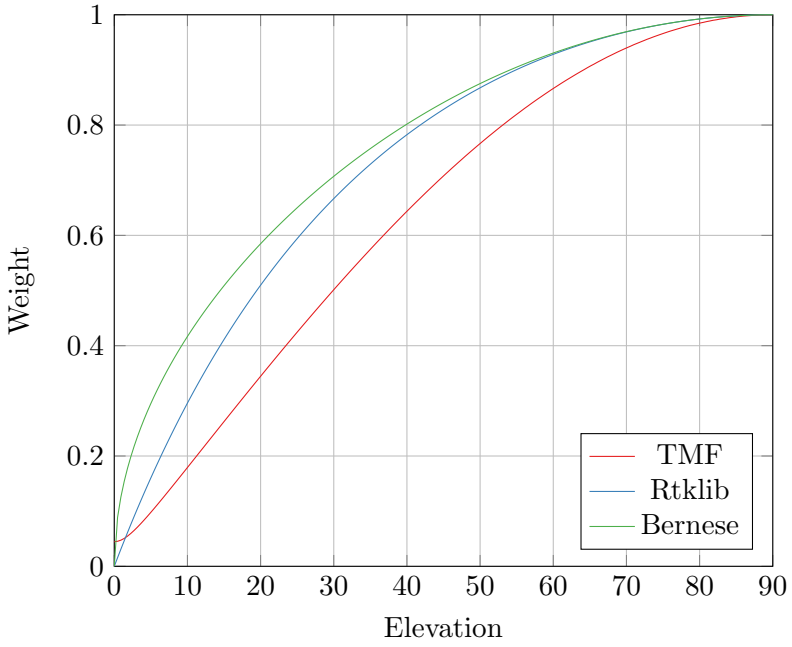


Figure 12.2.: Comparison of the different elevation based weighting schemes

There exists also alternatives to predict the measurement error based on the satellite elevation and some empirically derived constants for a specific receiver and antenna association. For example, Jin [62], uses the following exponential func-

tion:

$$y = a_0 + a_1 * \exp\left(\frac{-e}{e_0}\right) \tag{12.7}$$

where e, the elevation, and e_0 are in degrees and a_0, a_1 are in metres. The variables a_0, a_1 and x_0 are empirically derived for each receiver and antenna association. Due to the requirement of specific testing for each receiver-antenna combination, this solution has not been selected to meet the easy accessibility and practicability requirements of the sport.

Another alternative to reduce the impact of low elevation satellites measurements is to use a higher elevation cut-off angle. Whilst this alternative indeed excludes all, possibly noisier, measurements made to low elevation satellites, it has the drawback of reducing the number of available measurements which also weakens the geometry of the GPS observations.

Signal-To-Noise Ratio Weighting Scheme

Another family of weighting schemes is based on the SNR measurement (Section 10.1.2) output by the GPS receiver. The SNR, as it quantifies the strength of the signal received, has been found to be a potentially useful estimate of the measurement noise. Some original work on SNR based weighting schemes can be found in Collins & Langley [26], Satirapod et al. [113] and Hartinger & Brunner [47]. More recent contributions have been built on top of the original work by Wieser et al. [146, 147]. An extended analysis of the different SNR based weighting schemes has been made by Luo [79].

In general, the SNR is available from GPS receivers as a noise power density ratio C/N_0 in dB-Hz where C is the carrier power and N_0 the noise power density. As the C/N_0 ratio is an indicator of how well a GPS signal is tracked by the GPS receiver tracking loops, it also is an indicator of the potential noise in the GPS observation [70]. In comparison of the elevation based weighting schemes,

SNR weighting is theoretically more appropriate to model the measurement error as it is a measure made directly by the GPS receiver and associated with an observation. For this reason, it is the weighting scheme used in Blake [15] to weight the different GPS observations in his Temporal Double Differencing (TDD) carrier phase positioning.

In order to convert the C/N_0 to a measurement error, the initial contribution has been made by Langley [70] who published the following relation between the phase variance and C/N_0:

$$\sigma^2_{\phi_i} = C_i \times 10^{-\frac{C/N_0}{10}} \tag{12.8}$$

where i denotes the band (i.e. L1 or L2), $\sigma^2_{\phi_i}$ is the carrier phase variance in m^2, C/N_0 is the carrier to noise power density ratio in dB-Hz and finally C_i is a factor in m^2Hz depending on the carrier tracking loop bandwidth (Hz) and a wavelength λ_i related quadratic term in metres.

Based on this, the SNR weighting scheme used in Blake [15] is defined as:

$$\frac{1}{w^s} = \sqrt{\frac{1}{10^{\frac{C/N_0}{10}}}} \tag{12.9}$$

with w^s is the weight associated to the measurement made for a given satellite s.

More recently, an extended weight model based on the C/N_0 ratio has been proposed by Wieser & Brunner [146] where the variance of the phase observation is described as:

$$\sigma^2_{\phi_i} = V_i + C_i \times 10^{-\frac{C/N_0}{10}} \tag{12.10}$$

where C_i and V_i are model parameters respectively in m^2Hz and m^2 that are dependant on the receiver-antenna combination. For the same reasons as for elevation based weighting schemes, this model has not been applied as the specific calibration of each receiver-antenna combination does not meet the sport requirements. As a result, the SNR based weighting scheme investigated is the one presented in

Equation (12.9).

12.1.4. Ephemeris

In the existing literature of time-differenced carrier phase positioning techniques, the impact of precise ephemerides and precise satellites clock has not been fully investigated. Below is a summary of the different contributions regarding the impact of the type of ephemerides used.

The studies of Ding & Wang [32], Wendel et al. [144] and Zhao et al. [151] don't provide any details on the type of ephemeris used and the motivating reasons. According to different clues in the text, the studies seem to be all using broadcast ephemeris.

Work from Blake [15] clearly mentions the use of broadcast ephemeris based on the ephemeris work conducted by Zhang et al. [149] but does not present any analysis of this choice in the context of time-differenced carrier phase processing.

Work from Grass & Soloviev [139] mentions that the use of broadcast ephemeris results in negligible errors in the satellite orbit besides highlighting the benefits of using precise ephemeris in the presence of abnormally large orbit errors before they are corrected by the ground segment. However, no further analysis is developed in the paper.

On the contrary, Serrano et al. [114] make an extended analysis on the use of broadcast ephemeris. The primary focus is to highlight the validity of broadcast ephemeris usage, however, no direct comparison of the results obtained with precise ephemeris is presented.

Finally, Freda et al. [40] highlights the presence of discontinuities in the satel-

lite positions and clock offsets when different broadcast ephemeris sets are used. The work also mentions and compares the use of precise ephemeris against the broadcast ones, however, no detailed results are presented as it claims that the positioning accuracy obtained with broadcast ephemeris is equivalent to the one obtained with precise ephemeris.

Section 13.1.3 and 13.2.6 will present the results of the investigation done in this work in order to provide a deeper comparison of the different ephemeris and clock products available and their impact on the positioning accuracy.

12.2. Graph-based Inertial Navigation

The fundamentals of strapdown inertial navigation and the underlying equations have been detailed in Chapter 11. All the principles and equations that were detailed have been known for years and extensively used. The iterative process of position, velocity and attitude estimation is usually achieved with an Extended Kalman filter to account for the non-linear errors that are part of the system.

The present work chose to use a different approach to process the measurements through the use of a non-linear least squares graph optimisation. The fundamentals as well as the reasons behind the choice alongside its advantages and disadvantages have been covered in Chapter 9.

This section aims to detail how strapdown inertial navigation can be formulated as a graph problem and how other sensors can be integrated. Starting from the original factor graph formulation of strapdown inertial navigation, as developed by the robotics community [29, 55, 57], novel approaches are being presented. The first contribution, presented in Section 12.2.1 is the graph formulation of the "tightly-coupled" integration of strapdown inertial navigation with the developed GPS processing (Section 12.1) . The second contribution is an alternative

and novel way of summarising IMU measurements (Section 12.2.2). Finally, new factors are proposed to take advantage of the specificities of the sprint canoeing and kayaking application to further improve the quality of the solution obtained (Section 12.3).

12.2.1. A novel tightly-coupled integration for GPS and IMU sensor fusion

This section describes how strapdown inertial navigation can be formulated as a graph and how it can be integrated with other sensors. The initial implementation is based on the original work of Indelman & Williams [55] which has been largely reused by the community. It also is the simplest approach which is to create an edge between each consecutive IMU measurement as shown on Figure 12.3.

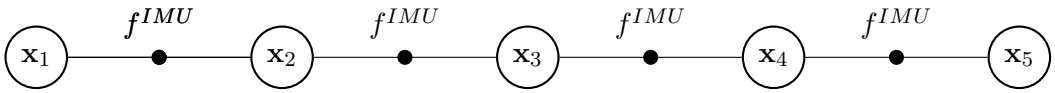


Figure 12.3.: Factor graph representation of strapdown inertial navigation

The next step is to incorporate GPS measurements. The approach used in the literature by Indelman et al. [57] is a "loosely-coupled" integration between IMU and GPS measurements. This means that raw GPS measurements are processed separately to obtain a position estimate. This position estimate is then used and applied as a new factor added to the graph and shown on Figure 12.4.

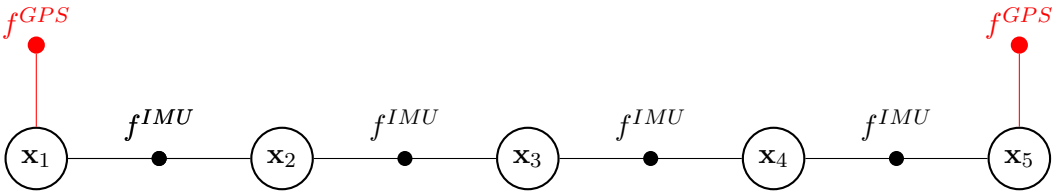


Figure 12.4.: Factor graph representation of strapdown inertial navigation integrating GPS positions in a loosely coupled form

A key contribution of this work is to propose a novel 'tightly-coupled' approach to integrate IMU and GPS measurements. Whilst such "tightly-coupled" ap-

proaches exist for inertial navigation based on Kalman Filters [51, 127], it has never been applied to the graph-based optimisation method. A different form of "tightly-coupled" integration in the context of graph optimisation has been done by Leutenegger et al. [73] to integrate vision sensors with an IMU.

The advantage of a "tightly-coupled" approach is that all the measurements are jointly optimised giving a global optimal solution. This implies that GPS measurements help the inertial integration as much as the IMU measurements help the GPS positioning algorithm. In our case, IMU measurements are "tightly-coupled" with time-differenced carrier phase measurements.

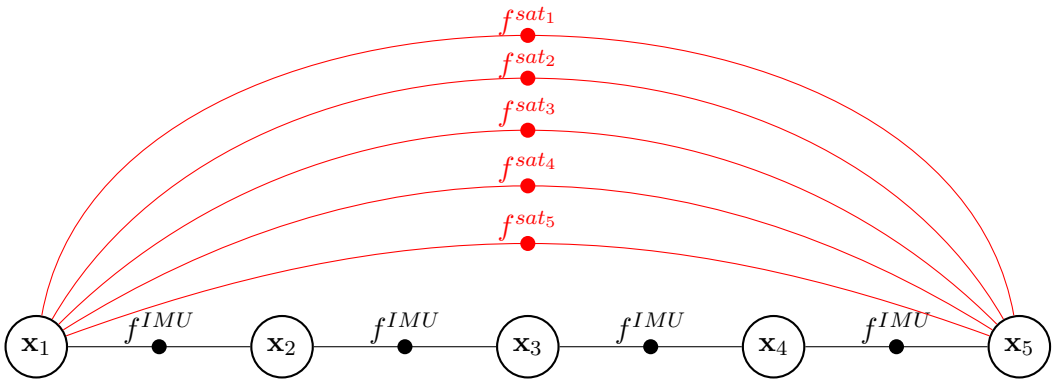


Figure 12.5.: Factor graph representation of the proposed tightly-coupled integration of strapdown inertial navigation and time-differenced carrier phase positioning

The factor graph modelling of the proposed tightly-coupled integration is visible on Figure 12.5. It can be seen that the position, velocity and attitude estimates $\mathbf{x}_2$, $\mathbf{x}_3$ and $\mathbf{x}_4$ are linked together only by IMU measurements. However, there is a GPS measurement linking $\mathbf{x}_1$ and $\mathbf{x}_5$. Thanks to the global optimisation approach, both systems will benefit from the information provided by the other one.

As a side note, the representation of one GPS measurement every 4 IMU measurements is quite unrealistic as a typical GPS update rate is 1Hz and an IMU update rate is 200Hz. As a result, the reader should imagine a chain of 200 IMU factors between two GPS measurements which leads us to the focus of the next section on how IMU factors can be summarised.

12.2.2. Summarised IMU Factor

The previous section proposed a novel factor graph formulation for the tightly-coupled integration of IMU and GPS. However, it has been shown that adding a factor for each IMU measurement dramatically increases the size of the problem which subsequently slows down the optimisation process. In addition, in between external aiding measurements, GPS in our case, there is no extra information that can be used to aid and constrain the inertial navigation solution. In other words, in between measurements from other devices, the inertial navigation graph subset is totally independent. Based on these facts, several alternative have been proposed all based on the seminal work of Lupton's PhD thesis [80] and his papers [81, 82].

Lupton has introduced the idea of pre-integrated IMU measurements. In his original work, the optimisation was done using a delayed state information smoother. Notably, due to the linearity of the equations, pre-integrated IMU measurements can be utilised within Extended Kalman Filters [81] to reduce the computational cost. More recently, the robotics community applied Lupton's technique to graph based SLAM optimisation which is now extensively used [21, 36, 54, 56, 57]. Figure 12.6 shows how the problem depicted on Figure 12.5 is transformed by using pre-integrated IMU measurements in a new factor denoted as f^{equiv}.

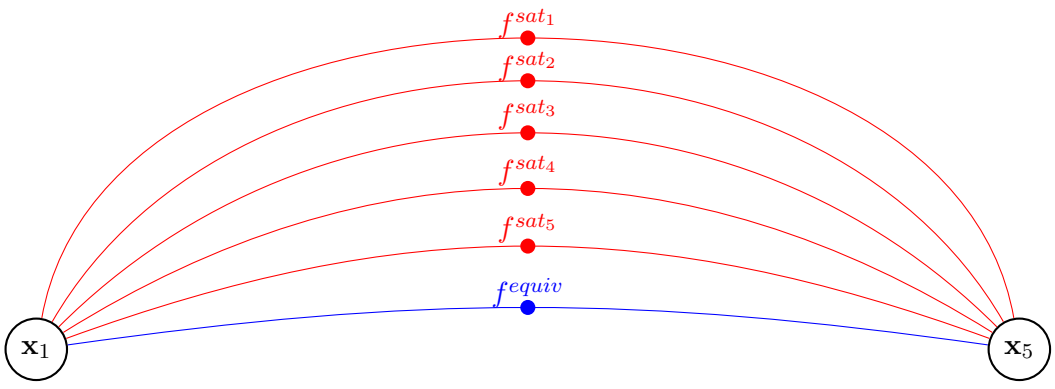


Figure 12.6.: Factor graph representation of the developed tightly-coupled integration with the use of pre-integrated IMU factors.

The validation of the theory in the context of graph optimisation is more recent with the analysis and mathematical proof done by Forster et al. [37, 38].

This approach computes a single update for position, velocity and attitude which is obtained by pre-integrating all the IMU measurements made between t_1 and t_2. This single update is then used in the problem to replace all the IMU measurements that have been pre-integrated. A clear advantage of this technique is a drastic reduction of the number of factors and parameters that are added to the graph and thus reducing the size of the problem to be solved by non-linear least squares.

More recently, the work of Forster et al. [37, 38] emphasised the validity of the IMU pre-integration, especially regarding the non-linearity induced by the rotations. It also proposed to incorporate bias updates (Section 12.2.3), by pre-computing the associated Jacobian describing how the pre-integrated update would change due to a change in the bias estimate. This technique allows to account for small variations of the bias estimate, however, for larger bias changes, the linearisation point used to generate the Jacobian may not be valid.

For this reason, an alternative and novel way of summarising IMU measurements has been developed in this work. The proposed solution is to aggregate all the IMU measurements made between t_1 and t_2 into one factor but not to perform any pre-integration. Hence, the number of factors and parameters composing the graph and the underlying problem is still minimised whilst a complete re-linearisation is still possible.

The graph representation of the proposed approach is strictly identical to the one obtained with pre-integrated IMU factors (Figure 12.6). The difference actually lies in the equations composing the IMU factor itself. Whilst the usual approach summarises a set of measurements into one fixed update, the proposed approach recomputes each time a new update based on the current linearisation point. As a result, the proposed approach can better accommodate changes in the bias

estimate.

In direct comparison to the pre-integration technique, it requires slightly more processing power since, at each Levenberg-Marquardt (Section 9.3) iteration, all the IMU measurements are mechanised to obtain the resulting update whilst the standard approach would have pre-computed a fixed update but with assumptions on the linearisation point. Notably, the computational cost is still significantly reduced compared to the original formulation (Figure 12.5) and allows the developed approach to be used in real-time.

12.2.3. Bias modelling

Commonly, IMU sensors biases are estimated as part of the inertial navigation problem. This is done with both Kalman Filters approaches [45, 51, 127] and graph optimisation approaches [29, 37, 57]. The present work uses the bias modelling as described by Indelman et al. [57] and presented on Figure 12.7.

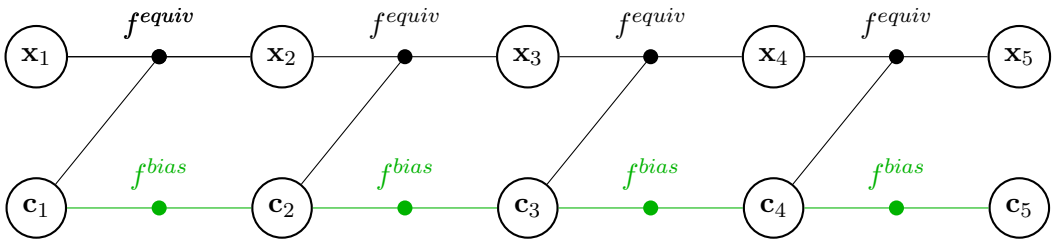


Figure 12.7.: Factor graph representation of strapdown inertial navigation with bias modelling factors

As it can be seen, new vertices are created to hold the IMU bias estimates. These estimates are also linked to the IMU factor as they are used to correct the IMU measurements. In addition, the bias factors are linked together as the bias estimates at t_2 are linked to the previous ones at t_1. In general, the time propagation of the calibration parameters can be described by a random walk process. A typical vertex holding bias estimates is composed of a 6-length vector. This vector contains the bias estimates for each axis of the accelerometers (3-axis) and

gyroscopes (3-axis).

12.3. Taking advantage of the application specificities

So far the positioning equations defined are completely general. This section aims to investigate the specificities of the sprint canoeing and kayaking application that could be advantageously by used by the positioning algorithm. As a result, two new factors have been developed both providing some constraints on the position. They are of great interest as it has been seen that the time-differenced carrier phase positioning technique used is subject to drift. The extra information provided by the developed factors can advantageously be used to constrain that drift.

12.3.1. Height constraint

The first specificity identified is the constant altitude of the athlete during sprint canoeing and kayaking. By definition, the surface of a lake has a constant altitude. Because the athlete is performing directly on the top of the lake surface, the altitude of the boat is also constant plus or minus small vertical oscillations during the effort.

This specificity has advantageously been used to develop a height constraint factor. The developed factor is directly based on the traditional formulation of GPS factors as used in the literature and as represented on Figure 12.1. Whilst a GPS factor provides a 3-dimensional position constraint, the developed height constraint factor is only 1-dimensional, on the vertical component.

Two questions are however still to be answered. The first one is how often the

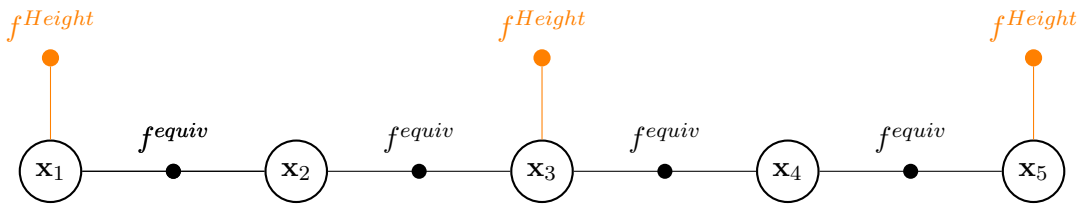


Figure 12.8.: Factor graph representation of strapdown inertial navigation integrating height constraint factors

height constraint factors should be inserted in the graph. This question has been investigated empirically during the trials (Section 13.1.3 and 13.2.6). Height constraint factors were inserted for each GPS measurement, then every 100 measurements and finally every 1000. It has been found that height constraint factors inserted every 100 GPS measurements provided comparable positioning performance to factors inserted at each single GPS measurement for a reduced computational cost. Inserting them only every 1000 GPS measurements however resulted in degraded positioning performance. As a result, the preferred insertion rate is every 100 GPS measurements.

The second question to be answered is how to accommodate the vertical oscillation movement of the boat in the factor. Typically, a boat is subject to vertical oscillations in the magnitude of $6-8$ cm during a sprint. Referring back to the probabilistic formulation of the problem (Section 9.1) and more precisely Equation 9.9, a factor is defined as an error function and an information matrix $\Omega_k = \Sigma_k^{-1}$. The information matrix is generally the inverse of the measurement covariance matrix. In our case, it can be used to incorporate the allowed vertical oscillation of the boat in the factor. In order to avoid over constricting the problem and possibly mask some vertical osculations, the height constraint factor has been implemented using $\sigma = 0.25$ metres and the information matrix is defined as:

$$\Omega_k = \Sigma_k^{-1} = \frac{1}{\sigma^2} \tag{12.11}$$

12.3.2. Loop closure

Another novel contribution is a loop closure factor highly inspired by the work of the robotics community. In the SLAM problem, it is common to have a robot visiting the same place several times whilst exploring an unknown area. Whenever a robot goes by a place more than once, its position at that place should be the same for each of its visit. To take advantage of this situation, the robotics community uses loop closure detection and applies a position constraint each time a loop closure is detected.

From a sport perspective, it is fairly common to start and finish a training session at roughly the same place. By simply marking down the starting place and putting the GPS unit back on it at the end of the session, a loop closure constraint can be applied. As the point marked does not need to be surveyed or to have known coordinates, it is a procedure that can be fairly easily deployed in the sport. Nevertheless, the developed solution does not require loop closures to be functional. However, if a loop closures are available, it can take advantage of them to constrain the solution. A typical factor graph modelling of loop closure is shown on Figure 12.9.

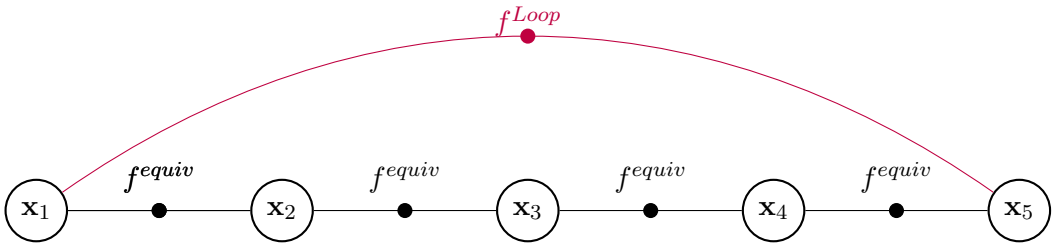


Figure 12.9.: Factor graph representation of strapdown inertial navigation integrating a loop closure factor

A loop closure factor applied between the start and end of a session is only beneficial for post-processing as the factor can not be applied until the session is finished. However, as demonstrated in Section 13.1.3 and 13.2.6, loop closure factors significantly improve the positioning quality.

13. Practical validation

In order to evaluate and validate the GPS and inertial processing techniques respectively developed in Chapter 10 and Chapter 11, different validation trials have been conducted. The first one (Section 13.1) focuses on the evaluation of the GPS technique in a static context whilst the second part (Section 13.2) evaluates the performance of the integrated GPS-inertial solution during on-water sprint kayak trials. In the second part, both the GPS positioning quality and inertial navigation solution are investigated.

