

**Resilient, Low-Cost Navigation in Urban Environments with GNSS
Precise Point Positioning, Inertial Measurement Unit and Precise Clock
Sensor Fusion**

Sihan Yang

**A Dissertation Submitted to the Faculty of Graduate Studies
In Partial Fulfilment of the Requirements
for the Degree of Doctor of Philosophy**

Graduate Program in Earth and Space Science

**York University
Toronto, Ontario**

**August, 2025
©Sihan Yang, 2025**

Abstract

The Precise Point Positioning (PPP) measurement processing technique for Global Navigation Satellite Systems (GNSS) is widely applied in scientific and commercial applications that require sub-metre level accuracy, particularly in areas with few obstructions. PPP is expanding into mass-market applications at the consumer level. However, the technique suffers from the inherent disadvantages of GNSS-based technologies. Obstructed environments, such as urban canyons, downgrade the solution. The study presents a novel solution that provides stronger resilience and improved performance in challenging environments by implementing clock-aided algorithms and augmenting the GNSS receiver with both an inertial measurement unit (IMU) and an external clock as a frequency reference. The PPP algorithm was augmented with clock-specific modifications, including the receiver clock coasting (RCC) and receiver clock modelling (RCM) algorithms. The study began with building a prototype triple-sensor fusion using higher-end hardware and was successfully tested by collecting and processing real-world, kinematic GNSS data. Horizontal rms errors have been reduced from 1.49 m with the traditional dual-sensor solution to 1.14 m with the novel solution. With multiple simulated outages, the horizontal rms errors were reduced from 2.07 m to 1.36 m using the RCM solution and 0.55 m using the RCC solution. To address the requirement of consumer-grade applications, a mass-market-friendly variant using low-cost hardware was built and tested in real urban canyons. The positioning accuracy over time was found to improve by 10–40% in most datasets when compared to traditional PPP/IMU fusion using identical hardware. In both fusion configurations, three-satellite GNSS solutions were verified. Finally, the study delved into position-domain Integrity Monitoring (IM) and showed an improved integrity performance of the novel triple-sensor solution. The number of unavailable epochs under the IM criteria was reduced by half.

The percentage of safe operations in the cross-track and along-track components, as determined by the criteria, increased from 97.20% and 82.34% to 99.45% and 92.67%, respectively, indicating a more trustworthy solution for safety-of-life applications. The result has proven the novel triple-sensor solution as a viable option in improving the performance and integrity for liability-critical applications such as intelligent vehicles. Further research is recommended to optimise the characterisation of the external oscillators and to achieve deeper integration among the three sensors.

Acknowledgement

My deepest gratitude goes to my supervisor, Professor Sunil Bisnath, who introduced me to the world of GNSS research and provided invaluable guidance in academics and in life throughout the past seven years. He was there to help me go through some of the most challenging times I have ever experienced. Beyond that, Dr. Bisnath has been tremendously helpful in building my soft skills from my first day in the lab. It was inconceivable back then to see myself confidently giving academic talks.

I would also like to express my gratitude to Garrett Seepersad and Professor Jianguo Wang, who have served on my supervisory committee and whom I have known for a long time. I have had the pleasure of working with Garrett on a project that spanned several years, during which I learned a great deal. Professor Wang has been giving me invaluable advice since my first introduction to the GNSS community and has taught me in the courses. I thank them, as the completed work would not have been the same without their efforts in my studies.

My lab colleagues through the years have been awesome friends and peers. While I have been happy working with everyone, a few names stand out because they have been particularly helpful to the completion of the dissertation. My thanks go to John, Sudha, Nacer, Ding and Jiahuan.

Finally, I would like to acknowledge and express my gratitude to my wife, Jing, for her unwavering and steadfast support over the years. Lastly, but not least, I would like to thank my mom and dad, as well as my mother-in-law and father-in-law, for their unwavering support and understanding, which enabled me to complete the work.

Contents

Abstract	ii
Acknowledgements	iv
Table of Contents	v
List of Tables	ix
List of Figures	x
1 Introduction	1
1.1 The Challenge of GNSS Technology in Novel Applications	3
1.2 Problem Statement	5
1.2.1 The Crucialness of the GNSS PPP Solution	6
1.2.2 Improving PPP in Challenging Environment through Sensor Fusion	7
1.3 Research Statement	8
1.4 Novelty and Significance	10
1.4.1 Hardware Integration – External Oscillator	10
1.4.2 Algorithmic Improvement	11
1.5 Dissertation Overview	13
2 Estimation Techniques in of York-PPP GNSS Processing	15
2.1 GNSS-PPP Observation Model and Error Sources	17

2.2	Kalman Filter	23
2.2.1	Key Definitions	24
2.2.2	Kalman Filter in Steps	25
3	PPP Sensor Fusion Techniques	32
3.1	Inertial Measurement Unit	34
3.1.1	IMU Measurements and Errors	35
3.1.2	IMU Navigation Equations	40
3.2	GNSS/IMU Fusion	48
3.2.1	GNSS-PPP/IMU Fusion Models	51
3.2.2	Measurement Model	53
3.3	External Clock as Aiding Sensors	56
3.3.1	Clocks in GNSS Receivers	57
3.3.2	Clock Frequency and Stability	59
3.3.3	Clock Coasting and Clock Modelling	65
3.3.4	Triple-Sensor Fusion	76
3.4	Hardware Assembly and Integration	77
3.4.1	Clock Discipline	79
4	Building the Prototype Triple-Sensor Fusion Solution	83
4.1	Chip-Scale Atomic Clock	84
4.1.1	Operational Principles	85
4.1.2	Physical Principles and Environmental Factors	85
4.1.3	Clock Discipline	88
4.2	Road Tests Set-up and Methodology	90
4.2.1	Experimental Set-up	91
4.2.2	Road Test Conditions	93
4.3	Results and Analysis	94
4.3.1	Baseline PPP Solution	95

4.3.2	Results with Tightly-Coupled Integration and Clock Aiding	97
4.3.3	Results with Simulated Outages	101
4.3.4	Summary	106
5	A Mass-Market Navigation System: PPP/IMU/OCXO Integration using Low-cost Hardware	107
5.1	Oven-controlled Crystal Oscillator (OCXO)	108
5.2	Setup, Road Tests, Results and Analysis	122
5.2.1	Experimental Setup	122
5.2.2	Results and Analysis	124
6	Interpretation in Real Life: Position-Domain Integrity Analysis	140
6.1	Overview of GNSS Evaluation Parameters	142
6.1.1	Parameters for Evaluating GNSS Solutions	142
6.1.2	Defining an IM Framework	146
6.2	Building an Integrity Monitoring Framework	147
6.2.1	Classic and Recent Innovations in RAIM	147
6.2.2	An Alternative Position-Domain Approach	149
6.2.3	Applicability of the Position-domain IM Framework