13.1. GNSS - Static validation

This section aims to present the results obtained with the proposed time-differenced carrier phase positioning technique in a static context. The data has been collected using two different GPS receivers connected to the same antenna left static (in the same location) in a clear-sky environment. With the data collected being static, its quality and more importantly its drift can easily be assessed. It will first present the results obtained by a low-cost GPS receiver against a survey grade one using the optimal configuration. It will then present the impact of the different corrections detailed in Section 10 using the low-cost GPS receiver's data.

13.1.1. Equipment and location facility

The data has been collected at the Nottingham Geospatial Building at the University of Nottingham. A Leica AS10 GPS antenna was left static on a surveyed pillar on the roof of the building. Thanks to the elevated location of the antenna, a perfect sky visibility was ensured during the complete data collection. Data was collected for over 90 minutes which include changes in the geometry of the satellite constellation. Two GPS receivers were used and connected to the same GPS antenna through a splitter. This set-up ensured that both GPS receivers were receiving the same signals as an input.

The low-cost GPS receiver used is a u-blox LEA-6T with carrier phase measurements capability which is a requirement to be able to use the time-differenced carrier phase positioning technique. The u-blox receiver was configured to make measurements at 5 Hz which is the maximum allowed. Previous experience with this equipment suggests that the 5 Hz measurement rate is not always ensured and can sometimes drop to 4 Hz. However, the time-differenced carrier phase positioning technique is robust to missed measurements or varying measurement rate by its formulation.

The survey grade GPS receiver used is a Leica GS10 which also has carrier phase measurements capability. The Leica receiver was configured to make measurements at 10Hz, its maximum rate.

13.1.2. Methodology

The two GPS receivers were connected to the same GPS antenna through a splitter to ensure that the same signal was received. As the data was collected on a surveyed pillar, the known coordinates of the pillar were used as the initial position for the time-differenced carrier phase positioning algorithm. With the data being

static, the positioning error and its drift rate against the pillar position have been generated and the results presented below.

Let's first define the optimal configuration which gives the best results using the time-differenced carrier phase positioning technique developed as:

- Elevation mask : 20°
- Weighting scheme : Rtklib
- Precise ephemeris + precise clock
- Height constraint : every 100 epochs
- Loop closure

This optimal configuration has been used to compare the performance of the two GPS receivers in Section 13.1.3.

The following methodology has been used to investigate the impact of the different corrections and optimisations individually. Starting from the optimal configuration defined above, only the correction or optimisation investigated is changed so that the impact of that correction or optimisation is clearly isolated. It also allow for direct comparison against the optimal solution.

Let's now detail the different metrics used to assess the positioning quality. Because the height constraint factor is being applied in the optimal configuration, the vertical error becomes irrelevant a part from when the impact of the height constraint itself is being evaluated (Section 13.1.3). For this reason, the horizontal 2D positioning error has been preferred over the 3D positioning error as a first metric to assess the positioning quality. Moreover, the 2D positioning error is more

relevant for the sprint canoeing and kayaking application as it takes place on flat water with a constant altitude by definition. Additionally, the primary interest of the sport is to accurately measure the sprint distance which is also planar by definition. An analysis of the vertical error is however presented in Section 13.1.3 where the vertical component is the actual focus of interest.

The second metric used to assess the positioning quality is the drift rate. This choice is based on the specificities of the time-differenced carrier phase positioning technique which suffers from drift as it is conceptually accumulating small errors inevitably leading to a drift (Section 10.2 and 12.1).

In more details, let's define e_{t_1} the positioning error at t_1 and e_{t_2} the positioning error at $t_2 = t_1 + \Delta_t$. The drift rate of the positioning error between t_1 and t_2 is defined as:

$$\text{drift}(t_1, t_2) = \frac{e_{t_2} - e_{t_1}}{\Delta_t} \tag{13.1}$$

The drift rates have been computed over different periods Δ_t to cover the different race distances and time subdivisions commonly used in the sport.

- 0'10" : One of the shortest split time used in the sport
- 0'30" : Approximately a 200 metres sprint
- 1'30" : Approximately a 500 metres sprint
- 3'00" : Approximately a 1000 metres sprint

13.1.3. Results

Receivers comparison with optimal configuration

This section presents the results obtained using the developed time-differenced carrier phase positioning technique using the optimal configuration. Figure 13.1 presents the positioning error over the 90 minutes of the trial obtained with the low-cost u-blox LEA-6T receiver and the survey grade Leica GS10 receiver.

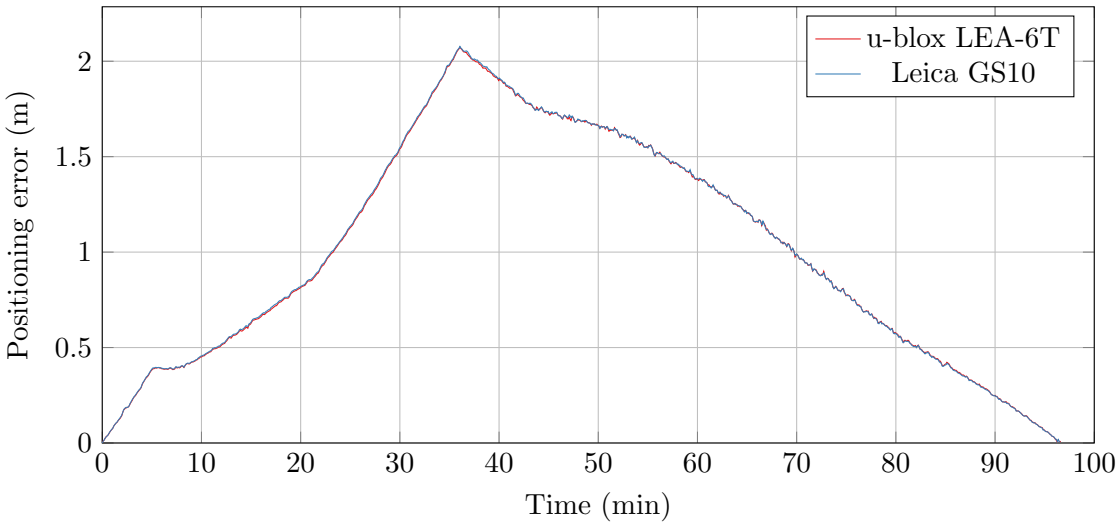


Figure 13.1.: Positioning error in metres obtained with the u-blox LEA-6T and the Leica GS10 receivers

The first significant result is the absence of difference in the positioning quality between the u-blox and the Leica receiver. This is an important finding as a low-cost solution is key to be deployed in the sport environment.

The second result that arises from the figure is the presence of different trends in the positioning error drift rate. These different trends can be explained by changes in the satellite constellation. In order to better quantify the impact of satellite constellation changes on the positioning error, drift rates have been computed. Figure 13.2 presents the drift rate obtained with the low-cost receiver over the different periods Δ_t of interest. The drift rates obtained with the Leica receiver

are not represented for clarity as they are almost identical to the u-blox ones (Table 13.1 and 13.2).

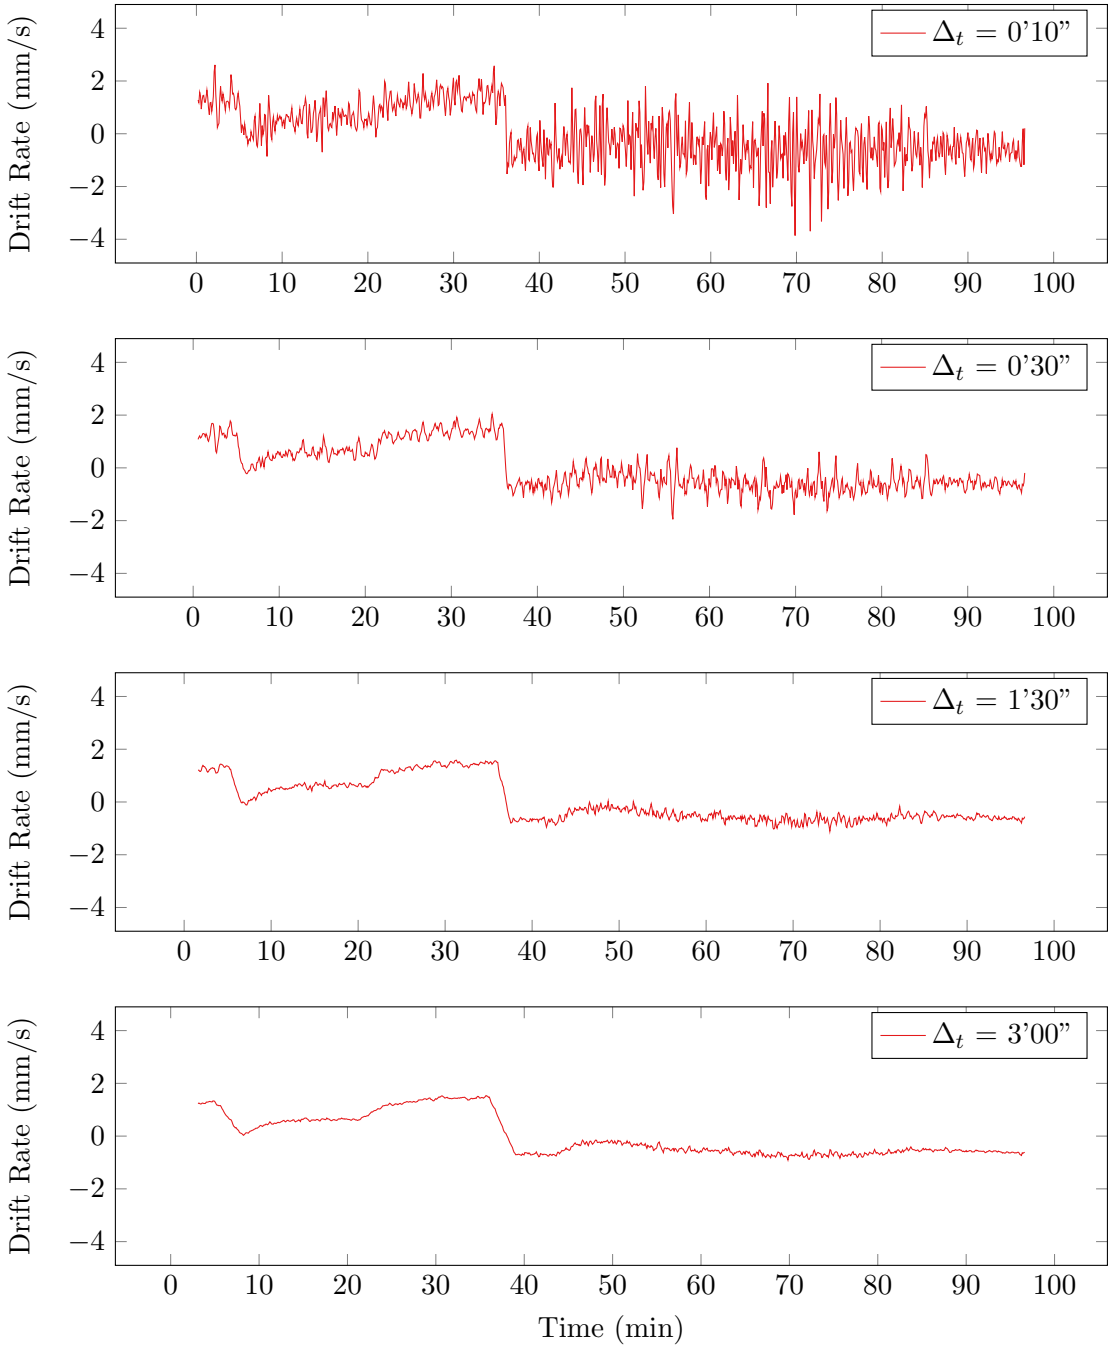


Figure 13.2.: Drift rates of the positioning error in millimetres per second obtained with the u-blox LEA-6T over different periods during the static validation trial

Once again the different trends are clearly visible, especially for periods Δ_t over 30 seconds. An interpretation of the results is that for periods below 30 seconds,

the noise of the receiver is prevalent over the drift induced by the atmospheric propagation error and satellite constellation geometry.

Table 13.1 and Table 13.2 respectively presents a summary of the drift rates obtained with the Leica GS10 and the u-blox LEA-6T receivers.

Table 13.1.: Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the Leica GS10 receiver

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	0.862	0.590	1.980	3.751
0'30"	0.731	0.422	1.527	2.145
1'30"	0.705	0.356	1.448	1.609
3'00"	0.692	0.334	1.430	1.544

Table 13.2.: Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the u-blox LEA-6T

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	0.864	0.598	1.970	4.059
0'30"	0.726	0.423	1.532	2.093
1'30"	0.701	0.357	1.452	1.600
3'00"	0.689	0.335	1.435	1.545

It can be seen that drift rates obtained with the u-blox LEA-6T are comparable to the drift rates obtained with the Leica GS10 GPS receiver. This supports the capability of using a low-cost GPS receiver with carrier phase measurements for time-differenced carrier phase positioning.

Additionally, Table 13.1 and Table 13.2 reflect the sensitivity of drift rates to the receiver measurement's noise for short time periods. This sensitivity decreases

with larger time periods where the errors due to atmospheric propagation and the satellite constellation geometry takes over.

Loop Closure

The first option to be investigated is the impact of the loop closure constraint on the positioning error. Whilst the idea and implementation behind loop closure is simple, it represents a significant contribution to constrain the drift rate. Figure 13.3 presents the 2D positioning error obtained with and without loop closure.

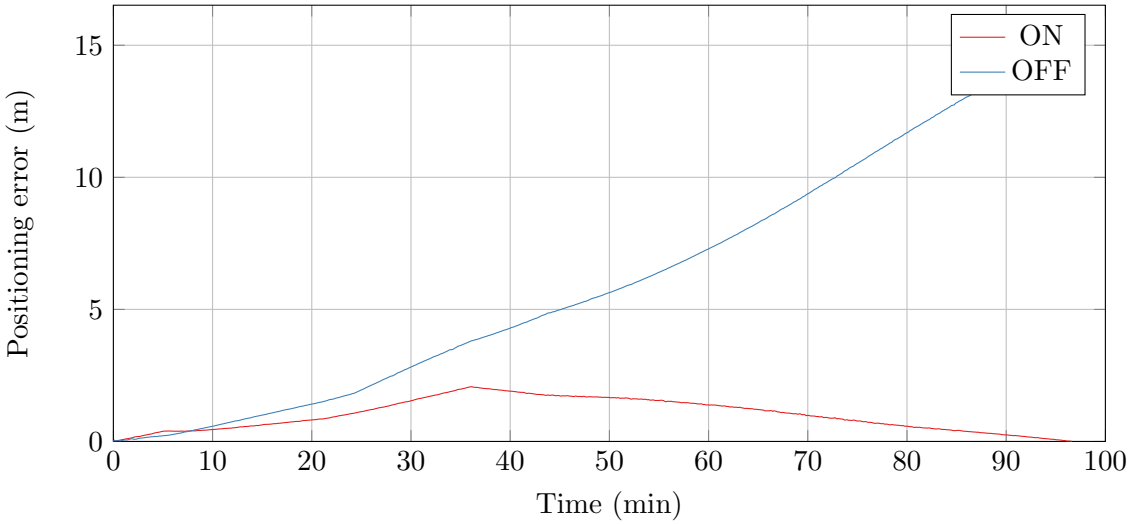


Figure 13.3.: Positioning error in metres with and without the loop closure constraint

By using the loop closure constraint, the maximal positioning error is reduced from 15.02 metres to 2.07 metres, a significant improvement. As detailed in Section 12.3.2, the loop closure constraint fully benefits from having all the states as part of the optimised problem.

To further highlight the benefits of the loop closure constraint, the drift rates have been computed and are presented in Table 13.3. It can be seen that the magnitude of the drift rates also significantly degrades without the loop closure

Table 13.3.: Mean (μ), standard deviation (σ), 2-sigma (2σ) and maximum drift rates in millimetres per second over different periods (Δ_t) obtained with the u-blox LEA-6T without loop closure

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	2.550	1.122	4.501	7.393
0'30"	2.582	0.978	4.029	4.880
1'30"	2.605	0.939	3.894	4.250
3'00"	2.622	0.923	3.863	4.050

constraint.

Elevation Mask

The second option investigated is the elevation mask and its impact on the positioning quality. Whilst a 15° mask is quite commonly used for traditional GPS processing, it may not be the optimal value for time-differenced carrier phase positioning. More precisely, it is known that atmospheric perturbations are larger for measurements made on satellites with lower elevations. Whilst this already represents a challenge for real-time kinematic processing where a base station is used to correct these perturbations, it is even more sensitive in the context of time-differenced carrier phase positioning where no external corrections are used and errors accumulated. This would inevitably lead to an increased drift rate and thus a reduced positioning quality.

Figure 13.4 presents the results obtained with the following elevation masks : 10°, 15°, 20° and 25°.

The best results are obtained with an elevation mask of 20°. This suggests that lower elevation satellites have a significant contribution to the drift rates as it can

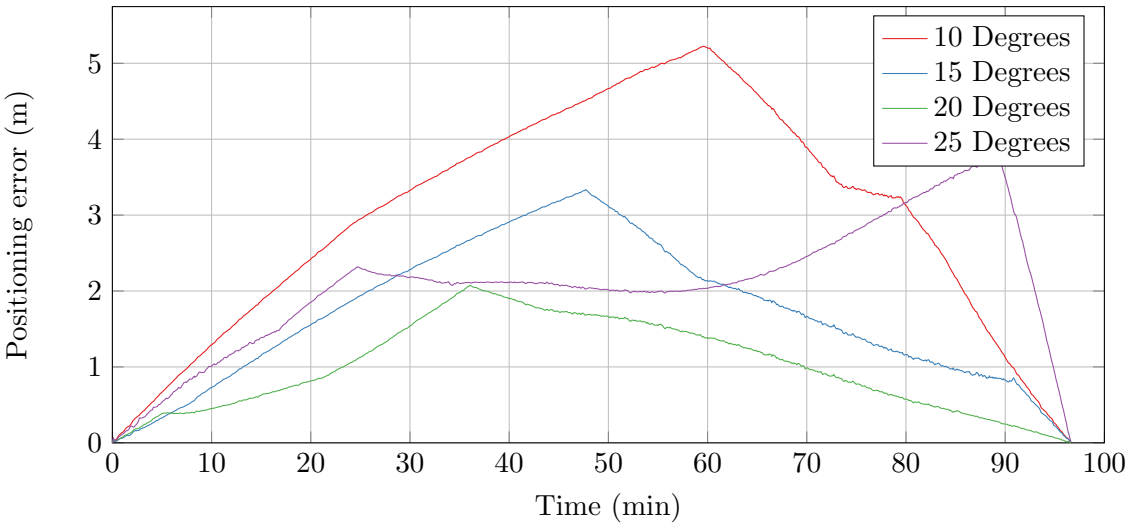


Figure 13.4.: Elevation Mask

be seen with the results obtained with 10° and 15° masks. However, increasing the mask to 25° results in a decreased positioning quality due to the limited number of satellites remaining with only five used for most of this trial.

Finally, the effects of satellite constellation changes are clearly visible with different change times depending on the elevation mask applied leading to different drift trends.

Weighting Scheme

This section presents the results obtained with different weighting schemes whilst using time-differenced carrier phase positioning. As detailed in Section 12.1.3, the weighting scheme helps to reduce the impact of measurements made on lower elevation satellites more prone to be affected by larger noise and errors.

Figure 13.5 presents the results obtained with different weighting schemes and without any. Three weighting schemes have been investigated, with two based on elevation and one on the signal-to-noise ratio.

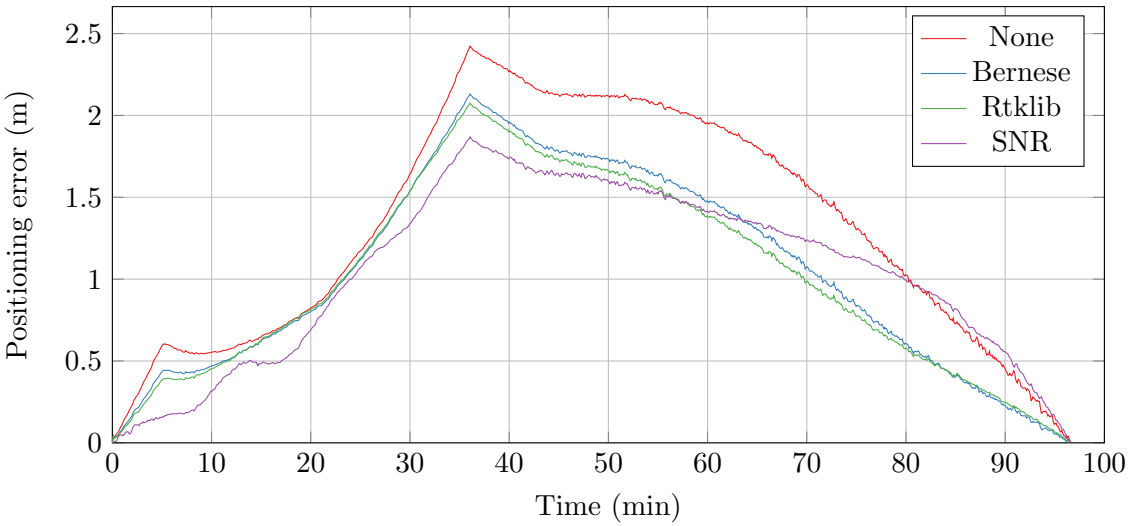


Figure 13.5.: Positioning error in metres obtained with different weighting schemes

Amongst the different elevation based weighting schemes, the best results are obtained with the Rtklib one. Regarding the SNR weighting scheme, good performance is also obtained with a positioning quality fairly similar to the Rtklib scheme.

Elevation based weighting schemes have the advantage to be always applicable whilst the SNR based one can only be applied with receivers outputting the signal to noise ratio.

Height Constraint

The next option investigated is the height constraint which has been presented in Section 12.3.1. As mentioned in the preamble of this chapter, the 3D positioning error will be considered in this section. Figure 13.6 presents the 3D positioning error obtained with and without applying height constraint. In addition, several constraint rates have been tested with :

- One constraint every measurement

- One constraint every 100 measurements

- One constraint every 1000 measurements

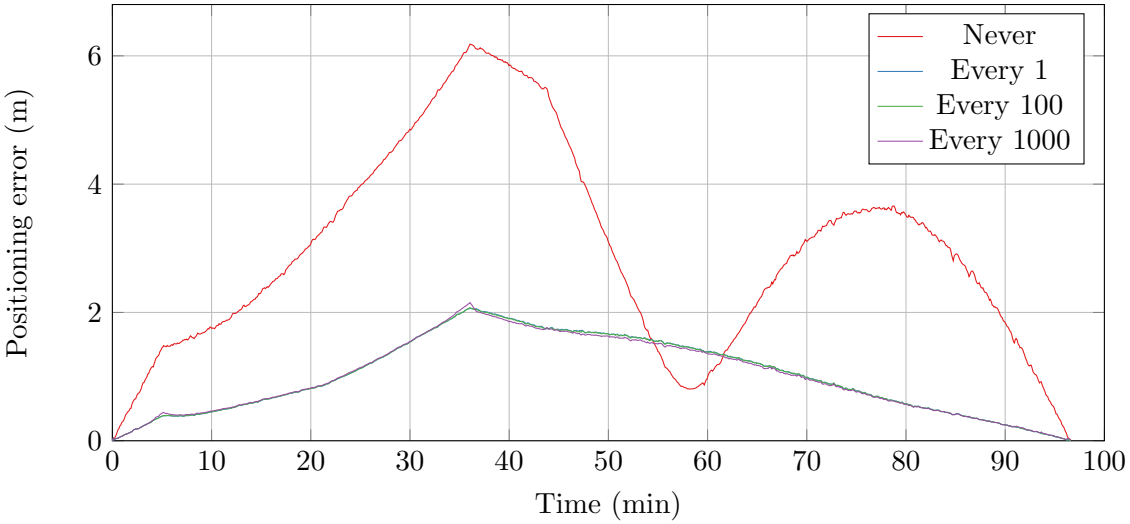


Figure 13.6.: 3D Positioning error in metres obtained with height constraint factors inserted at different rates

First of all, let's highlight that the 3D positioning quality, when height constraint is applied, is comparable to the 2D positioning accuracy presented above with the optimal configuration. This can be explained by the fact that height constraint factors force the vertical error to be close to zero. However, it is clear that the 3D positioning accuracy degrades when height constraint factors are not applied.

Figure 13.7 shows the vertical positioning error only. It can be seen that when height constraint is applied, the vertical error is close to zero. Additionally, it tells us that the biggest contribution of the 3D positioning error without height constraint originates from the vertical component.

Loss Function

This section presents the results obtained by using loss functions as described in Section 9.3.3 and 12.1.2. The role of a loss function is to scale the residuals to

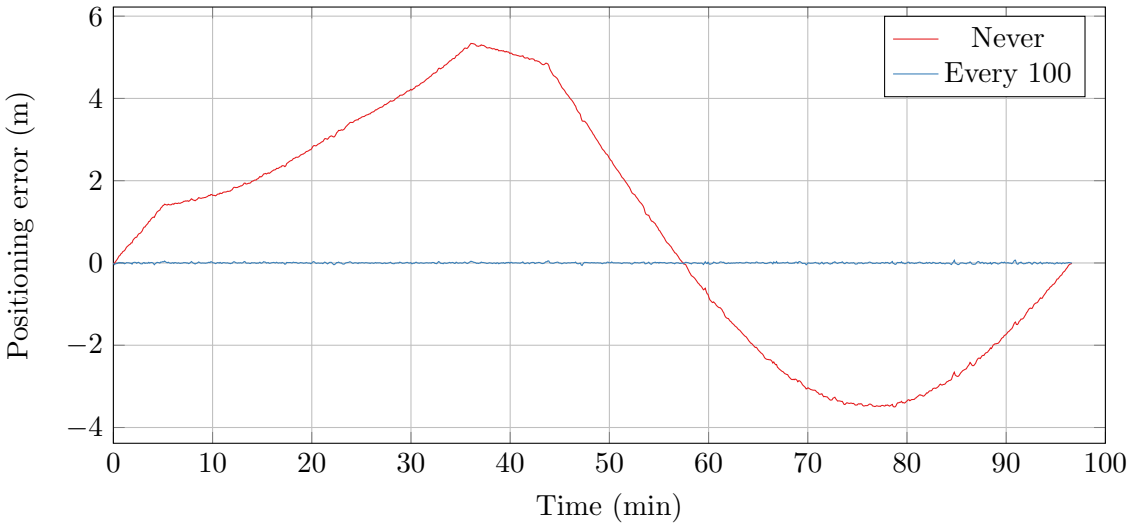


Figure 13.7.: Vertical Positioning error in metres obtained with and without height constraints

reduce the cost of large residuals. As a consequence, it makes the optimisation more robust to outliers.

As detailed in Section 12.1.2, a Cauchy loss function has been applied to measurements made on satellites with an elevation below 30° as they are more prone to be affected by outliers.

Figure 13.8 presents the results obtained with and without loss function applied.

As it can be seen, there is no significant difference between the two configurations with results even slightly better without loss functions applied. This can be explained by the fact that this validation dataset has been obtained in near optimal conditions with a static antenna in a perfect clear-sky environment thus limiting the presence of outliers. However, it will be shown in Section 13.2.6 that loss function becomes valuable in more challenging environments such as on-water kayaking.

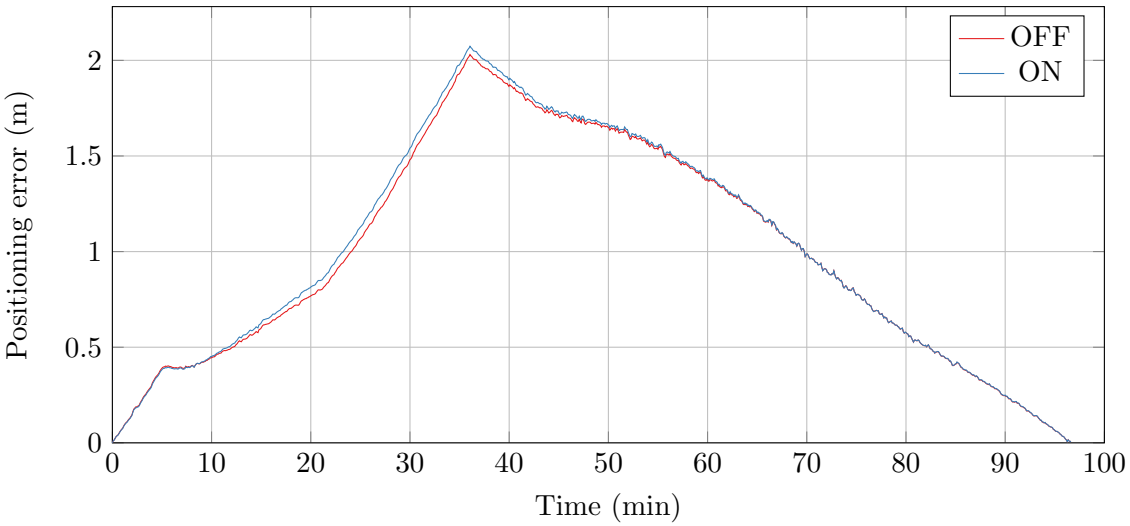


Figure 13.8.: Positioning error in metres obtained with and without the use of loss function

Precise Ephemeris

This section presents the results obtained with different types of ephemeris and clock products (Table 13.4).

Table 13.4.: Different types of ephemeris and clock products investigated and their availability

Product	Latency	Ephemeris	Clock
Broadcast	Real-Time	✓	
Ultra-Rapid	Real-Time	✓	
Rapid	17-41 hours	✓	✓
Final	11-17 days	✓	✓

The broadcast ephemeris is available as a standard on a GPS receiver as it is the one used on-board to obtain a position solution. The IGS products namely, Ultra-Rapid, Rapid and Final have different latency with only one, the Ultra-Rapid product, available in real-time. As all the IGS products are accessible online, if the Ultra-Rapid product was to be used in real-time, an internet connection would be required so it can be downloaded.

Figure 13.9 presents the results obtained with the different ephemeris and clock products.

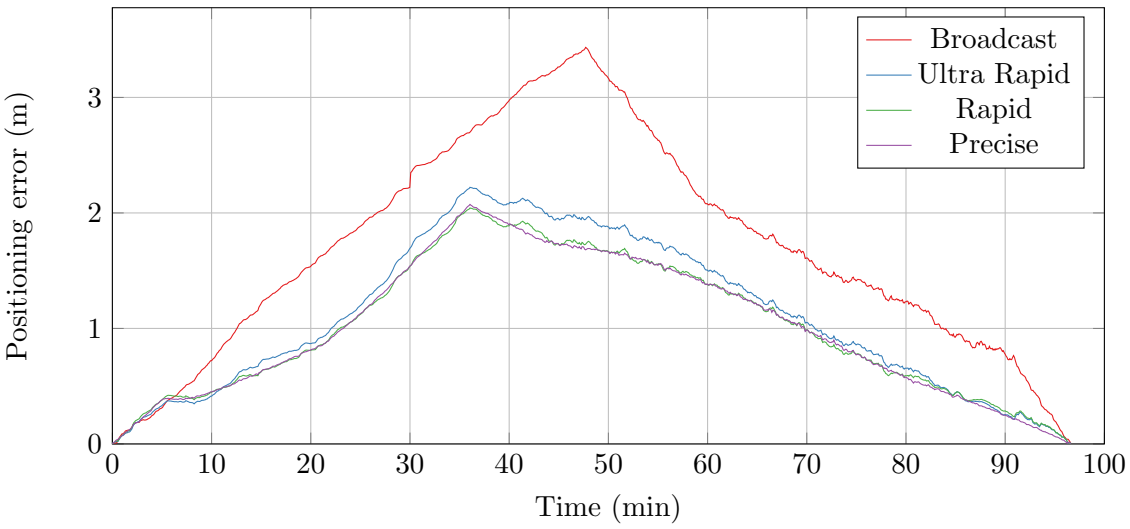


Figure 13.9.: Positioning error in metres obtained with and without the use of precise satellite ephemeris and clocks

It can be seen that IGS products improve the positioning quality as they get more refined. However, the Ultra-Rapid product is already a significant improvement compared to the broadcast ephemeris. However, the difference between the results obtained with the Rapid and Final products is limited, with reduced local fluctuation using the Final product.

13.2. On-water sprint kayak validation

Following the GPS practical validation presented in Section 13.1, the next step was to undertake a practical validation of the complete GPS + IMU system with on-water sprint kayak trials. This validation is the final validation of the work which will validate the performance of the system and algorithms developed in the context which they have initially been designed for.

13.2.1. Trials and location facility

The on-water trials were conducted at Dorney Lake, England. This facility, also known as Eaton Dorney, was used during the 2012 Olympic Games for the rowing, sprint canoe and sprint kayak events. The lake has a 2.200 km straight length with still water offering consistent water conditions. It has 8 rowing lanes, 13.5 metres wide each.

The trials were supported by the English Institute of Sport providing access to a K-1 kayak racing boat with an elite athlete. Thanks to the athlete and its coaches, the trial session was entirely designed to support the technical trials.

The complete trial session lasted 30 minutes including six sprints of different length at medium intensity organised as :

Table 13.5.: Sprint efforts of the on-water trial session and their associated distance

Sprint	Distance
Sprint 1	100 metres
Sprint 2	100 metres
Sprint 3	100 metres
Sprint 4	500 metres
Sprint 5	100 metres
Sprint 6	200 metres

13.2.2. Commercial Equipment

The on-water trial had several objectives which dictated and orientated the choice of the equipment used. First of all, it targeted to collect some gold-standard reference data. The motivation behind this was to collect high quality data, regardless

of price, size and weight constraints to present to the English Institute of Sport what can be best obtained.

The second purpose was to collect some data to be processed by the technique developed in this work. A system has been specifically developed for that purpose and is presented in details in the next Section (13.2.3). Importantly, having the developed system running aside the gold-standard reference allowed to have a quality control on the developed solution.

Finally, two commercial devices were also used for quality assessment purposes with the results presented in Section 8.1. The complete set-up can be seen on Figure 13.10.

As the gold-standard reference solution, a NovAtel SPAN system has been used. This choice was mainly directed by the equipment available at the Nottingham Geospatial Institute and its form factor. The SPAN system is composed of a survey-grade GNSS receiver [96] associated with a tactical grade IMU [97].

The survey-grade GNSS receiver is dual-frequency (L1 + L2), multi-constellation (GPS + Glonass) with a 1 Hz update rate. It was connected to a high-performance NovAtel Pinwheel antenna [99] which supports multi-constellation as well as offering strong multipath rejection.

The data was post-processed using the commercial software Inertial Explorer [98] which is associated with the SPAN system. Inertial Explorer is a reference in the world of GPS+IMU processing and will confidently provide a reference solution that can be used to assess the quality of the other systems. The positioning quality of the reference system is below the centimetre level. Regarding orientation, angles are estimated with an accuracy of a few $1/100^{\text{th}}$ of a degree.



Figure 13.10.: Fully equipped K1 boat for the on-water trials

13.2.3. Developed system

This section describes the hardware used to go along with the processing solution developed in this work. Different constraints have orientated the hardware choices. The general constraints were to come up with a low-cost solution which could be used by the sport. Due to the application, it also required to be lightweight and as a consequence power-efficient as batteries are heavy. Additionally, a small form factor was targeted so the system can easily be mounted on a kayak racing boat with very limited space.

Due to the processing technique and algorithms developed in this work, there was a need for a system combining a GPS receiver with an IMU. Regarding the GPS, the principal criteria was to use a receiver with raw carrier-phase measurement capabilities. Furthermore, as the IMU has its own internal clock to timestamp the measurements, a GPS receiver with timing capability is required to relate the IMU time with the GPS time. These two requirements added to the low-cost constraint resulted in the choice of a u-blox LEA-6T GPS receiver which was used in its evaluation kit form factor. The u-blox LEA-6T has raw carrier phase measurement capability and offers the possibility to timestamp event with the GPS time by sending impulses on a specific pin of the receiver.

The choice for the second main component of the system, the IMU, was to use a MicroStrain 3DM-GX3-25 inertial measurement unit [91]. This choice was guided by the availability of the system at the Nottingham Geospatial Institute alongside its cost, quality and form factor. Additionally, successful previous work experiences with this sensor strengthened the decision.

Whilst the choice of the sensors is critical, their connection and interaction is also of primary importance especially in terms of timing. GPS and IMU measurements need to be precisely time synchronised so they can be processed together to generate an inertial navigation solution. If the two sensors are not precisely synchronised, the IMU measurements will not relate to the GPS measurements made at the same time. In order to obtain a clock synchronisation and also add a data logging feature, it was decided to use a Raspberry Pi [39] as the central piece of the system. The Raspberry Pi is a low cost, credit card sized computer which was primarily designed to support eduction by providing an affordable and easy access to computers for children all over the world. The device is also widely used by the do-it-yourself community as it also is a flexible platform with different connection capabilities that can be used to plug external sensors or cards to add extra functionalities to the system. This is also the reason why it was used in this work thanks to the availability of USB ports and GPIO pins on the board. The USB ports were used to plug both the GPS and IMU sensors and log their measurements. Additionally, the GPIO pins were used to synchronise the IMU clock with the GPS time utilising the timing capability of the u-blox LEA-6T GPS receiver. Software written in C was specifically developed to configure, synchronise and log the data from the GPS and IMU sensors. The output produced by the logger is saved on the Raspberry Pi's SD card and results in two files, one with GPS measurements and one with IMU measurements. The GPS log is formatted using the binary u-blox format (.ubx) whose specifications are available online [133]. The u-blox binary log can be converted in the standard Receiver Independent Exchange (RINEX) format [46] using different tools such as TEQC [135] or RTKLIB [125]. On the other side, the IMU log is formatted as a text file containing the

accelerometers and gyroscopes measurements timestamped with GPS time. The GPS time-stamping is obtained through the event timing capability of the GPS receiver as described above.

The whole system was powered using a single USB power bank that was linked to the Raspberry Pi. Both GPS and IMU sensors were powered through their USB connection with the Raspberry Pi. The complete system has a very low power consumption which allowed it to run for more than a day using a single USB power bank.

Regarding weight, the developed system is very light with the two critical components being the USB power bank and the GPS survey grade antenna. However, both elements are still in the hundreds of grams range. If weight was more critical, a smaller USB power bank could be used as well as a different GPS antenna. Survey grade antennas are bulkier and heavier as they embed a choke ring element. However, there exists other antenna designs such as helix antennas which are particularly suitable for water environments.

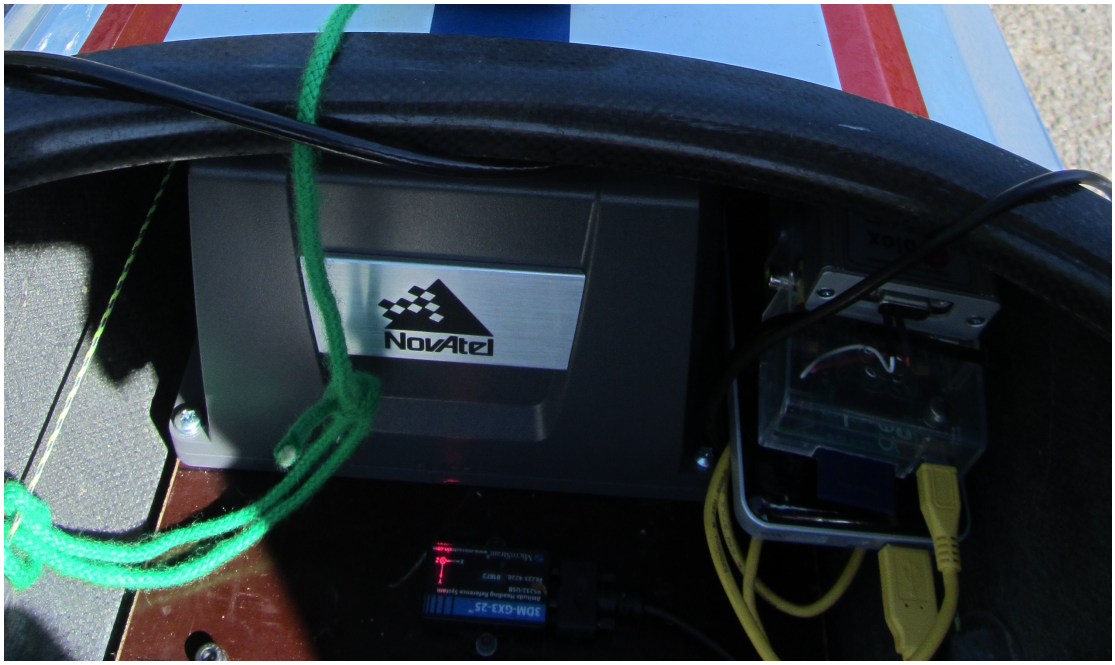


Figure 13.11.: Overview of the developed system installed in the K1 boat

Figure 13.11 shows the developed system installed in the K1 boat just before it

was sealed into waterproof bags. The Microstrain IMU can be seen at the bottom centre of the picture. On the right hand side, the rest of the system is stacked together. At the bottom is the USB power bank powering the Rasbperry PI, seating just on its top, with the yellow USB cable. On top of the Raspberry PI, the u-blox LEA-6T GPS receiver can be distinguished with thin red and white cables used for the time synchronisation of the IMU time with the GPS time.

13.2.4. Data collection

To support the trials, the English Institute of Sport and GB Canoeing lent a K-1 racing boat a few days prior to the trials. This helped to find a mounting protocol of the equipment in the kayak including the design and manufacturing of a wooden platform to support the inertial sensors of the reference and developed system. The wooden platform holding the IMU was mounted in the hull just behind the seat of the athlete and securely screwed to the boat using the seat adjustment rails. It was decided to use a thick wood platform to ensure that the two inertial sensors were attached to a rigid body so that their measurements can be more closely related and compared.

A single survey-grade NovAtel Pinwheel antenna [99] was used for both the reference and the developed system. The antenna was connected to both systems through the use of a GPS splitter. The advantage of this set-up was to have two GPS receivers of different quality (survey-grade vs low-cost) receiving the exact same signal allowing for a direct comparison of their performance. The antenna was attached on the rear deck of the boat about 1 metre behind the athlete. The lever arms between the two inertial sensors and the GPS antenna were determined using a Leica total station.

The two commercial devices were also mounted on the rear deck of the kayak between the seat and the survey-grade GPS antenna. These devices need to have

a clear sky environment as the GPS antenna is embedded in the casing.

All the equipment of the reference and developed system, including batteries and data-loggers but excluding the GPS antenna, were put in dry bags to offer water protection. The two commercial devices were designed for this application and thus are waterproof.

An extended static period was observed at the beginning and end of the trial to ensure the perfect initialisation of the reference system based on a fine alignment procedure (Section 11.5.1) using stationary data.

13.2.5. Methodology

The methodology used for quality assessment during the on-water trials is similar to the one used for the static trials detailed in Section 13.1.2.

The ground truth reference used is the inertial navigation solution obtained by post-processing the data of the NovAtel SPAN system with Inertial Explorer.

For each variable of interest, position, velocity or attitude, the estimates obtained by the solution developed in this work have been compared to the estimates obtained by the reference system and presented as an estimate error. The first variable of interest is the quality of the positioning which can be seen as the continuity of Section 13.1 with dynamic data this time. The second variable investigated is the quality of the velocity estimation using the combined GPS-inertial solution. Finally, the quality of the attitude estimation is presented.

The position used to initialise the time-differenced carrier phase positioning algorithm has been extracted from the NovAtel SPAN solution. This has the advantage of enabling direct comparison with the reference solution as it started from the

same point and made any drift clearly visible. As a reminder, the developed solution is not expected to have such a high absolute positioning quality for its initial position as it generally uses the standalone C/A code GPS solution which has an accuracy of a few metres rather than centimetres for its initialisation. However, this would not significantly affect the positioning quality as it mainly means that the solution is translated by a few metres.

This part of the validation focuses on a quality assessment regardless of the sport context. As such, the results presented below summarise the system's performance over the whole 30 minutes of the trial. A sport specific analysis is presented in Chapter 14 where the performance of the system is evaluated during the different sprint efforts of the trial.



Figure 13.12.: Fully equipped K1 boat during the on-water validation trial

13.2.6. Results

Positioning error

In order to assess the positioning quality of the developed time-differenced carrier phase technique, the IMU data was not used in this section. As a result, the positions obtained correspond to the GPS antenna location. For this reason, the reference solution of the NovAtel SPAN system has been translated to the GPS antenna location to be used as a reference in this section. An overview of the position trace during this trial is shown on Figure 13.13.

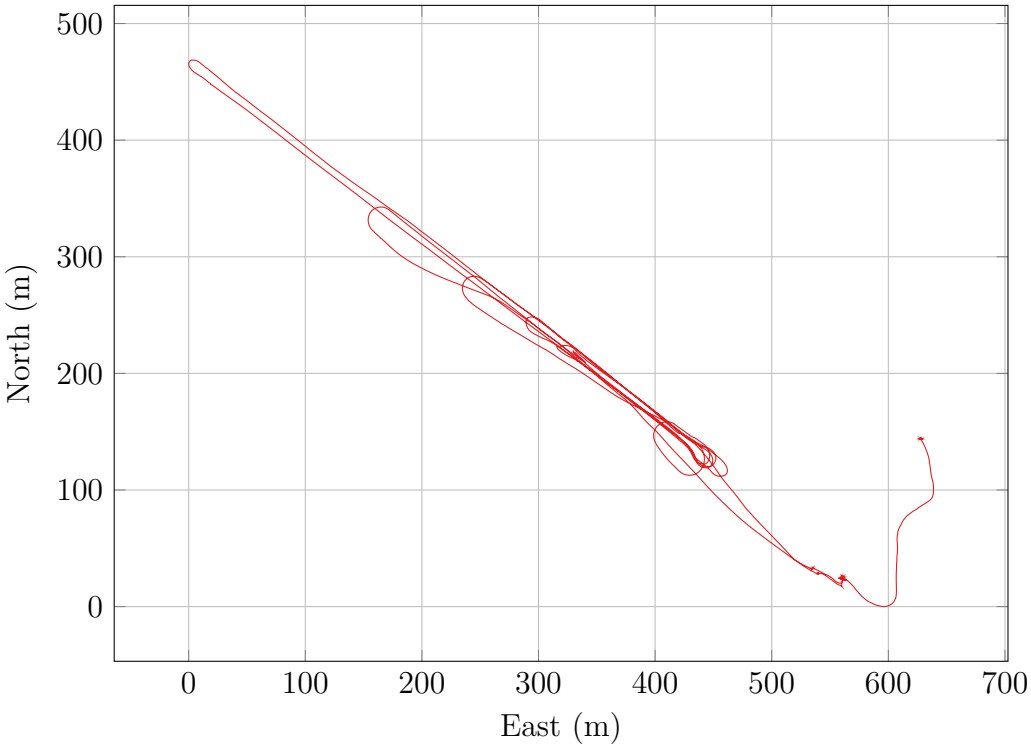


Figure 13.13.: Position trace of the on-water validation trial

Developed solution vs RTK This section presents the results obtained by the developed GPS processing with the optimal configuration (Section 13.1.2) using the u-blox LEA-6T receiver.

Additionally, a traditional real-time kinematic (RTK) solution has been generated

with Rtklib [125, 126] using the same u-blox LEA-6T data and a base station. This will highlight the performance of the developed technique against traditional RTK using the same data.

Figure 13.14 shows the results obtained with the two different processing techniques. The positioning accuracy of the developed solution reported a mean horizontal error of 4.9 cm with a standard deviation of 2.9 cm and a maximal error of 20.6 cm over the 30 minutes of the trial. The positioning quality obtained by the developed solution challenges the RTK processing without the need for a base station.

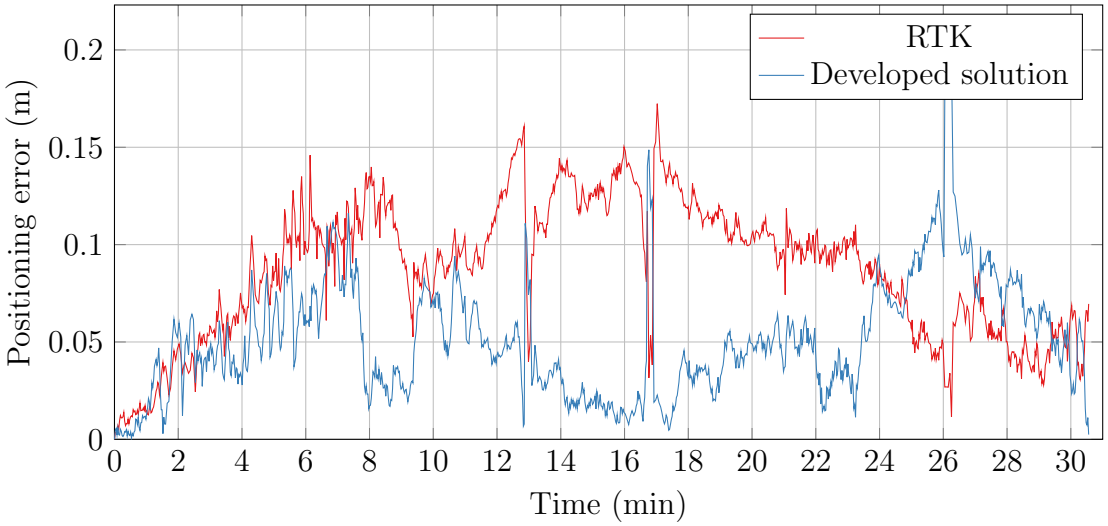


Figure 13.14.: Positioning error in metres obtained by the u-blox LEA-6T with the developed time-differenced carrier phase positioning technique and the RTK solution

Noticeable spikes appear for both solutions around the following times : 13 minutes, 17 minutes and 26 minutes. After investigation, these spikes are due to a fault in the u-blox data logging explaining why they affect both processing techniques.

Similarly to Section 13.1.3, drift rates have been computed over different time periods with results presented on Figure 13.15. In order to avoid biasing the results, the data where the logging bug is present has been omitted which explains

the flat line sometimes visible on Figure 13.15.

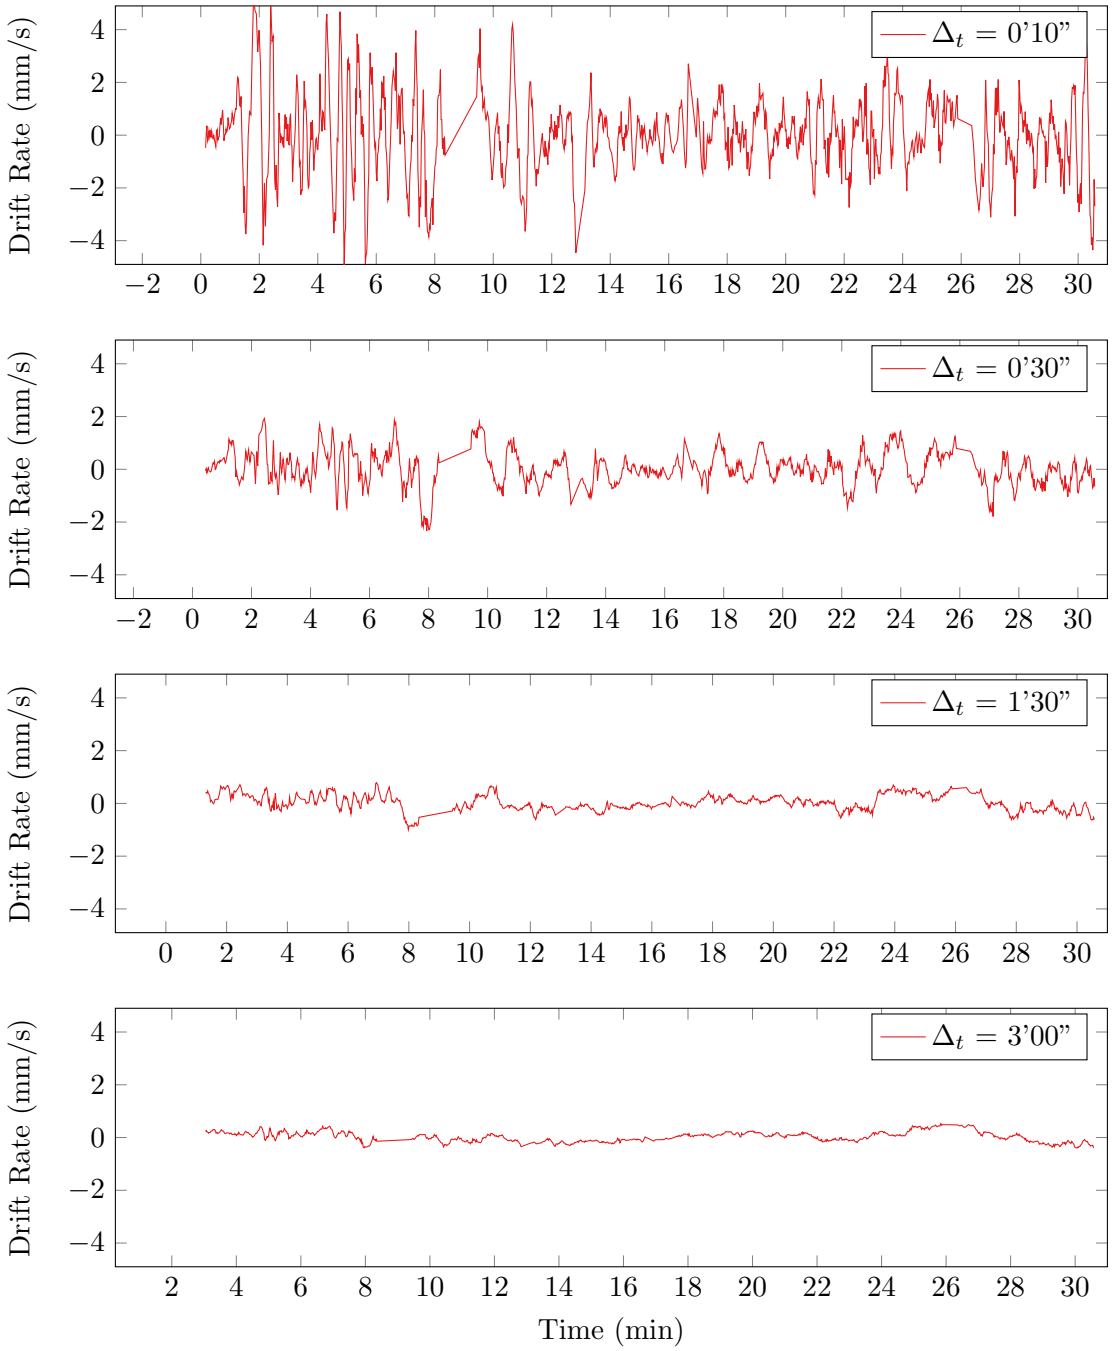


Figure 13.15.: Drift rates of the positioning error in millimetres per second obtained with the u-blox LEA-6T over different periods during the on-water validation trial

As it can be seen the drift rate over a 10 seconds period is noisier than what was obtained in static condition (Section 13.1.3). This can be explained by the more challenging environment, involving dynamics and an increased likelihood

of multipath reflection due to the water. However, similarly to the findings of the static validation, the noise contribution decreases for time periods over 30 seconds.

Additionally, the reader may notice that the results do not exhibit different drift trends as opposed to the static trial. This can be explained by the shorter duration of the trial (30 minutes vs 90 minutes) which results in a less varying satellite constellation.

Table 13.6 and 13.7 present summaries of the drift rates obtained with the developed technique and the traditional RTK processing.

Table 13.6.: Drift rates in millimetres per second over different periods obtained with the u-blox LEA-6T with the developed technique

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	1.079	0.990	3.244	6.235
0'30"	0.497	0.423	1.391	2.865
1'30"	0.233	0.182	0.597	1.068
3'00"	0.155	0.105	0.377	0.554

Table 13.7.: Drift rates in millimetres per second over different periods obtained with the u-blox LEA-6T with RTK processing

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	0.931	0.916	2.919	7.022
0'30"	0.423	0.361	1.210	2.285
1'30"	0.204	0.150	0.493	1.076
3'00"	0.143	0.102	0.334	0.622

It can be seen that drift rates of the same magnitude are obtained with both techniques. With a deeper look, the developed GPS processing results in slightly larger drift rates compared to traditional RTK processing without it being a significant

difference.

It is also interesting to compare these drift rates to the ones obtained in the static validation trial (Table 13.2). As mentioned above, drift rates over the 10 seconds period are larger in the dynamic trial. This can be explained by the more challenging environment. This is nicely reflected by the σ, 2σ and Max measures which are increasing in this order whilst the mean μ is almost comparable to what was obtained in the static trial.

In other words, whilst the average drift rate over 10 seconds is comparable between static and dynamic data, the later is more likely to be affected by outliers resulting in a larger drift rate due to the more challenging environment.

Interestingly, this statement does not hold when considering periods over 30 seconds with drift rates obtained in the dynamic trial being smaller than during the static one. This may sound surprising at first but it can actually be explained by the duration of the trial. The static data collection was 3 times longer than the dynamic trial. This has consequences when applying the loop closure constraint as the shorter the trial is, the bigger the constraint represented is. This shorter duration also results in fewer changes in the satellite constellation. All these differences result in a problem that is better globally optimised.

Loop Closure Similarly to the static validation trial, the first option investigated is the impact of the loop closure constraint as it has a significant impact on the positioning error. Figure 13.16 presents the positioning error obtained with and without loop closure.

As expected, the positioning error increases when the loop closure constraint is not applied with a maximal positioning error increasing from 0.21 metres to 1.37 metres.

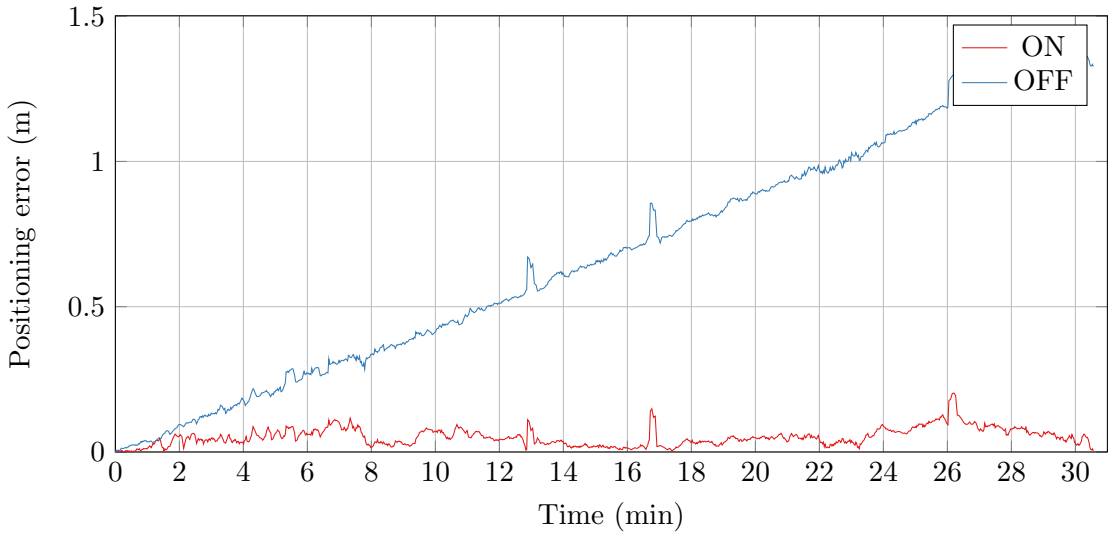


Figure 13.16.: Positioning error in metres with and without the loop closure constraint

Table 13.8 presents a summary of the drift rates obtained without loop closure.

Table 13.8.: Drift rates millimetres per second over different periods obtained with the u-blox LEA-6T without the loop closure constraint

Δ_t	Drift rates (mm/s)			
	μ	σ	2σ	Max
0'10"	1.109	0.942	2.892	6.474
0'30"	0.785	0.484	1.715	2.944
1'30"	0.745	0.231	1.111	1.778
3'00"	0.753	0.144	0.994	1.297

Interestingly, the drift rate over the 10 second time period is unchanged which highlights that, in this situation, the drift correction resulting from the loop closure constraint is negligible compared to the measurement noise over 10 seconds. However, this is not longer true for periods over 30 seconds where drift rates are larger and thus highlighting the benefits obtained from the loop closure constraint.

As a side note, a similar trend was obtained in the static validation trial, with the

average drift rates for periods over 30 seconds being almost constant when loop closure was not applied.

Weighting Scheme The next option investigated is the impact of the different weighting schemes. Findings from the static validation trial (Section 13.1.3) suggested that Rtklib was the best elevation based weighting scheme with performance almost comparable to the SNR based one. Figure 13.17 presents the results obtained this time with the dynamic data.

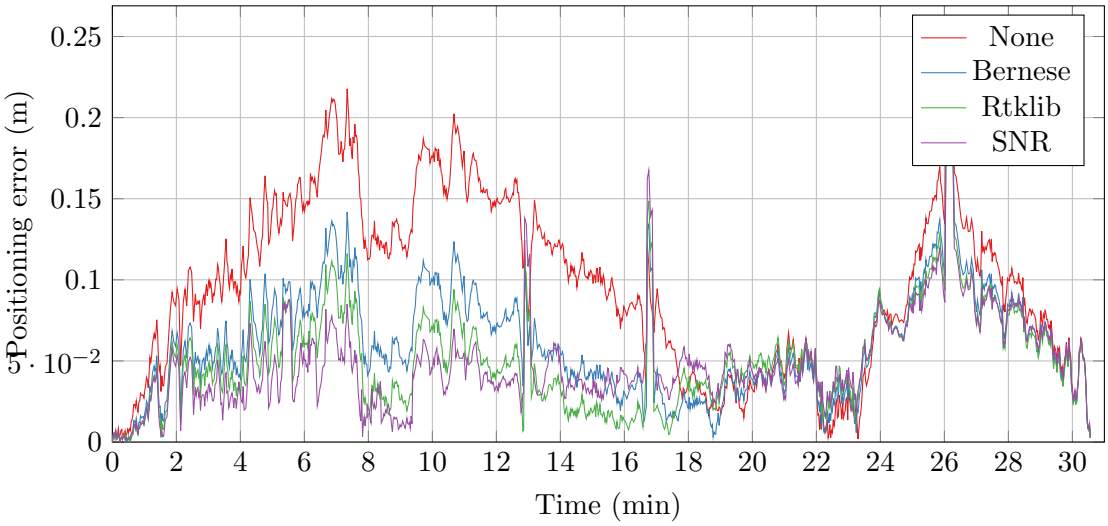


Figure 13.17.: Positioning error in metres obtained with different weighting schemes

The results are supporting the findings of the static validation trials with Rtklib being the best elevation based weighting scheme. Similarly, the SNR based one is again the best option if the signal-to-noise ratio is available from the GPS receiver.

Elevation Mask Following the same investigation order, the next focus is on the elevation mask. Findings from the static validation trials suggested that the positioning quality improved with elevation masks up to 20° before degrading with higher masks due to less satellites available and a weaker satellite geometry. Figure 13.17 presents the results obtained during the on-water validation trial.

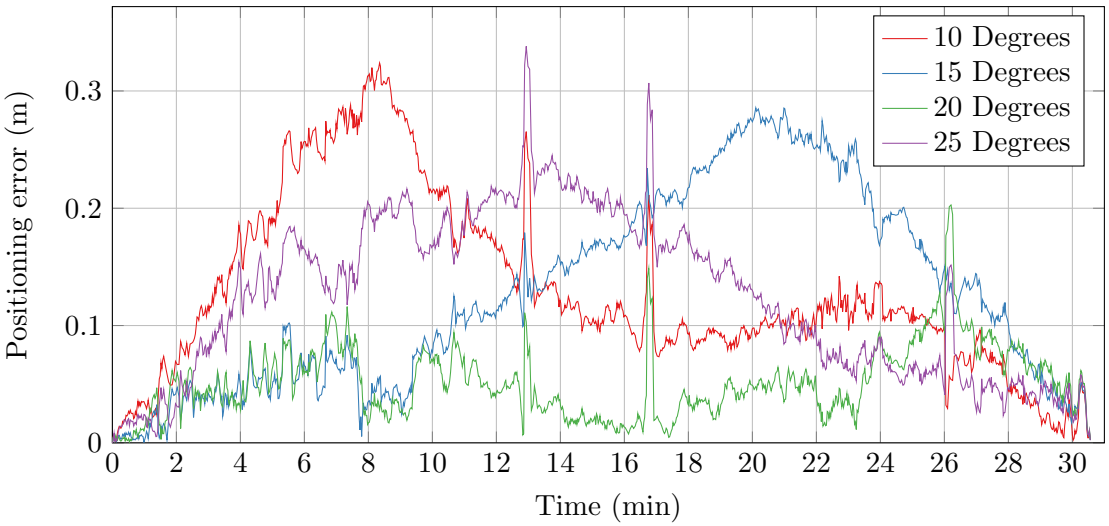


Figure 13.18.: Elevation Mask

The results confirm once again the findings of the static validation trial with the optimal elevation mask being 20°.

Height Constraint The next option investigated is the impact of the height constraint. As a reminder, this is the only section focusing on the 3D positioning error and its vertical component. Figure 13.19 presents the results obtained with the height constraint applied at different rates.

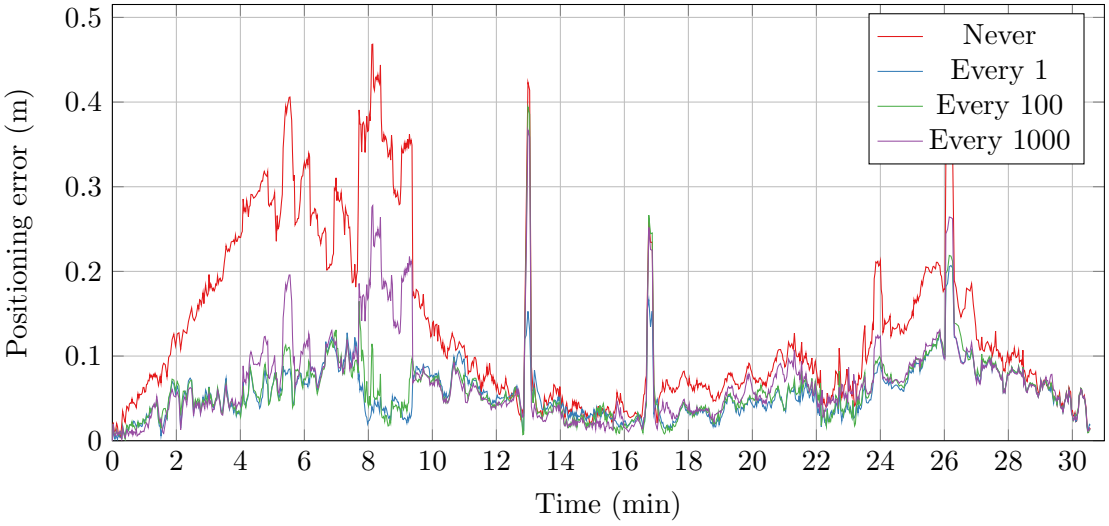


Figure 13.19.: 3D Positioning error in metres obtained with different height constraints

Similarly to the static validation trial, the 3D positioning error is improved with higher insertion rates for height constraint factors. Best and almost equivalents results are obtained when inserting height constraint factors every 1 or 100 measurements. Dropping the insertion rate to every 1000 measurements does result in a noticeable degradation of the positioning accuracy.

Figure 13.20 presents the results of the vertical positioning error. It indicates that most of the 3D positioning error is originating from the vertical component.

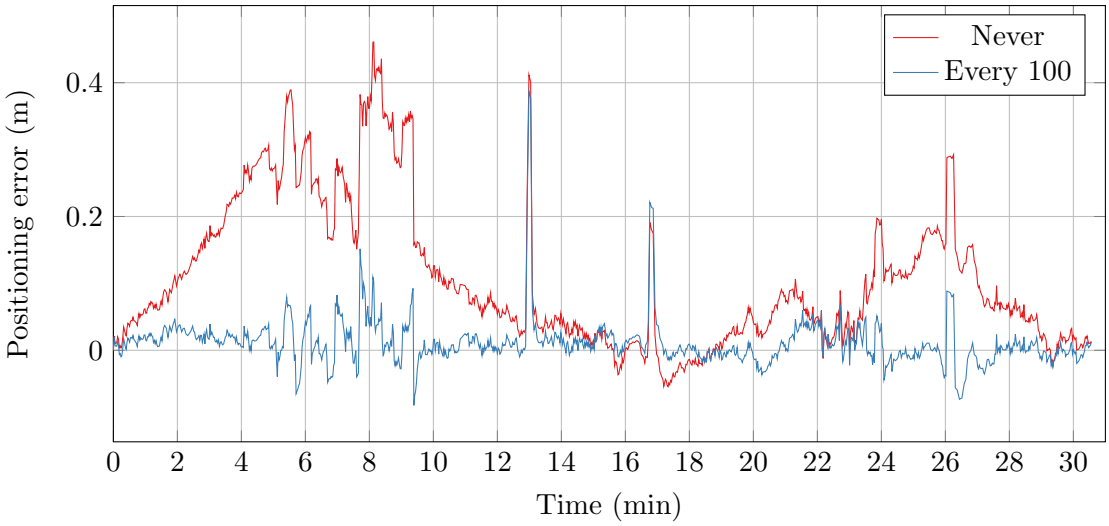


Figure 13.20.: Vertical Positioning error in metres obtained with and without height constraints

These findings confirm what emerged from the static validation trial.

Loss Function This section focuses on the impact of applying loss functions for measurements made on satellite with an elevation lower than 30°. The static validation trial did not show any improvement when loss functions were applied (Section 13.1.3). Figure 13.21 presents the results obtained with the dynamic data this time.

Contrarily to the static validation trial, applying a loss function to measurements made on satellite with an elevation below 30° is this time beneficial. This can

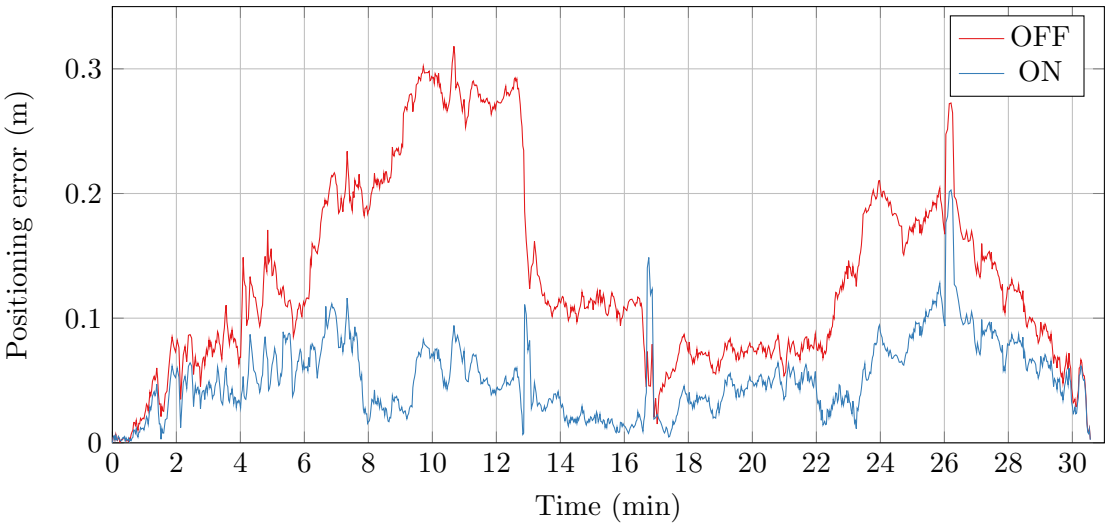


Figure 13.21.: Positioning error in metres obtained with and without the use of loss function

be explained by the very different environment of the two trials. The first one was made in a static, perfect clear-sky visibility environment which greatly limits the risk of outliers and measurements affected by multipath. The second one was made on a boat, rolling by $\pm 15°$ and thus an increased likelihood of multipath. This means that the static dataset was unlikely to contain lots of outliers to be scaled by the loss function contrary to the dynamic one.

Precise Files The last focus is on the impact of using the precise ephemeris and precise satellite clocks instead of the broadcast one. Results are presented on Figure 13.22.

It can be seen that the best results are obtained with the Final IGS products, closely followed by the Rapid ones. This is consistent with the findings of the static validation trial. A surprising result is the slight degradation of the positioning quality obtained with the Ultra Rapid over the Broadcast ephemeris.

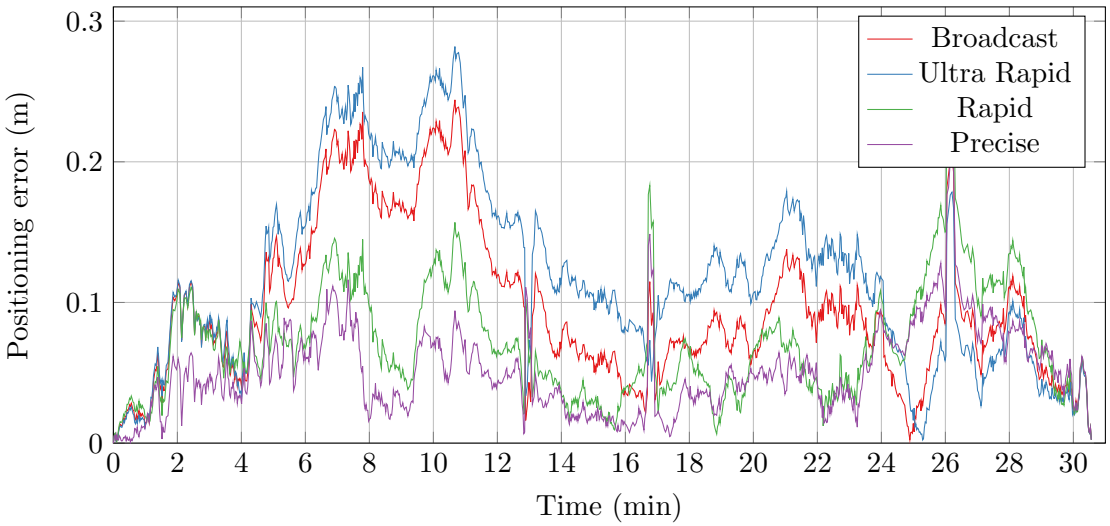


Figure 13.22.: Positioning error in metres obtained with and without the use of precise satellite ephemeris and clocks

Velocity error

The second variable investigated is the quality of the velocity estimation. The velocity is obtained from the integration of GPS and inertial measurements and is expressed in the navigation frame (Section 11.2) for both the developed system and the reference system. Velocity estimates are available at 200 Hz from both systems.

The performance of the velocity estimates has been investigated in 3-dimensions as well as in 2-dimensions (i.e. planar velocity) which is of significant interest for the sprint canoeing and kayaking application.

The 3 dimensional velocity estimate error is presented on Figure 13.23. The distribution reports a mean error of -0.757 mm/s with a standard deviation of 6.777 mm/s.

Similarly the planar velocity estimate error is presented on Figure 13.24. The error distribution reports a mean error of -0.759 mm/s with a standard deviation of 6.785 mm/s. In other words, estimates of the boat velocity are available at

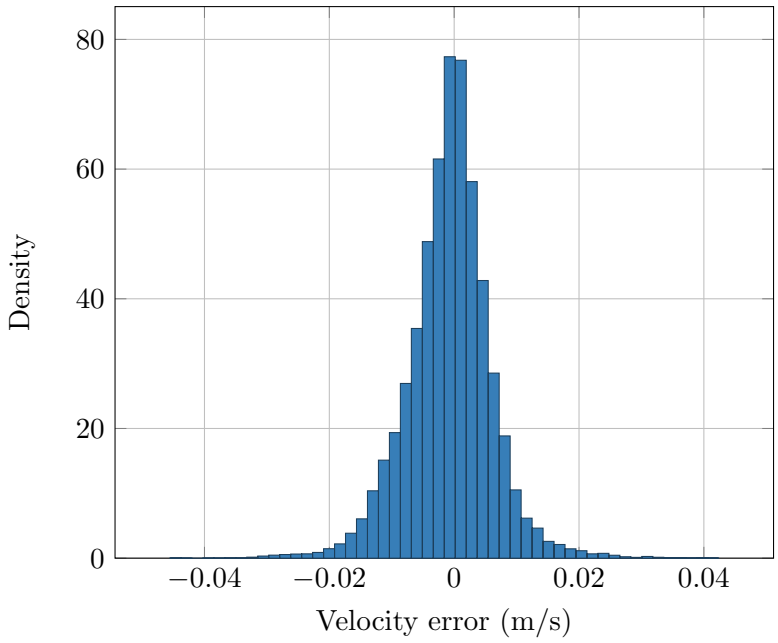


Figure 13.23.: Histogram of the 3 dimensional velocity error over the 30 minutes trial

200Hz with a standard deviation equivalent to 0.02 km/h or 0.006 m/s.

Attitude error

The last output of interest is the attitude estimation. Its quality has been investigated through a direct comparison with the attitude obtained with the reference system. Although the two inertial measurement units were mounted and screwed to the same rigid plate there remains an alignment error between the units. The direct consequence is a small offset in the angles obtained by the two systems. Table 13.9 summarises the mean error on the three angles mainly reflecting the units misalignment.

The mean error has been subtracted as it is largely due to the misalignment between the units. Figure 13.25 presents histograms of the attitude error of the developed system against the reference.

The results show the benefits of an inertial navigation solution on the attitude

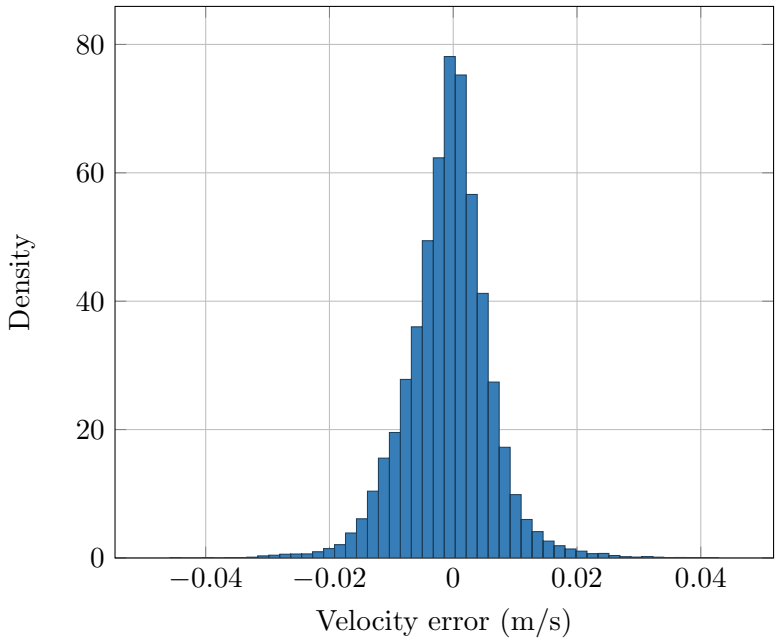


Figure 13.24.: Histogram of the planar velocity error over the 30 minutes trial

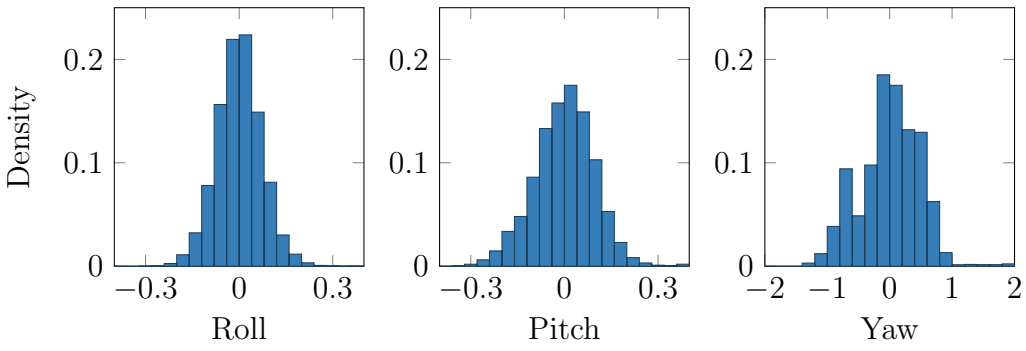


Figure 13.25.: Attitude error in degrees over the 30 minutes trial

Table 13.9.: Mean attitude error largely due to misalignment between the units

	Misalignement
Roll	0.073°
Pitch	0.084°
Yaw	0.425°

accuracy. According to the specifications of the MicroStrain 3DM-GX3-25 [91], the IMU used in the developed solution, an accuracy of ±2.0° is expected for roll, pitch and yaw in a dynamic situation. The standard deviations of the attitude error using the processing developed in this work are summarised in Table 13.10.

Table 13.10.: Standard deviation of the attitude error

	Standard dev.
Roll	0.071°
Pitch	0.096°
Yaw	0.336°

The developed processing largely outperformed the attitude accuracy that is obtained by the on-board attitude estimation. Roll and pitch angles are estimated with a standard deviation of 0.071° and 0.096° respectively. The yaw angle, more challenging to be estimated, is reported with a standard deviation of 0.336°.

14. Sport specific metrics

This Chapter focuses on the measurements and metrics that have been developed to specifically support the sprint canoeing and sprint kayaking application. Different aspects need to be taken into consideration, from how the sprint distance is computed to the detection of sport specific patterns.

14.1. Distance estimation

The distance estimation is critical for the application within sprint canoeing and sprint kayaking as it is the measure determining the end of a sprint. Any error in the distance estimation is directly transformed into a timing error and thus affecting the efficiency of the system. An inaccurate distance estimation would thus critically affect the validity of the system for its targeted application. For these reasons, accurate distance estimation was a key objective from the English Institute of Sport and GB Canoeing.

Apart from during competitions, coaches and athletes do not have a way of precisely measuring sprints distance. Commonly, they rely on the coloured buoys delimiting the rowing lanes or, if present, some markers on the sides of the lake. From there, the coach would start a stop watch which is then stopped when he thinks the target distance has been covered. It is obvious that this process has a limited accuracy.

The system developed in this work produces position estimates which can be used to estimate the distance covered by the athlete and thus marking the end of a sprint. However, the computation of the distance covered has to be done carefully due to the Earth's shape. This section describes the most common ways of computing distance estimates, each of them accounting for the Earth's shape in a different way.

14.1.1. Haversine

A simple way of accounting for the Earth's shape is to use the Haversine formula. This formula gives the great-circle distance between two points on a sphere. When used to compute distance on the Earth, a spherical assumption is made on the Earth's shape. The Haversine function is defined as:

$$\operatorname{hav}(\theta) = \sin^2\left(\frac{\theta}{2}\right) = \frac{1-\cos(\theta)}{2} \tag{14.1}$$

Based on the Haversine function, the great-circle distance, D, between two points can be computed with:

$$\begin{aligned} d &= 2R \arcsin\left(\sqrt{\operatorname{hav}(\varphi_2-\varphi_1) + \cos(\varphi_1)\cos(\varphi_2)\operatorname{hav}(\lambda_2-\lambda_1)}\right) \\ &= 2R \arcsin\left(\sqrt{\sin^2\left(\frac{\varphi_2-\varphi_1}{2}\right) + \cos(\varphi_1)\cos(\varphi_2)\sin^2\left(\frac{\lambda_2-\lambda_1}{2}\right)}\right) \end{aligned} \tag{14.2}$$

where:

d	is the great-circle distance between the two points on a sphere (m)
R	is the radius of the sphere (m)
φ_1, φ_2	are the respective points latitude (rad)
λ_1, λ_2	are the respective points longitude (rad)

Usually, R is defined as the average Earth radius with $R = 6371$ km. However,

the Earth is not a perfect sphere and its radius varies. When cut along a constant latitude plane, the radius varies from 6356 km at the poles to 6378 km at the equator. Similarly, and more importantly, when cut along a constant longitude plane the earth radius varies by more than 1% with 6399 km at the poles compared to 6335 km at the equator. As a consequence, the Haversine formula can not guarantee to provide a distance estimate better than 0.5%.

14.1.2. Vincenty's formulae

The Vincenty's formulae [140, 141] is a reference method used in geodesy to calculate distance between two points on the Earth. More generally, it is an iterative method which computes the distance between two points on the surface of a spheroid. One of the best assumption that can be made on the Earth's shape is an oblate spheroid which explains why this method is widely used in geodesy applications. As the formulae computes a distance between two points on an ellipsoid, it is always associated with an ellipsoid model. When dealing with points on the surface of the Earth, a common choice is to use the WGS-84 model.

On the Earth ellipsoid, distances obtained with Vincenty's formulae are accurate to within 0.5 mm.

One downside of the Vincenty's formulae is its expensive computation budget due to the complex operations involved and its iterative nature.

14.1.3. Local tangent plane

The distance on the local tangent plane is the distance used in this work. The local tangent plane is closely linked to the navigation frame defined in Section 11.2.3. The three axis composing the local tangent plane are defined as the local east e,

the local north n and the local up u. It actually is the navigation frame kept fixed at a certain point. Different choices can be made regarding the origin point also referred to as the projection point. During real-time operations, the fixed point is usually chosen as the first point of the trial. For post-processing purposes, two popular choices are to use the central point of the trial to distribute the induced projection error on the whole session. A second choice can be to use the point with the minimal latitude and minimal longitude of the trial so that any point of the trial have positive local coordinates. The conversion from geographic coordinates to local tangent plane coordinates is defined by:

$$\begin{aligned} e &= (\lambda - \lambda_0) \times R_\lambda \times \cos(\varphi) \\ n &= (\varphi - \varphi_0) \times R_\varphi \\ u &= h - h_0 \end{aligned} \tag{14.3}$$

where:

φ_0 is the latitude of the projection point (rad)

λ_0 is the longitude of the projection point (rad)

h_0 is the altitude of the projection point (rad)

R_φ is the latitudinal Earth radius (m)

R_λ is the longitudinal Earth radius (m)

14.1.4. Comparison

Based on the on-water trials (Section 13.2), a 100, 200 and 500 metres sprints have been investigated and used to present a comparison of the different distance calculation and their performance. As described above, the Vincenty's formulae is the reference and thus is used as a comparison point. The results are presented in Table 14.1.

Table 14.1.: Comparison of the different distance calculations with distance in metres and the error percentage

Vincenty	Haversine		Local Tangent Plane	
	Distance	Error	Distance	Error
100 m	99.810	−0.190%	100.000	±0.000%
200 m	199.618	−0.191%	200.000	±0.000%
500 m	499.033	−0.193%	500.006	+0.001%

By assuming a sprint velocity of 4 m/s, this distance error can then be converted into an induced timing error which is presented in Table 14.2

Table 14.2.: Comparison of the timing error induced by the different distance calculations

Distance	Sprint time	Haversine	Local Tangent Plane
100 m	0'25.000"	+0.047"	−0.000"
200 m	0'50.000"	+0.096"	−0.000"
500 m	2'05.000"	+0.242"	−0.002"

Both Vincenty's formulae and local tangent plane distance calculation are acceptable for sprint canoeing and kayaking timing. The Haversine distance affects the timing accuracy by a magnitude which can not be neglected for this application. It also has a lower accuracy than the positioning accuracy (Section 13.2.6) and the timing accuracy (Section 14.2) obtained with the processing developed and thus would degrade the overall performance of the system.

The results obtained are specific to the location of the trials, in this case Dorney Lake in England. A similar analysis done on another part of the globe would give different results. The discrepancy originates from the Earth's shape and how well the different assumptions are modelling it at a given location.

14.2. Sprint timing

As previously stated, an accurate sprint timing was one the main requirements for this research. Because the processing developed produces timestamped position estimates, a sprint time can be extracted by monitoring the distance covered by the boat. The different ways to compute distance alongside the precision of each method have been detailed in the previous section. It has been shown that the distance on the local tangent plane, the one used in this work, produces distance estimates at Dorney Lake, England with an accuracy greater than 0.001% of the distance. This accuracy is a level of precision better than the positioning quality obtained with the developed solution (Section 13.1 and Section 13.2.6). As a result, the timing accuracy presented in this section, largely reflects the positioning quality of the developed solution.

In order to evaluate the timing quality of the proposed solution, all the sprint efforts of the trials have been investigated. For each sprint, a reference time has been computed using the gold-standard system. Whilst the detection of the end of a sprint is a direct consequence of the distance estimation, there is still a need to define what constitutes the beginning of an effort. The assumption made in this work is to use the beginning of the first stroke to mark the effort's start. A reliable stroke detection algorithm has been developed as part of this work using inertial measurements. The algorithm itself and its performance and reliability are presented in Section 14.4.2.

The reference times have been computed using the positioning solution of the SPAN system, our gold-standard system. Using the inertial measurements of the equipment, strokes were detected (Section 14.4.2) and the beginning of the first stroke used as the sprint start time. The distance covered by the athlete was then computed using the Vincenty's formulae to provide the best distance monitoring possible (Section 14.1.2). Finally, the sprint time was obtained once the target distance had been covered by the athlete.

A similar process has been applied to the developed solution. The beginning of the first stroke was detected based on the inertial measurements and used as the start time of the effort. The distance covered by the athlete was computed using the distance on the local tangent plane due to its accuracy and computation efficiency. Finally, the sprint time was obtained once the target distance was reached by the athlete.

Table 14.3.: Comparison of sprint timing performance using the efforts of the on-water trial session

Sprint	Distance	Reference Timing	Developed Timing	Timing Error
Sprint 1	100 metres	00'29.862"	00'29.825"	−0.037"
Sprint 2	100 metres	00'24.219"	00'24.210"	−0.009"
Sprint 3	100 metres	00'23.107"	00'23.110"	0.003"
Sprint 4	500 metres	02'21.378"	02'21.363"	−0.015"
Sprint 5	100 metres	00'21.981"	00'21.969"	−0.012"
Sprint 6	200 metres	00'49.275"	00'49.249"	−0.026"

Table 14.3 presents the results of the two systems performing totally independently. As a consequence, the reference start time is based on the first stroke detection using the inertial measurements of the reference system whilst the developped solution start time is based on the first stroke detection using the inertial measurements of the proposed solution.

A second analysis was made by using the reference start time for both solutions. This would in effect compensate for any error made in the detection of the first stroke. In practice, as the two solutions use two different inertial measurement units, there is an inevitable misalignment of the units which are also subject to slightly different dynamics although they are attached to a rigid body. Hence, the detected start times have small differences. Table 14.4 presents the difference between the start detection for the two systems alongside a compensated timing for the proposed solution obtained by using the reference system start time as the start of the sprint.

Table 14.4.: Comparison of sprint timing performance using the efforts of the on-water trial session after compensation of the start time difference

Sprint	Distance	Start Difference	Corrected Time	Timing Error
Sprint 1	100 metres	−0.002”	00’29.827”	−0.035”
Sprint 2	100 metres	−0.002”	00’24.212”	−0.007”
Sprint 3	100 metres	+0.001”	00’23.109”	+0.002”
Sprint 4	500 metres	−0.002”	02’21.365”	−0.013”
Sprint 5	100 metres	−0.004”	00’21.973”	−0.008”
Sprint 6	200 metres	+0.000”	00’49.249”	−0.026”

As it can be seen, compensating for the start time difference marginally improves the timing accuracy, however, the impact is still within the general positioning accuracy of the proposed solution and thus does not affect the general suitability of the solution.

Additionally, the start time difference is below ± 0.005 seconds and thus lower than the IMU measurement rate (200 Hz). This result supports the validity and quality of the sprint start time detection based on the beginning of the first stroke.

A global analysis of the six sprints gives an average timing accuracy of -0.016 ± 0.013 seconds which is marginally improved by compensating for the error induced by the start time detection to -0.014 ± 0.012 seconds.

In addition to the principal requirement of providing an accurate sprint timing, having accurate split times is also of great interest from a sport performance analysis perspective. Using the six sprints performed during the on-water trials, split times every 10 metres have been computed constituting a set of 110 splits. For each 10 metre split, a timing error has been obtained by comparing the reference split time to the proposed solution split time. Results are presented as an histogram of the split timing error on Figure 14.1.

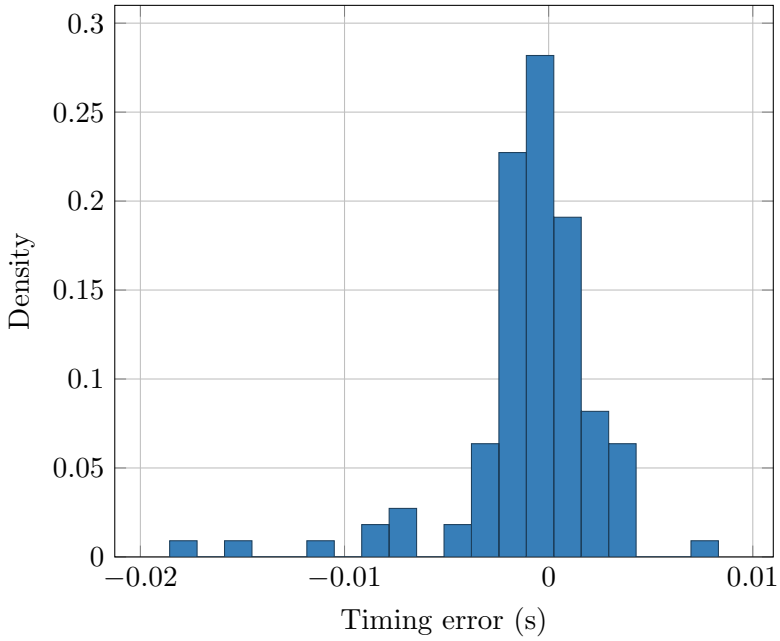


Figure 14.1.: Histogram of the timing error on 10 metres splits using 110 splits

The distribution of the 10 metres split timing error can be described with a mean error of -0.9 milliseconds ($\sigma = 3.5$ ms). A few occurrences of larger timing errors (> 10 ms) are present. After investigation, it has been found that these larger timing errors occurred when the u-blox logging fault (Section 13.2.6) was encountered.

The results indicates that the proposed solution might have a tendency to report marginally fast times. Nevertheless, an accuracy under the 1/100th of a second range is robustly achieved by the proposed solution for 10 metres split times. Sprint timing for distances between 100 and 500 metres is achieved to an accuracy of a few 1/100th of a second, largely meeting the initial expectations and requirements of the sport (Section 7.3).

14.2.1. Potential distance error

With the procedure detailed above, a sprint is considered finished when the distance between the start point of the sprint and the current position of the boat

reaches the targeted distance. This can result in a sprint distance error if the athlete does not race aligned with the sprint lane.

According to the regulations [17], each sprinting lane has a width of 9 metres. In the worst case scenario, the athlete will start from one side of the lane and finish by the other side which means that the sprint will be marked as completed slightly before it actually is. Based on Pythagoras' theorem, the actual sprint distance covered by the athlete is:

$$D = \sqrt{d^2 - w^2} \tag{14.4}$$

where d is the distance measured and w is the cross-track deviation. Table 14.5 presents the distance error in the worst case scenario of the athlete deviating by 9 metres.

Table 14.5.: Distance error induced if the athlete starts from one side of the sprint lane and finishes on the other side

Measured Distance	Real Sprinting Distance
100 m	99.594 m
200 m	199.797 m
500 m	499.919 m
1000 m	999.959 m

The resulting distance error is quite large, especially for shorter distances and need to be kept in mind. However, regulations [17] states that athletes must keep the whole of their boat within the four-metre wide central area of their lane. Based on this 4 metres allowance, the potential distance error is reduced and presented on Table 14.6

Whilst the potentially induced distance error is negligible compared to the developed system's accuracy for distance over 200 metres it can be the bottleneck point of the system for shorter sprints. This highlights the precision obtained with the processing developed and the need to carefully monitor other aspects as they can degrade the overall system's performance.

Table 14.6.: Distance error induced if the athlete deviates by 4 metres between the start and the finish line

Measured Distance	Real Sprinting Distance
100.000 m	99.920 m
200.000 m	199.960 m
500.000 m	499.984 m
1000.000 m	999.992 m

14.3. Aligned boat velocity

The processing technique developed in this work produces an inertial navigation solution including velocity estimates expressed in the navigation frame. Figure 14.2 presents the velocity as obtained from the processing for the 500 metres sprint of the trial.

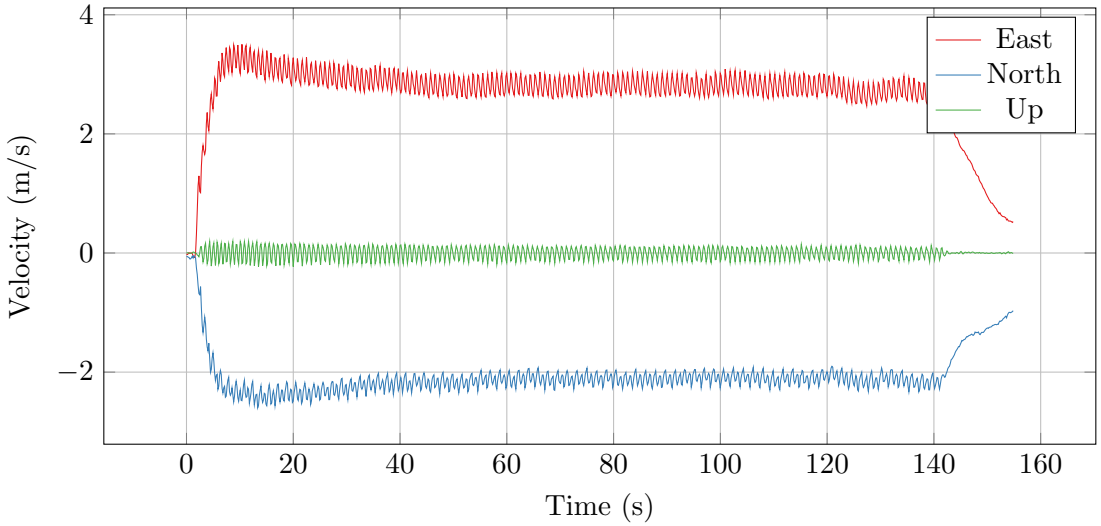


Figure 14.2.: 3D Velocity as obtained from the processing developed with East, North and Up components

Whilst the east, north and up components of the velocity can be of interest for certain applications, they are generally not meaningful for sprint canoeing and kayaking. A first alternative is to generate a global velocity, v_{3D}, representing the

total velocity of the boat and defined by:

$$v_{3D} = \sqrt{v_E^2 + v_N^2 + v_U^2} \tag{14.5}$$

This velocity is similar to what is obtained using GPS doppler measurements [16, 22, 124]. A second option, of bigger interest for the sprint canoeing and kayaking application, is to generate a planar velocity, v_{2D}, representing the horizontal velocity of the boat and thus linked to how fast it progresses on the race course and defined by:

$$v_{2D} = \sqrt{v_E^2 + v_N^2} \tag{14.6}$$

The two options described above can easily be computed and provide interesting information that can be used for performance analysis. However, they do not take fully advantage of the developed solution capabilities. From a sport perspective, an even more interesting information would be to have velocity estimates related to the sprint itself. In other words, having the velocity divided into forward, lateral and vertical components would allow for a deeper level of analysis of the athlete's paddling quality and performance. This can be achieved if the heading of the sprint course is known.

A first approach would be to "hard-code" the heading of the different rowing courses around the world which is not infeasible thanks to the limited set of training and competition locations. However, this is not very practical since preliminary work is required and it would require to have the heading information stored and available on the equipment.

An alternative approach which has been developed as part of this work is to extract a heading estimate based on the position data. Using the position at the beginning of the sprint and the position at the end of it, a heading estimate can be extracted with :

$$\psi_{\text{sprint}} = \operatorname{atan} 2\left(\Delta_E, \Delta_N\right) \tag{14.7}$$

where Δ_E is coordinates difference between the finish position and start position

on the east component and Δ_N is the coordinates difference on the north component. Figure 14.3 presents the aligned boat velocity obtained on the 500 metres sprint. The interest of having an aligned velocity from a sport performance perspective is clear when compared to the standard velocity components as shown on Figure 14.2. The different components now directly relate to the boat movements and ultimately to the athlete's paddling which is of great interest for performance analysis.

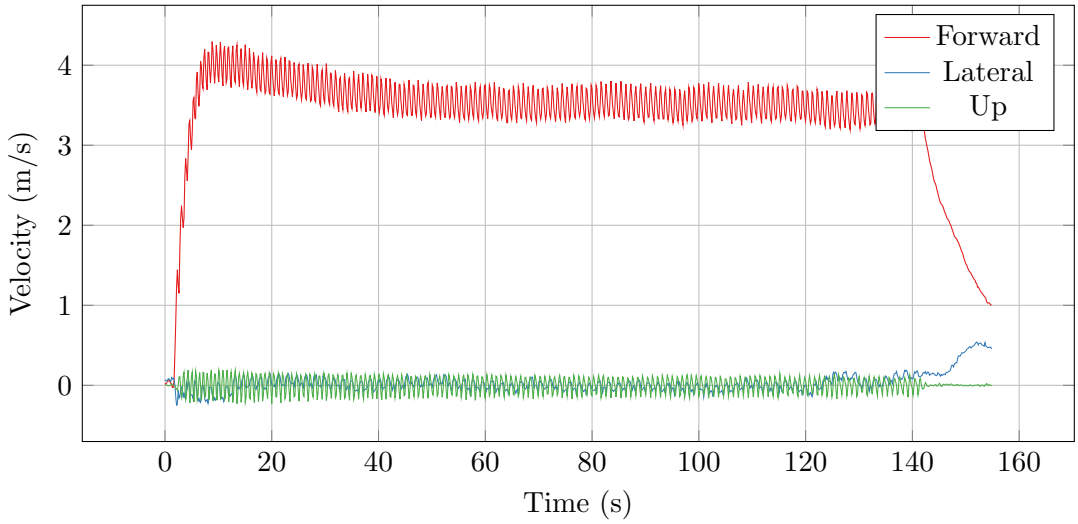


Figure 14.3.: Velocity aligned to the sprint heading with Forward, Lateral and Up components

One limitation of not hard-coding the heading of a course is that aligned boat velocity can not be computed until the sprint is finished so that an estimate of the course heading can be extracted. However, based on discussions with GB Canoeing, there is currently no need for real-time aligned boat velocity feedback as some analysis is needed to make sense of the measurements.

14.4. Stroke analysis

This last section focuses on stroke analysis as it is what drives the boat and eventually lead to a good sprint time. Knowing what happens during a stroke is key and, as such, any stroke information that can be extracted is of significant

value for performance analysis and to understand what makes an athlete fast.

Logically, the first step is the development of an algorithm for accurate stroke detection based on inertial measurements followed by its validation. Once strokes are accurately detected, it opens up a lot of possibilities for performance analysis starting from the traditional stroke duration and the derived stroke rate. However, having an inertial navigation solution can provide a lot more detailed information about each stroke. Stroke metrics such as distance covered, stroke efficiency, resulting speed change or phases durations have been developed as part of this work and are presented below.

14.4.1. Kayaking dynamics

Before detailing the stroke detection algorithm, it is key to understand the dynamics involved in sprint canoeing and kayaking. Figure 14.4 presents a typical 3-axis accelerometers output during 5 strokes.

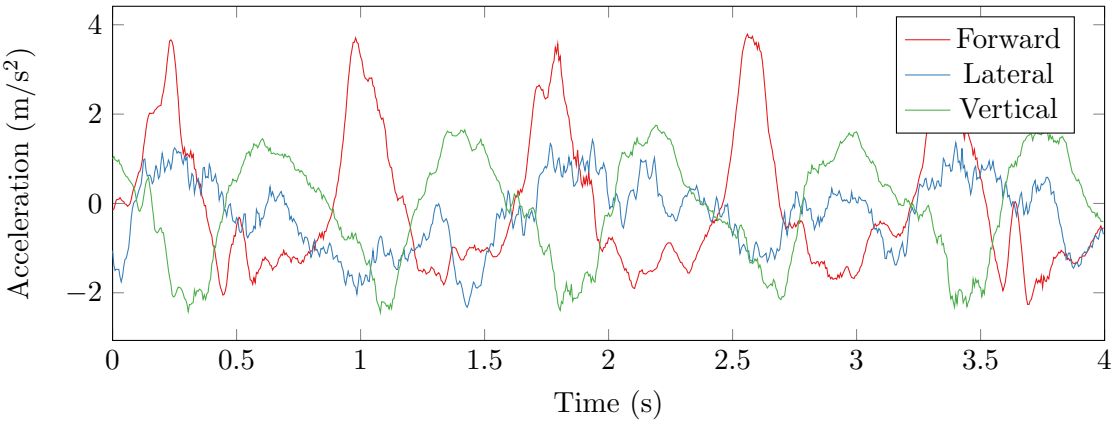


Figure 14.4.: 3-Axis acceleration during 5 strokes

As expected, strokes are clearly visible on the forward axis with a nice acceleration pattern. However, it can be seen that there is also a lot happening on the lateral and vertical axis which is of great interest from a sport performance perspective. Indeed, understanding the origin of lateral and vertical acceleration can result

in performance improvements as these forces are potentially slowing the boat down.

The second output coming from the IMU are the gyroscope measurements which measure how quickly the boat rotates during strokes. Figure 14.5 presents the 3-axis output of the gyroscopes during the same 5 strokes. This time it is clear that pitch and yaw movements are limited compared to the roll axis.

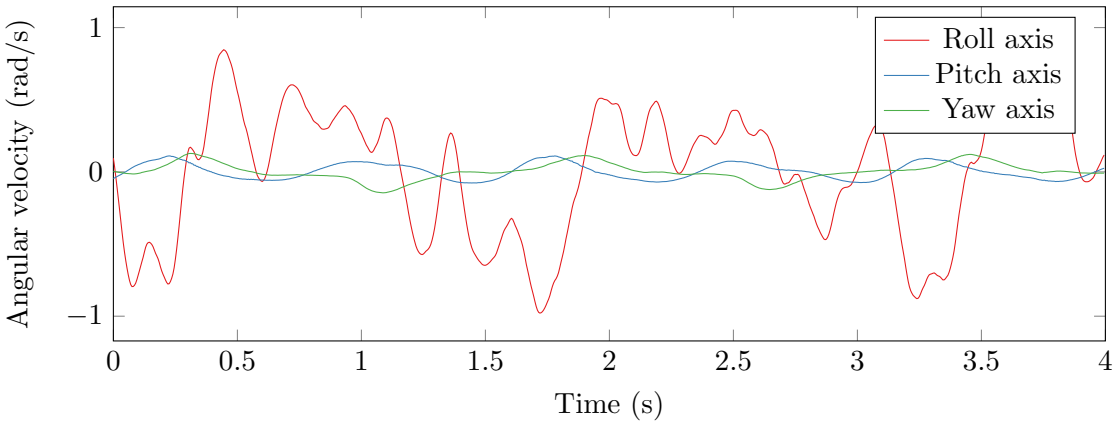


Figure 14.5.: 3-Axis angular velocity during 5 strokes

As quantifying a raw gyroscope output is not easy, Figure 14.6 presents the orientation of the boat resulting from gyroscopes integration. It is once again clear that most of the rotation happens on the roll axis with an amplitude of over 15°. Limited amounts of rotation is also present on the pitch and yaw axis.

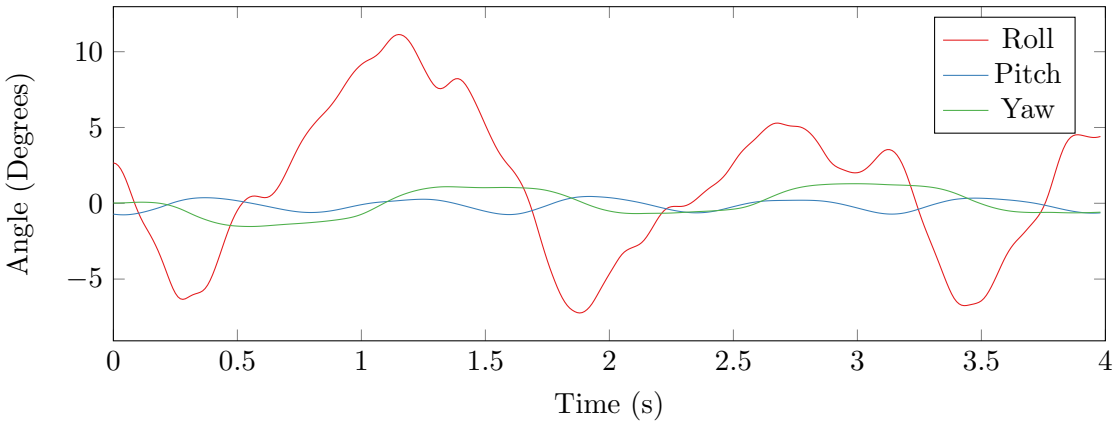


Figure 14.6.: Roll, Pitch and Yaw angles during 5 strokes

It may be noticed that the roll output is not zero-centred. This difference can partially be explained by the misalignment between the IMU and the boat during the installation but mainly by the offset generated once the athlete seats in the boat. Whilst ideally you would like the athlete to maintain the boat to a 0° roll once seated, this is unrealistic and the magnitude of this offset can be another interesting area of investigation for the future.

14.4.2. Stroke detection

With the kayaking dynamics being presented in the previous section, the next focus is on the stroke detection algorithm to identify when a stroke starts and finish. Several definitions of a stroke have been used in the literature [89, 137, 138]. All of them are conceptually correct as they simply consider different part of the stroke as the start and finish marker. After literature comparison and discussions with GB canoeing members, it has been agreed that a stroke starts when the forward acceleration becomes positive and ends with the detection of the next stroke. Hence, one stroke contains both the pull (in-water) and glide (over-water) phases.

In order to detect them, the most singular measurement available is forward acceleration. The technique used is based on rising-edge zero crossings coupled with an acceleration threshold. The threshold represents a minimal forward acceleration to be achieved to consider the stroke valid. This has been done to avoid the detection of small boat movements when almost static or the detection of the adjustment strokes before the start. Based on practical trials, the acceleration threshold has been set to 1.5 m/s^2.

Figure 14.7 shows an example of 5 strokes detection using the developed algorithm.

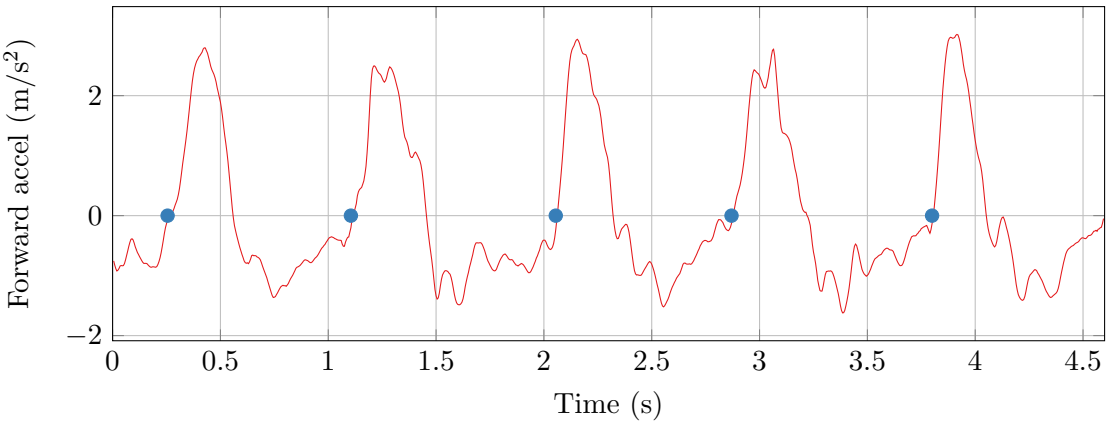


Figure 14.7.: Example of stroke detection using the forward acceleration (5 strokes)

In order to assess the detection rate of the developed algorithm, it has been applied to a paddling dataset recorded over several days by two athletes each of them using a different unit. Detection rates have been obtained by manual verification looking for both missed strokes and wrongly-detected strokes. Results are summarised on Table 14.7.

Table 14.7.: Summary of the stroke detection for 2 athletes using different units over multiple days

	Session	Detected Strokes	Detection Rate
Athlete 1	Day 1	2965	99.7%
	Day 2	3286	99.9%
	Day 3	2552	99.5%
	Day 4	3597	99.6%
Athlete 2	Day 1	2358	99.6%
	Day 2	2330	99.3%
	Day 3	3362	99.4%

The developed algorithm reported an average detection rate of 99.5% over more than 20000 strokes. The algorithm has also been tested using the on-water validation trial data (Section 13.2) using both the MicroStrain IMU and the SPAN IMU data. Equivalent detection rates were obtained and more importantly, the detection timing accuracy has been investigated by comparing the detection time

of each stroke for both systems. The time difference was found to be below 0.005 seconds, the measurement rate of both systems.

14.4.3. Stroke rate

With the strokes being detected, strokes duration and the derived stroke rate can easily be obtained. Figure 14.8 presents strokes duration during the 500 metres sprint and Figure 14.9 the associated stroke rate.

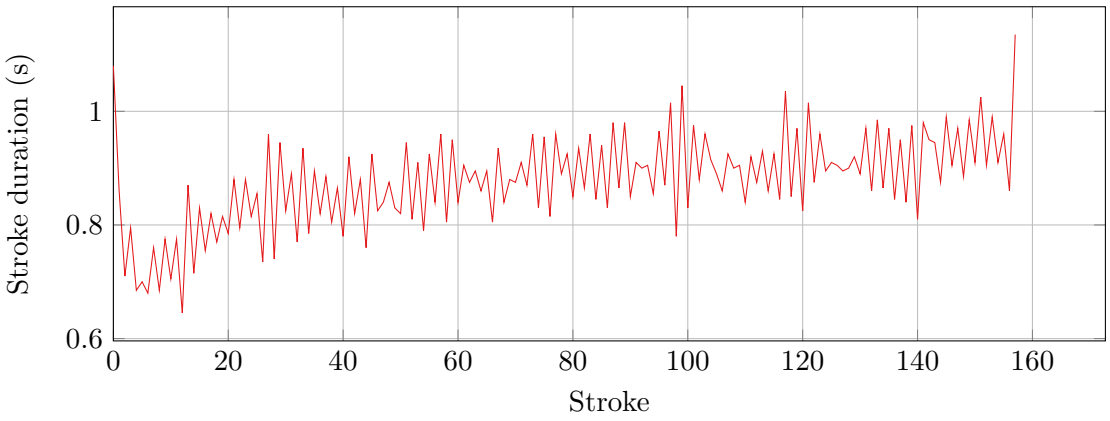


Figure 14.8.: Stroke duration over a 500 m sprint

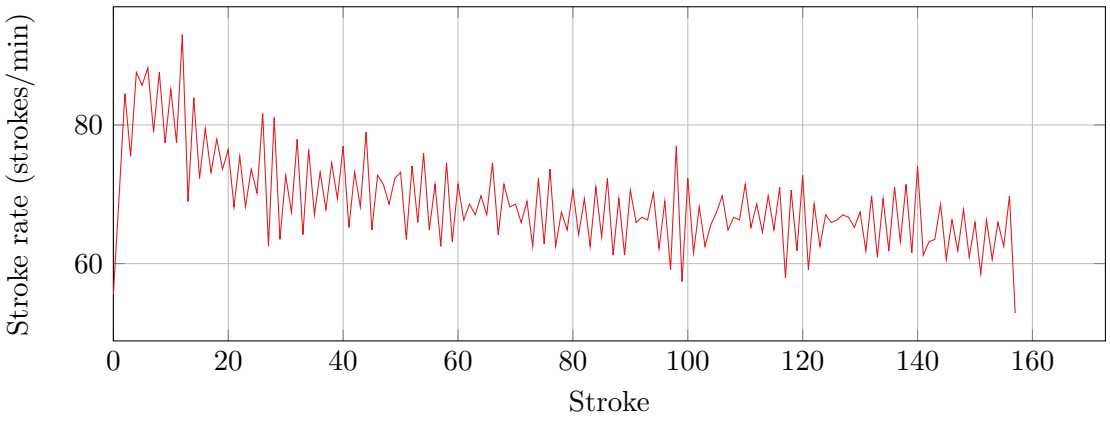


Figure 14.9.: Stroke rate over a 500 m sprint

14.4.4. Stroke distance

Having both an accurate stroke detection and an inertial navigation solution opens up a lot of possibilities to extract performance information for each stroke. A simple example is presented on Figure 14.10 with the distance in metres covered by each stroke.

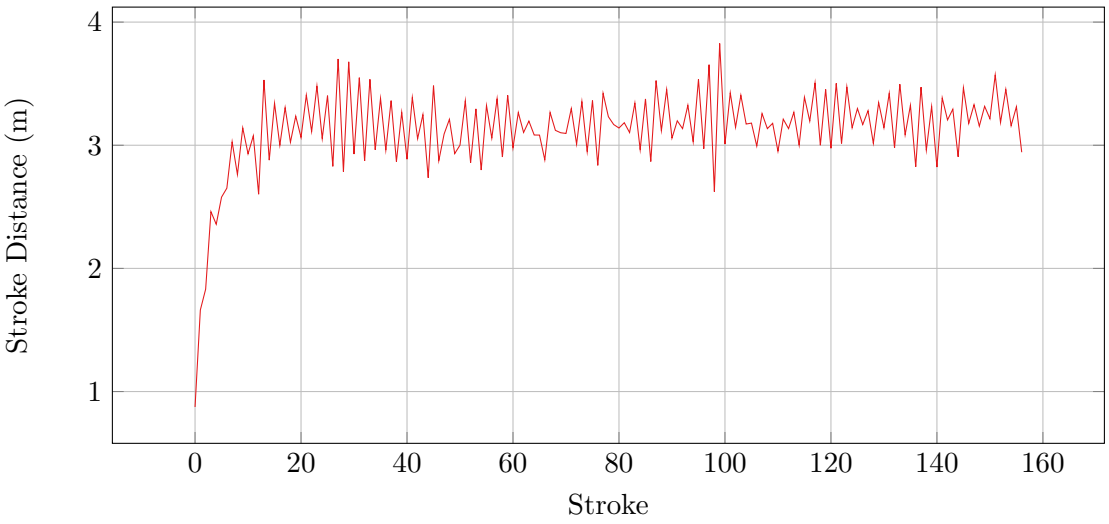


Figure 14.10.: Distance covered by each stroke over the 500 m sprint

14.4.5. Stroke efficiency

In addition to the distance covered by each stroke, the stroke duration is also known. By representing stroke distance against stroke duration, a measure of stroke efficiency is obtained. Figure 14.11 presents the stroke efficiency during the 500 metres sprint excluding the first 40 strokes.

Interestingly, by separating the two paddling sides, clusters can be identified. One side is composed of shorter strokes with less distance covered than the other. Figure 14.12 presents the same measure but including the first 40 strokes this time.

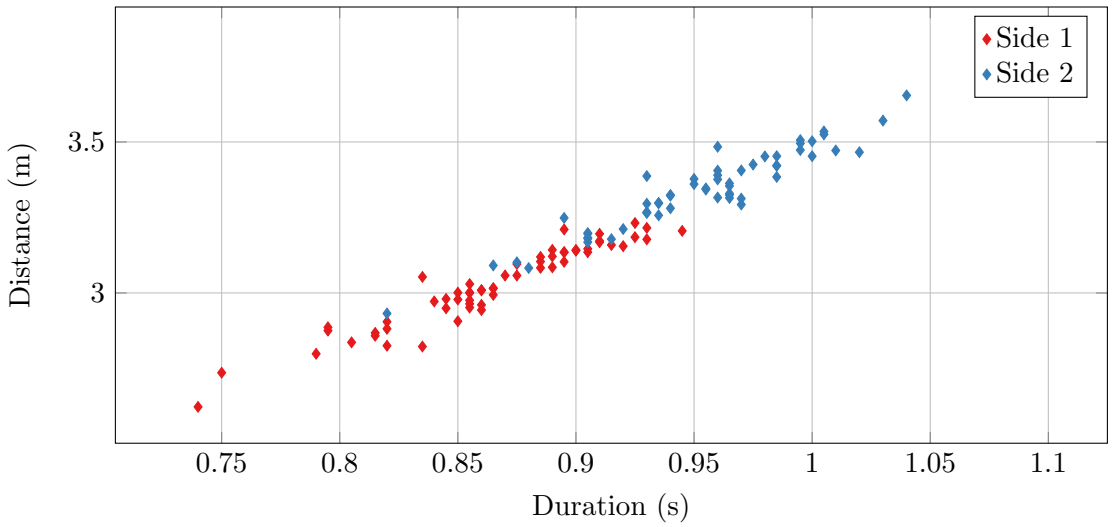


Figure 14.11.: Stroke distance against stroke duration for each side excluding the first 40 strokes

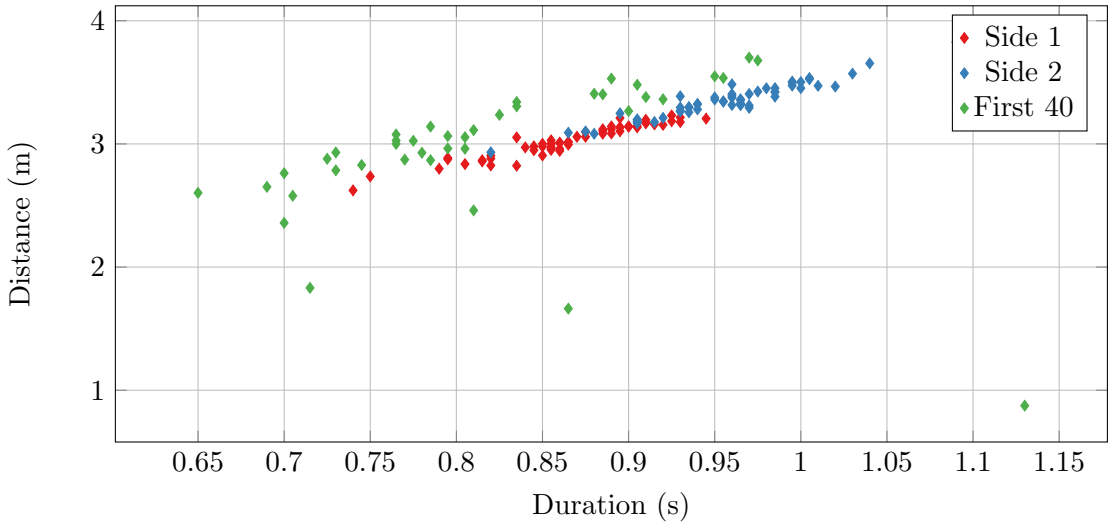


Figure 14.12.: Stroke distance against stroke duration for each side including the first 40 strokes

It can be seen that the first 40 strokes are more irregular. The majority of the first 40 strokes cover more distance in shorter durations than the strokes of the second part of the sprint and thus have a bigger efficiency. This can be interpreted by the athlete making a push for the first 40 strokes before settling down to a more constant pace that can be maintained till the finish line. The four lower efficiency strokes are, not surprisingly, the first 4 strokes of the sprint.

14.4.6. Stroke velocity change

Another measure that can be obtained thanks to the inertial navigation solution and the stroke detection, is the change in forward velocity (Section 14.3) of each stroke. This is obtained by comparing the forward velocity of the boat at the beginning and at the end of each stroke. Figure 14.13 presents the stroke velocity change during the 500 metres sprint.

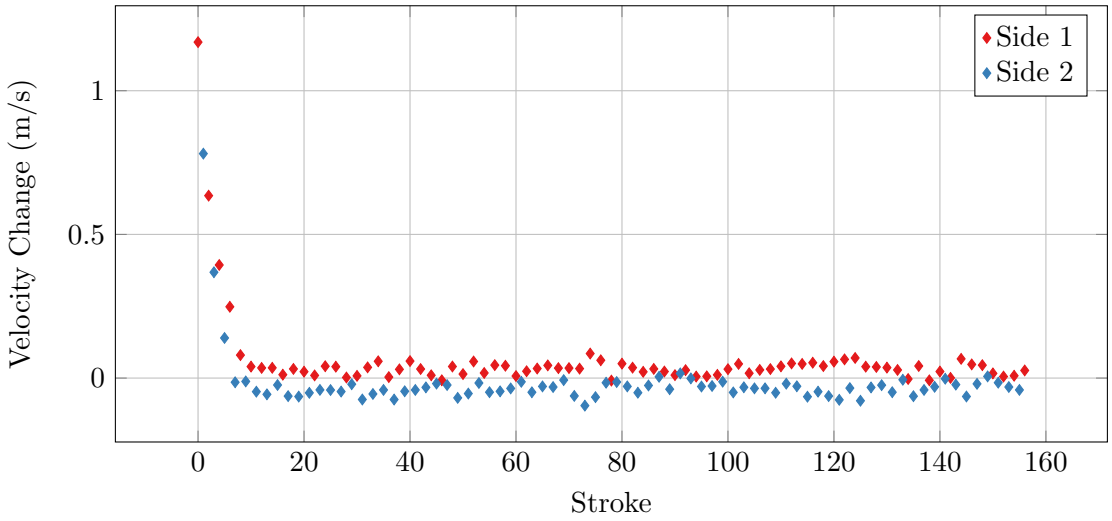


Figure 14.13.: Velocity difference between the start and the end of each stroke

The two paddling sides have been separated using the same labelling as in the previous section. Once again, this leads up to an interesting finding. On one side, the athlete is able to increase its boat velocity whilst on the other side it is loosing speed. More interestingly, the side where the athlete is able to increase its boat velocity is the side of the shorter strokes covering less distance.

14.4.7. Stroke phases

In addition to stroke detection, there is the possibility to detect different phases in each stroke. Based on the work of Wainwright et al. [142] each stroke has been separated into 4 different phases . The first phase starts from the beginning of the stroke until reaching the maximal forward acceleration. The second phase follows until null acceleration is reached. The third phase follows until maximal deceleration is reached which occurs when the paddle is taken out of the water [142]. Finally, the rest of the stroke is the glide phase. Figure 14.14 illustrates the detection of these different phases over 5 strokes.

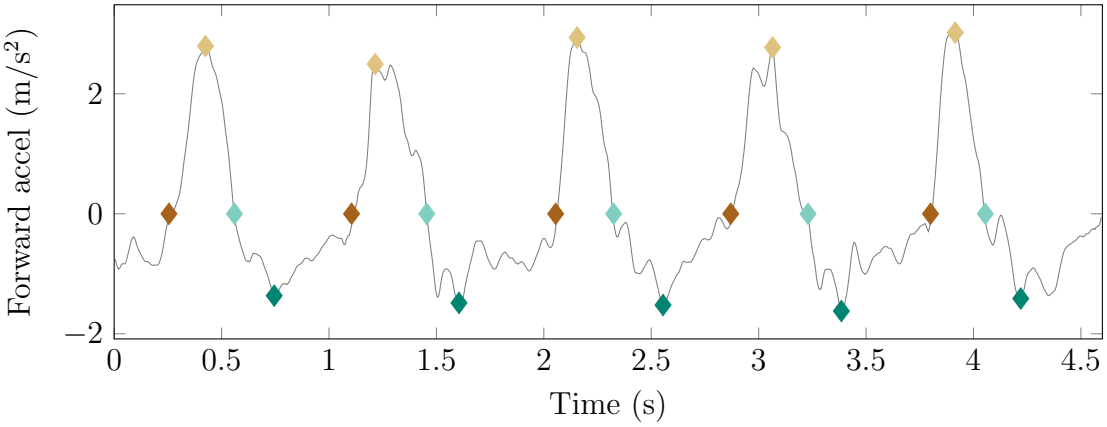


Figure 14.14.: Detection of the different phases of a stroke

Once these different phases are detected, it is possible to plot the duration of each of them for each stroke. Figure 14.15 and 14.16 presents the phases duration for the strokes taken on Side 1 and Side 2 respectively whilst Figure 14.17 concatenates both sides on the same plot.

Different proportions are obtained between both sides. Side 2 visually seems to be more irregular compared to Side 1. It also has a reduced duration of the phase 3 and an increased duration of the glide phase. Whilst these results have been generated to show the possibilities offered by the developed system, a collaborative work with GB Canoeing is needed to interpret them from an athlete and paddling technique perspective.

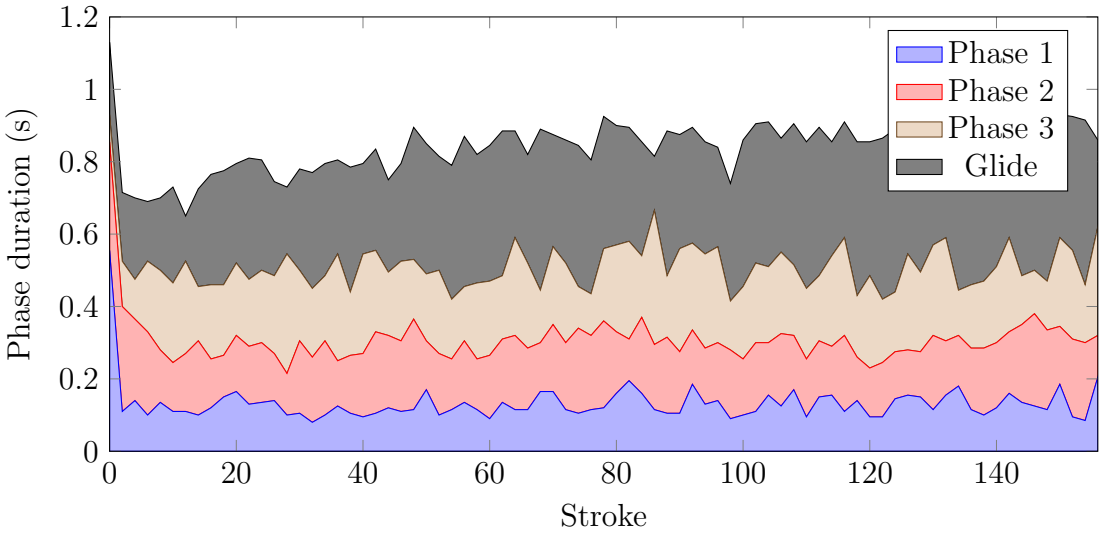


Figure 14.15.: Time duration of the different stroke phases for strokes on side 1

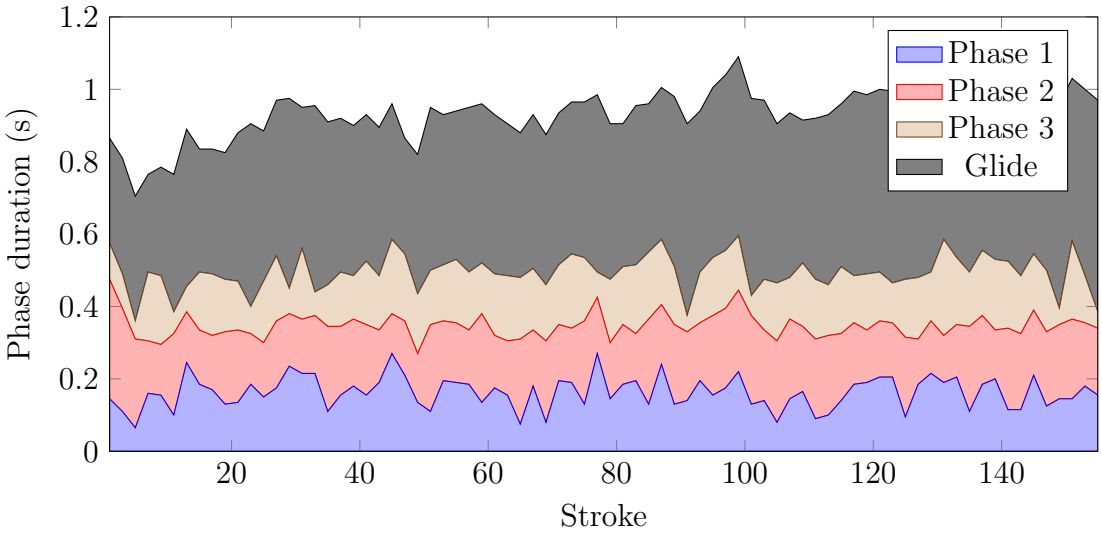


Figure 14.16.: Time duration of the different stroke phases for strokes on side 2

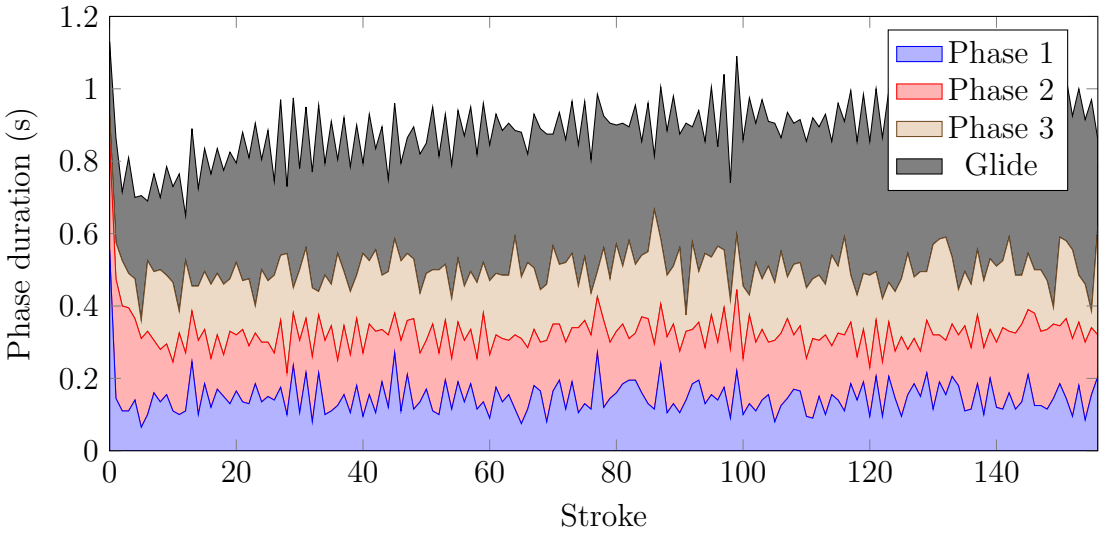


Figure 14.17.: Time duration of the different stroke phases for all the strokes

15. Conclusion

This part of the thesis aimed to develop an outdoor tracking system to be used for sprint canoeing and kayaking disciplines. After summarising the requirements of the sport (Section 7.3), a critical review of the existing commercial systems has been undertaken (Section 8.1). Based on the requirements of the sport and the limitations of the existing solutions, a review of the state-of-the-art positioning technologies and techniques that can be applicable has been completed (Section 8.2). Gathering the information from these three sections, specific positioning technologies and techniques have been identified (Section 8.3).

As the primary sport requirement is position tracking to derive sprint distance and ultimately sprint timing, the GPS technology has been used as the base of this work. A novel GPS processing technique has been developed, allowing relative precise positioning without the need for a base station. Additionally, the technique does not require to solve for integers ambiguity and has no initialisation period which are two significant advantages for this application.

To overcome the limited measurement rate of GPS receivers, the developed solution also uses an IMU to obtain an inertial navigation solution. Not only is the measurement rate increased, the inertial navigation solution also provides velocity and attitude estimates at a very high-rate. This is of significant interest from a sport performance analysis perspective as this technology gives a lot of insights on the dynamics of an athlete during strokes.

In order to combine the IMU and GPS measurements into an integrated inertial navigation solution, a sensor fusion technique is required. Whilst the navigation community traditionnaly relies on Kalman Filtering to achieve the sensor fusion, a novel technique emerging from the robotics community has been used in this work. The technique uses a factor graph to represent the problem and its structure in a convenient way. The factor graph is then optimised via non-linear least squares to obtain an optimal solution. This technique has the advantage of allowing complete re-linearisation as opposed to traditional Kalman filters. It also offers a convenient way do develop new factors which has advantageously been used in this work (Section 12.3).

The solution developed in this work presents novelties and contributions. The first contribution is the formulation of the time-differenced carrier phase positioning technique as a factor graph (Section 12.1). The second contribution is the fusion of the time-differenced carrier phase positioning technique with an IMU with a tightly-coupled approach which is a novelty in the context of factor graph problem optimisation (Section 12.2). In the literature, the integration of GPS measurements with an IMU has always been performed with a loosely-coupled approach in the context of factor graph based problem optimisation. In addition, sport specific factors have been developed and applied (Section 12.3)

An extensive validation of the developed solution has been performed, first in a static controlled environment and then in the targeted application environment with on-water sprint kayak validation. As part of the validation, an analysis of the impact of different parameters on the time-differenced carrier phase positioning techniques has been performed contributing to the knowledge of this technique. The on-water trial highlighted the potential of the developed solution for its sport application. The positioning accuracy reported a mean horizontal error of 4.9 cm with a standard deviation of 2.9 cm and a maximal error of 20.6 cm over the 30 minutes of the trial. In addition, the proposed approach produced velocity estimates with a mean error of -0.759 mm/s and a standard deviation of 6.785

mm/s. In other words, estimates of the boat velocity are available at 200Hz with a standard deviation equivalent to 0.02 km/h or 0.006 m/s. Finally, roll and pitch angles were estimated with a standard deviation of 0.071° and 0.096° respectively. The yaw angle, more challenging to be estimated, was reported with a standard deviation of 0.336°.

In addition, a timing with an accuracy below the $1/100^{th}$ of a second has been robustly achieved for 10-metres splits (Section 14.2). Sprint timing for distances ranging from 100 to 500 metres was obtained with an accuracy of a few $1/100^{th}$ of a second, largely meeting the initial expectations and requirements of the sport (Section 7.3).

An extensive review of the different ways of computing the sprint distance has been presented (Section 14.1) highlighting the importance of choosing the right formulae to obtain accurate sprint timing.

More sport specific considerations have been addressed starting with the potential error induced by the athlete not sprinting parallel to its line (Section 14.2.1). A second focus was on converting the North-East-Down components of the velocity obtained with the developed solution into an aligned boat velocity providing forward, lateral and vertical velocity (Section 14.3). This velocity transformation is of great interest from a sport performance analysis perspective.

In order to fully take advantage of the IMU and the inertial navigation solution, different stroke metrics have been developed (Section 14.4). Knowing the typical dynamics involved in kayaking (Section 14.4.1), a stroke detection algorithm, based on the forward accelerometers output, has been developed reporting an average detection rate of 99.5% over more than 20000 strokes (Section 14.4.2). Directly derived from the stroke detection, strokes duration and stroke rate were presented (Section 14.4.3). Merging the stroke detection with the spatial information provided by the inertial navigation solution, the distance covered by each

stroke has been obtained (Section 14.4.4). Representing the distance covered by each stroke against its duration provided a potential measure of the stroke efficiency (Section 14.4.5). The details of the inertial navigation solution also enables the determination of the velocity change resulting of each stroke (Section 14.4.6). Finally, a stroke itself was divided into 3 phases providing additional insights to be used by performance analysts. Whilst all these sport specific metrics are of great value, a collaborative work with athletes, coaches and performance analysts is necessary to obtain a better understanding and refine them. This process is also likely to result in the development of new metrics.

In terms of the limitations of this work, only relative precise positioning is achieved. The developed solution is not suitable for any application interested in absolute positioning on the globe. From a real-time usage perspective, it is clear that loop closures are unlikely to be applicable. The impact of this limitation has however been investigated and presented in Section 13.1.3 and 13.2.6. The different drift rates computed during the on-water trials suggest that even without loop closure and considering the worst case scenario over the whole session, sprint timing can still be achieved with a few $1/100^{th}$ of a second accuracy. Still on a real-time usage, the developed solution would benefit from specific improvements to improve its computational efficiency. At the moment, the whole problem is solved at each iteration. However, with new measurements coming in, generally only a small subset of the factor graph is modified and need to be optimised. Such a partial update has been developed by the robotics community to solve the SLAM problem online [64]. This approach is also termed as incremental optimisation in the literature and provide an efficient way of performing real-time optimisation for a reduced computational cost.

15.1. Contribution to knowledge

Several contributions to knowledge have been made in this part of the PhD both to the navigation and sport science communities. Below is a list of the key contributions and novelties of this work.

- Factor graph formulation of the time-differenced carrier phase positioning technique (Section 12.1.1)

- Sensor fusion of GPS and IMU measurements with a tightly-coupled approach which has never been done in the context of factor graph based optimisation (Section 12.2)

- Extensive validation of the developed solution giving new insights on the impact of some parameters of the GPS technique such as ephemerides and weighting schemes (Section 13.1 and 13.2.6)

- Use of robust loss functions to reduce the impact of potential outliers on measurements made on low elevation satellites (Section 12.1.2)

- Development of sport specific factors (Section 12.3.1 and 12.3.2)

- In-depth review of sprint distance and timing aspects (Section 14.1, 14.2 and 14.2.1)

- Development of sport specific measurements and metrics:

  - Aligned boat velocity (Section 14.3)

  - Stroke detection (Section 14.4.2)

- Stroke duration and stroke rate (Section 14.4.3)
- Stroke distance (Section 14.4.4)
- Stroke efficiency (Section 14.4.5)
- Stroke velocity change (Section 14.4.6)
- Intra-stroke phases (Section 14.4.7)

The developed solution also has a big potential contribution to the sport science field. It has been identified that existing devices and solutions are a limiting factor for sport scientists to conduct studies on sprint canoeing and kayaking. The inertial navigation solution obtained from the developed solution allows for deeper levels of analysis as well as the accessibility to new metrics and insights on strokes dynamics.

15.2. Impact for elite sports

In addition to the potential contribution to the sport science research community, the work has already delivered great impact for elite sports. The developed stroke detection algorithm is now used by GB Canoeing to analyse their training sessions and competitions. In addition, the knowledge on distance computation has been shared to ensure that appropriate formulas were used within the sport to compute the distance covered by an athlete. Finally, the different metrics developed in this work should incrementally be implemented and introduced in the sport in an iterative process to find which ones are meaningful for athletes, coaches and performance analysts and how they can be refined.

15.3. Discussion : Other Sports Applicability

Whilst the work had a strong emphasis on sprint canoeing and kayaking disciplines, the developed solution can be applied to any outdoor sport offering a fairly good GPS satellite visibility. As a consequence, water sports such as rowing and sailing could directly benefit from the developed solution. In addition, sports such as BMX, horse racing or even automotive racing are fulfilling all the requirements for developed solution to be applicable. The solution also has the potential to be used during skiing events whether they are downhill or cross-country.

Part IV.

Summary and Conclusions

This thesis investigated and developed state-of-the-art positioning technologies and techniques for precise location measurement in high performance sports. Two specific investigations have been conducted, firstly indoors and then outdoors.

Indoors, the use of a commercially available Ultra-Wide Band (UWB) positioning system to track wheelchair court sports players has been investigated. The literature review shown that an increasing number of studies use UWB systems for sport tracking applications. It also highlighted the need for a quality assessment to validate the capabilities of such systems. Existing studies in the literature did not evaluate the performance of UWB systems in a sport context against a gold-standard reference or only considered pre-defined patterns (Objective 1). The possibilities of the development of a novel positioning algorithm were limited as the commercial UWB system did not allow for internal modifications. However, a specific filtering has been proposed to significantly improve the computation of the distance covered by an athlete. (Objective 2). Through the use of state-of-the-art surveying techniques, a quality assessment of the UWB system has been performed in an ecologically valid environment with a particular emphasis on dynamic tracking. The challenges induced by the wheelchair environment have also been considered. It has been demonstrated that with a quick and easy deployment procedure, dynamic tracking can be achieved with a mean horizontal positioning error of 0.37 m ($\sigma = \pm 0.24$ m). Additionally, distance covered can be determined with an error below 0.5% of course length with adequate data processing. Tag update rates did not have a significant impact on the positioning quality; however, higher rates ($\geqslant 8$ Hz) provided a greater number of points to more closely record the high dynamic movements. It was also found that having many wheelchairs on the court did not have a significant effect on the positioning. Finally, several tag mounting places have been tried with the smallest error obtained for the tag worn by the athlete in a GPS vest (Objective 3). The second contribution is the development of sport specific metrics converting the system's output into meaningful metrics for performance analysis (Objective 4). This work has supported

a number of publications in the sport science community. It has also supported the use of the UWB system during international wheelchair rugby competitions as well as at Team GB training camps. Finally, a critical review of the research has been conducted including the applicability of the UWB system to other sports (Objective 5).

Outdoors, sprint canoeing and kayaking was the targeted sport application. The first requirement of the sport was the development of a solution allowing precise sprint timing. The second requirement was for the solution to deliver insights on the sprinting dynamics to be used for performance analysis. Based on these objectives, the development of a solution combining GPS and IMU measurements has been investigated for its potential to fulfil both requirements without the need for any external infrastructure (Objective 1). As a result, a novel approach has been proposed with the use of an innovative, factor graph based, sensor fusion technique to combine IMU and GPS measurements. A key advantage of the proposed approach over traditionally used Kalman Filters is to allow full re-linearisation of the problem resulting in an improved positioning quality. The proposed approach contributes advances in two major ways. The first main contribution is a novel factor graph formulation of an advanced GPS technique identified as particularly suitable with regards to the sport requirements. The second main contribution is the proposal of a novel tightly-coupled approach to fuse IMU and GPS measurements. Such an approach had never been done in the context of factor graph optimisation. In addition, other contributions have been made through the use of loss functions to down weight potential GPS outliers as well as the development of sport specific factors to further enhance the positioning quality of the solution (Objective 2). An extensive validation and quality assessment of the developed solution has been conducted during static and on-water trials (Objective 3). During the 30 minutes of the on-water validation trial, the proposed approach achieved positioning with a mean horizontal error in the centimetre range. Similarly, the boat velocity was estimated at 200 Hz within a few mm/s. Finally, roll and pitch angles were estimated with an accuracy better than a $1/10^{\text{th}}$ of a degree whilst yaw

was reported with a standard deviation of 0.336°. Thanks to the positioning accuracy of the proposed approach, 10-metres split timing is robustly achieved with an accuracy below the $1/100^{th}$ of a second. Similarly, sprint timing for distances ranging from 100 to 500 metres is achieved with an accuracy of a few $1/100^{th}$ of a second, largely meeting the initial expectations and requirements of the sport. Sport specific considerations have also been addressed with the potential error induced by the athlete not sprinting parallel to its line and the development of an aligned boat velocity. The potential of the proposed approach to obtain new sprinting dynamics insights was additionally investigated. A robust stroke detection algorithm was developed followed by the development of a number of stroke metrics providing a whole range of new insights for performance analysis (Objective 4). Finally, a critical review of the research has been conducted including the applicability of the proposed approach to other sports (Objective 5).

Bibliography

[1] Aeroscout, 2015. URL http://aeroscout.com/asset-tracking.

[2] Zebra, 2016. URL https://www.zebra.com/gb/en/solutions/location-solutions.html.

[3] ZXY, 2016. URL http://www.zxy.no.

[4] S. Agarwal, K. Mierle, and Others. Ceres Solver. http://ceres-solver.org.

[5] A. Alarifi, A. Al-salman, M. Alsaleh, A. Alnafessah, S. Alhadhrami, A. Mai, and H. S. Al-khalifa. Ultra Wideband Indoor Positioning Technologies: Analysis , Recent Advances , and Outlook. *Submetted to IEEE Communications Survey and Tutorials Journal, Second quater 2015*, 2015.

[6] C.-o. Andrei, M. Vermeer, H. Kuusniemi, and H. Koivula. Evaluation of Absolute and Relative Carrier Phase Positioning Using Observations From Navigation-Grade U-Blox 6T Receiver. 2011.

[7] M. Aquino, T. Moore, and A. Dodson. Implications of ionospheric scintillation for GNSS users in Northern Europe. *Journal of Navigation*, 2005.

[8] A. E. Aubert, B. Seps, and F. Beckers. Heart rate variability in athletes. *Sports medicine*, 2003.

[9] R. J. Aughey. Applications of GPS technologies to field sports. *International Journal of Sports Physiology and Performance*, 2011.

[10] S. Banerjee, W. Suski, and A. Hoover. Sensor set switching noise in UWB indoor position tracking. *2012 IEEE International Conference on Ultra-Wideband*, 2012.

[11] J. C. Barbero-Alvarez, A. Coutts, J. Granda, V. Barbero-Alvarez, and C. Castagna. The validity and reliability of a global positioning satellite system device to assess speed and repeated sprint ability (RSA) in athletes. *Journal of science and medicine in sport / Sports Medicine Australia*, 2010.

[12] K. N. K. I. B.E Thompson, W. A Friess. Aerodynamics of speed skiers. *Sports Engineering*, 2001.

[13] R. Bini, P. Hume, and A. Cerviri. A comparison of cycling SRM crank and strain gauge instrumented pedal measures of peak torque, crank angle at peak torque and power output. *Procedia Engineering*, 2011.

[14] H. Black and A. Eisner. Correcting satellite Doppler data for tropospheric effects. *Journal of Geophysical Research: Atmospheres*, 1984.

[15] S. Blake. *Heave compensation using time-differenced carrier observations from low cost GPS receivers*. PhD thesis, 2007. URL `http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.538084`.

[16] D. Borio, N. Sokolova, and G. Lachapelle. Doppler measurements and velocity estimation: A theoretical framework with software receiver implementation. *Proc., ION/GNSS*, 2009.

[17] British Canoe Union. Lane Regulations — Canoe Sprint Racing UK, 2016. URL `http://canoeracing.org.uk/sprint/using-the-lane/`.

[18] K. Britting. *Inertial navigation systems analysis*. 1971.

[19] A. M. Bruton. *Improving the Accuracy and Resolution of SINS / DGPS Airborne Gravimetry*. PhD thesis, 2000.

[20] M. E. Cannon, G. Lachapelle, C. M. Szarmes, J. M. Hebert, J. Keith, and S. Jokerst. DGPS Kinematic Carrier Phase Signal Simulation Analysis for

Precise Velocity and Position Determination. *Journal of The Institute of Navigation*, 1997.

[21] L. Carlone, Z. Kira, C. Beall, V. Indelman, and F. Dellaert. Eliminating conditionally independent sets in factor graphs: A unifying perspective based on smart factors. *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[22] D. Chalko. High accuracy speed measurement using GPS (Global Positioning System). *Scientific Engineering Research P/L*, 2007.

[23] W. Chen. Integration of GPS and INS for precise surveying applications. 1992.

[24] Y. Chen, T. a. Davis, W. W. Hager, and S. Rajamanickam. Algorithm 887 : CHOLMOD , supernodal sparse Cholesky factorization and update / downdate. *Science*, 2006.

[25] Codamotion. Codamotion, 2016. URL `http://www.codamotion.com/`.

[26] J. Collins and R. Langley. Possible weighting schemes for GPS carrier phase observations in the presence of multipath. *Final contract report for the US Army*, 1999.

[27] R. Conker, M. ElArini, C. Hegarty, and T. Hsiao. Modeling the effects of ionospheric scintillation on GPS/SatelliteBased Augmentation System availability. *Radio Science*, 2003.

[28] H. Croft and D. C. Ribeiro. Developing and Applying a Tri-axial Accelerometer Sensor for Measuring Real Time Kayak Cadence. *Procedia Engineering*, 2013.

[29] D. A. Cucci. *A General Sensor-fusion and Parameters Self-calibration Framework with Applications in Mobile Robotics*. PhD thesis, 2014.

[30] P. Dabnichki and E. Avital. Influence of the postion of crew members on aerodynamics performance of two-man bobsleigh. *Journal of Biomechanics*, 2006.

[31] J. Diebel. Representing attitude: Euler angles, unit quaternions, and rotation vectors. *Matrix*, 2006.

[32] W. Ding and J. Wang. Precise Velocity Estimation with a Stand-Alone GPS Receiver. *Journal of Navigation*, 2011.

[33] A. Eryildirim and M. B. Guldogan. A Bernoulli Filter for Extended Target Tracking Using Random Matrices in a UWB Sensor Network. *IEEE Sensors Journal*, 2016.

[34] J. Farrell. *Aided Navigation: GPS with High Rate Sensors*. 2008.

[35] J. A. Farrell and M. Barth. *The global positioning system and inertial navigation*. 1999.

[36] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. IMU Preintegration on Manifold for Efficient Visual-Inertial Maximum-a-Posteriori Estimation. In *Robotics: Science and Systems XI, Rome, Italy*, 2015.

[37] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. On-Manifold Preintegration Theory for Fast and Accurate Visual-Inertial Navigation. 2015.

[38] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. Supplementary Material to : IMU Preintegration on Manifold for E cient Visual-Inertial Maximum-a-Posteriori Estimation. 2015.

[39] R. P. Foundation. Raspberry Pi 1 Model B, 2016. URL `https://www.raspberrypi.org/products/model-b/`.

[40] P. Freda, A. Angrisano, S. Gaglione, and S. Troisi. Time-differenced carrier phases technique for precise GNSS velocity estimation. *GPS Solutions*, 2014.

[41] H. P. Gavin. The Levenberg-Marquardt method for nonlinear least squares curve-fitting problems. *Department of Civil and Environmental Engineering, Duke University*, 2013.

[42] S. Goodwill and S. Haake. Why were spaghetti string rackets banned in the game of tennis? *The 4th international conference on the engineering of sport*, 2002.

[43] S. Goodwill and S. Haake. Ball spin generation for oblique impacts with a tennis racket. *Experimental mechanics*, 2004.

[44] M. Grewal, L. Weill, and A. Andrews. *Global positioning systems, inertial navigation, and integration*. 2007.

[45] P. Groves. *Principles of GNSS, inertial, and multisensor integrated navigation systems*. 2008.

[46] W. Gurtner. RINEX: The Receiver Independent Exchange Format Version 2.11, 2012. URL `https://igscb.jpl.nasa.gov/igscb/data/format/rinex211.txt`.

[47] H. Hartinger and F. K. Brunner. Variances of GPS Phase Observations: The SIGMA- Model. *GPS Solutions*, 1999.

[48] U. Hartmann and A. Mader. Peak force, velocity, and power during five and ten maximal rowing ergometer strokes by world class female and male rowers. *International journal of sports medicine*, 1993.

[49] R. Hedelin. Heart rate variability in athletes: relationship with central and peripheral performance. *Medicine and science in sports and exercise*, 2001.

[50] M. Hedley and J. Zhang. Accurate Wireless Localization in Sports. 2012.

[51] C. Hide. Integration of GPS and low cost INS measurements. 2003.

[52] B. Hoffman-Wellenhof, H. Liehtenegger, and J. Collins. *GPS - Theory and Practise*. Springer - Wien, New York, 2001.

[53] T. Humphreys. Modeling the effects of ionospheric scintillation on GPS carrier phase tracking. *IEEE Transactions on Aerospace and Electronic Systems*, 2010.

[54] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research*, 2015.

[55] V. Indelman and S. Williams. Factor graph based incremental smoothing in inertial navigation systems. In *Information Fusion (FUSION), 2012 15th International Conference on*, 2012.

[56] V. Indelman, A. Melim, and F. Dellaert. Incremental Light Bundle Adjustment for Robotics Navigation. *cc.gatech.edu*, 2013.

[57] V. Indelman, S. Williams, M. Kaess, and F. Dellaert. Information fusion in navigation systems via factor graph based incremental smoothing. *Robotics and Autonomous Systems*, 2013.

[58] I. Janssen and A. Sachlikidis. Validity and reliability of intra-stroke kayak velocity and acceleration using GPS-based accelerometer. *Sports Biomechanics*, 2010.

[59] I. Janssen and A. Sachlikidis. Validity and reliability of intra-stroke kayak velocity and acceleration using a GPS-based accelerometer. *Sports Biomechanics*, 2010.

[60] N. E. Jarboe and P. M. Quesada. The effects of cycling shoe stiffness on forefoot pressure. *Foot & ankle international / American Orthopaedic Foot and Ankle Society [and] Swiss Foot and Ankle Society*, 2003.

[61] C. Jekeli. *Inertial navigation systems with geodetic applications*. 2001.

[62] X. Jin. *Theory of Carrier Adjusted DGPS Positioning Approach and Some Experimental Results*. PhD thesis, 1996.

[63] C. Jones and C. Wilson. Defining advantage and athletic performance: The case of Oscar Pistorius. *European Journal of Sport Science*, 2009.

[64] M. Kaess, H. Johannsson, and R. Roberts. iSAM2: Incremental smoothing and mapping using the Bayes tree. *The International Journal of Robotics Research2*, 2012.

[65] E. Kaplan and C. Hegarty. *Understanding GPS: principles and applications*. 2005.

[66] J. Karvonen and T. Vuorimaa. Heart rate and exercise intensity during sports activities. *Sports Medicine*, 1988.

[67] J. Klobuchar. Ionospheric time-delay algorithm for single-frequency GPS users. *IEEE Transactions on aerospace and electronic systems2*, 1987.

[68] F. Kschischang. Factor graphs and the sum-product algorithm. *IEEE Transactions on information theory*, 2001.

[69] R. Kummerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. G2o: A general framework for graph optimization. *2011 IEEE International Conference on Robotics and Automation*, 2011.

[70] R. B. Langley. GPS Receiver System Noise. *GPS World*, 1997.

[71] Leica. Leica TS30 - Technical data, 2009. URL `http://w3.leica-geosystems.com/downloads123/zz/tps/TS30/brochures/TS30{_}Brochure{_}us.pdf`.

[72] R. Leser, A. Schleindlhuber, K. Lyons, and A. Baca. Accuracy of an UWB-based position tracking system used for time-motion analyses in game sports. *European journal of sport science*, 2014.

[73] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale. Keyframe-based visual-inertial odometry using nonlinear optimization. *The International Journal of Robotics Research*, 2014.

[74] H. Liu and H. Darabi. Survey of wireless indoor positioning techniques and systems. *IEEE Transactions on Systems, Man and Cybernetics*, 2007.

[75] H.-a. Loeliger. Least squares and Kalman filtering on Forney graphs. *Codes, Graphs, and Systems*, 2002.

[76] H.-a. Loeliger. An introduction to factor graphs. *Signal Processing Magazine, IEEE*, 2004.

[77] M. I. a. Lourakis. A Brief Description of the Levenberg-Marquardt Algorithm Implemened by levmar. *Matrix*, 2005.

[78] R. A. Lukes, S. B. Chin, and S. J. Haake. The understanding and development of cycling aerodynamics. *Sports Engineering*, 2005.

[79] X. Luo. *GPS Stochastic Modelling*. Springer Theses. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[80] T. Lupton. *Inertial SLAM with Delayed Initialization*. PhD thesis, 2010.

[81] T. Lupton and S. Sukkarieh. Efficient integration of inertial observations into visual SLAM without initialization. *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2009*, 2009.

[82] T. Lupton and S. Sukkarieh. Visual-inertial-aided navigation for high-dynamic motion in built environments without initial conditions. *IEEE Transactions on Robotics*, 2012.

[83] D. J. Macfarlane, I. M. Edmond, and A. Walmsley. Instrumentation of an ergometer to monitor the reliability of rowing performance. *Journal of sports sciences*, 1997.

[84] K. Madsen, H. Bruun, and O. Tingleff. Methods for non-linear least squares problems. 2004.

[85] R. Mandeljc, J. Per, M. Kristan, and S. Kova. An Alternative Way to Calibrate Ubisense Real-Time Location System via Multi-Camera Calibration Methods. In *19th International Electrotechnical and Computer Science Conference*, 2010.

[86] B. Mason, J. Lenton, J. Rhodes, and R. Cooper. Comparing the activity profiles of wheelchair rugby using a miniaturised data logger and radio-frequency tracking system. *BioMed research*, 2014.

[87] R. Mautz. Indoor Positioning Technologies. Technical Report February 2012, ETH Zurich, 2012. URL `http://e-collection.library.ethz.ch/eserv/eth:5659/eth-5659-01.pdf`.

[88] L. K. Mcdonnell, P. A. Hume, and V. Nolte. Place time consistency and stroke rates required for success in K1 200-m sprint kayaking elite competition. *International Journal of Performance Analysis in Sport*, 2013.

[89] L. McDonnell. The effect of stroke rate on performance in flat-water sprint kayaking. 2013.

[90] V. Mendes and R. Langley. Tropospheric zenith delay prediction accuracy for airborne GPS high-precision positioning. In *Proceedings of The Institute of Navigation 54th Annual Meeting*, 1998.

[91] MicroStrain. MicroStrain - 3DM-GX3 -25 - Product Datasheet. Technical report, 2014. URL `http://files.microstrain.com/3DM-GX3-25-Attitude-Heading-Reference-System-Data-Sheet.pdf`.

[92] M. Miller and J. Herniman. The effects of a 6-week plyometric training program on agility. *Journal of Sports Science and Medicine*, 2006.

[93] P. Misra and P. Enge. *Global Positioning System: Signals, Measurements and Performance Second Edition*. 2006.

[94] K. Muthukrishnan and M. Hazas. Position estimation from UWB pseudorange and angle-of-arrival: A comparison of non-linear regression and Kalman filtering. *Location and Context Awareness*, 2009.

[95] A. Niell. Global mapping functions for the atmosphere delay at radio wavelengths. *Journal of Geophysical Research: Solid Earth*, 1996.

[96] NovAtel. NovAtel - SPAN-SE. 1993.

[97] NovAtel. NovAtel - UIMU-LCI. Technical report, 1993. URL `http://www.novatel.com/assets/Documents/Papers/IMU-LCI.pdf`.

[98] NovAtel. NovAtel - Waypoint - Inertial Explorer. Technical report, 2015.

[99] NovAtel. NovAtel - Antennas. Technical report, 2015.

[100] B. Parkinson. *Progress in Astronautics and Aeronautics: Global Positioning System: Theory and Applications*. 1996.

[101] T. Paulson, B. Mason, and J. Rhodes. Individualized internal and external training load relationships in elite wheelchair rugby players. *Frontiers in*, 2015.

[102] B. Perrat, M. J. Smith, B. S. Mason, J. M. Rhodes, and V. L. Goosey-Tolfrey. Quality assessment of an UWB positioning system for indoor wheelchair court sports. *Proceedings of the Institution of Mechanical Engineers*, 2015.

[103] F. Ramirez-Mireles. On the performance of ultra-wide-band signals in Gaussian noise and dense multipath. *IEEE Transactions on vehicular technology*, 2001.

[104] T. Rappaport, J. Reed, and B. Woerner. Position location using wireless communications on highways of the future. *IEEE Communications Magazine*, 1996.

[105] J. Rhodes, B. Mason, B. Perrat, M. Smith, and V. Goosey-Tolfrey. The validity and reliability of a novel indoor player tracking system for use within wheelchair court sports. *Journal of sports sciences*, 2014.

[106] J. M. Rhodes, B. S. Mason, B. Perrat, M. J. Smith, L. A. Malone, and V. L. Goosey-Tolfrey. Activity Profiles of Elite Wheelchair Rugby Players During Competition. *International journal of sports physiology and performance*, 2014.

[107] J. M. Rhodes, B. S. Mason, L. A. Malone, and V. L. Goosey-Tolfrey. Effect of team rank and player classification on activity profiles of elite wheelchair rugby players. *Journal of sports sciences*, 2015.

[108] M. Robinson, L. Holt, T. Pelham, and K. Furneaux. Accelerometry measurements of sprint kayaks: The coaches' new tool. *International Journal of Coaching Science*, 2011.

[109] J. Saastamoinen. Contributions to the theory of atmospheric refraction. *Bulletin Géodésique (1946-1975)*, 1973.

[110] Sage. Journals Permissions — SAGE Publications Inc, 2016. URL `https://us.sagepub.com/en-us/nam/journals-permissions`.

[111] T. Sathyan and D. Humphrey. WASP: A system and algorithms for accurate radio localization using low-cost hardware. *IEEE Transactions on Systems, Man and Cybernetics*, 2011.

[112] T. Sathyan, R. Shuttleworth, M. Hedley, and K. Davids. Validity and reliability of a radio positioning system for tracking athletes in indoor and outdoor team sports. *Behavior research methods*, 2012.

[113] C. Satirapod and J. Wang. Comparing the quality indicators of GPS carrier phase observations. *Geomatics Research Australasia*, 2000.

[114] L. Serrano, D. Kim, and R. Langley. A single GPS receiver as a real-time, accurate velocity and acceleration sensor. *Proceedings of the 17th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*, 2004.

[115] E. Shin. *Accuracy improvement of low cost INS/GPS for land applications*. 2001.

[116] E. Shin. Estimation techniques for low-cost inertial navigation. *UCGE report*, 2005.

[117] K. Siwiak. Ultra-wide band radio: introducing a new technology. *Vehicular Technology Conference*, 2001.

[118] J. Skaloud. Strapdown INS orientation accuracy with GPS aiding. 1995.

[119] J. Sola. Course on SLAM. 2015.

[120] Stats. Sportvu, 2016. URL http://www.stats.com/sportvu/sportvu.asp.

[121] D. Stefanyshyn and C. Fusco. Increased Shoe Bending Stiffness Increases Sprint Performance. *Sports Biomechanics*, 2004.

[122] A. Stelzer. Concept and application of LPM-a novel 3-D local position measurement system. *IEEE Transactions on Microwave Theory and Techniques*, 2004.

[123] M. Susi and M. Aquino. Design of a robust receiver architecture for scintillation monitoring. In *IEEE/ION Position, Location and Navigation Symposium-PLANS 2014*, 2014.

[124] M. Szarmes. DGPS high accuracy aircraft velocity determination using Doppler measurements. *Proceedings of the international symposium on kinematic systems (KIS97)*, 1997.

[125] T. Takasu. RTKLIB: An Open Source Program Package for GNSS Positioning. URL http://www.rtklib.com/.

[126] T. Takasu. RTKLIB ver. 2.4.2 Manual. Technical Report C, 2013.

[127] D. H. Titterton and J. L. Weston. *Strapdown Inertial Navigation Technology - 2nd Edition*. 2004.

[128] J. Traugott, D. Odijk, O. Montenbruck, and G. Sachs. Making a Difference with GPSTime Differences for Kinematic Positioning with Low-Cost Receivers. *GPS World*, 2008.

[129] J. Traugott. *Precise Flight Trajectory Reconstruction Based on Time-Differential GNSS Carrier Phase Processing*. PhD thesis, 2011. URL http://d-nb.info/1013435400/.

[130] B. Triggs and P. McLauchlan. Bundle Adjustment - A Modern Synthesis. In *International workshop on vision algorithms*, volume 34099, pages 298–372, 1999.

[131] Trimble. Trimble - CenterPoint RTX, 2016. URL http://www.trimble.com/positioning-services/centerpoint-rtx.aspx.

[132] Trimble. Trimble - VRS Now, 2016. URL http://www.trimble.com/positioning-services/vrs-now.aspx.

[133] U-blox. u-blox 6 - Receiver Description, 2016. URL https://www.u-blox.com.

[134] Ubisense. How to set up a Ubisense system, 2014. URL http://eval.ubisense.net/howto/SystemSetup1{_}article/SystemSetup1.html.

[135] UNAVCO. TEQC — Software — UNAVCO. URL https://www.unavco.org/software/data-processing/teqc/teqc.html.

[136] J. Uren and W. Price. Control Networks. In *Surveying for engineers*, chapter 6.8, pages 241–245. Palgrave Macmillan, 4th edition, 2006.

[137] G. Vadai, G. Makan, and Z. Gingl. On-water measurement and analysis system for estimating kayak paddlers' performance. In *Information & Communication Technology Electronics & Microelectronics (MIPRO), 2013 36th International Convention on*, pages 131–136, 2013.

[138] G. Vadai and Z. Gingl. Performance estimation of kayak paddlers based on fluctuation analysis of movement signals. In *Noise and Fluctuations (ICNF), 2013 22nd International Conference on*, pages 1–4. Ieee, 2013.

[139] F. van Graas and A. Soloviev. Precise velocity estimation using a stand-alone GPS receiver. *Navigation*, 2004.

[140] T. Vincenty. Geodetic inverse solution between antipodal points. *Richard Rapp Geodetic Science Ohio State University*, 1975.

[141] T. Vincenty. Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations. *Survey review*, 1975.

[142] B. Wainwright, C. Cooke, and C. Low. Performance related technique factors in Olympic sprint kayaking. In *33rd International Conference on Biomechanics in Sport*, 2015.

[143] L. R. Weill. Conquering multipath: The GPS accuracy battle. *GPS World*, 1997.

[144] J. Wendel, O. Meister, R. Monikes, and G. Trommer. Time-Differenced Carrier Phase Measurements for Tightly Coupled GPS/INS Integration. *2006 IEEE/ION Position, Location, And Navigation Symposium*, 2006.

[145] A. Wieser. GPS based velocity estimation and its application to an odometer. 2007.

[146] A. Wieser and F. Brunner. SIGMA-F: variances of GPS observations determined by a Fuzzy system. *Vistas for Geodesy in the New Millennium*, 2002.

[147] A. Wieser and F. K. Brunner. An extended weight model for GPS phase observations. *Earth, Planets and Space*, 2000.

[148] E. Wikstrom and M. Tillman. A new force-plate technology measure of dynamic postural stability: the dynamic postural stability index. *Journal of Athletic Training*, 2005.

[149] J. Zhang, K. Zhang, R. Grenfell, and R. Deakin. GPS Satellite Velocity and Acceleration Determination using the Broadcast Ephemeris. *Journal of Navigation*, 2006.

[150] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney. UWB Systems for Wireless Sensor Networks. *Proceedings of the IEEE*, 2009.

[151] Y. Zhao, M. Becker, D. Becker, and S. Leinen. Improving the performance of tightly-coupled GPS/INS navigation by using time-differenced GPS-carrier-phase measurements and low-cost MEMS IMU. *Gyroscopy and Navigation*, 2015.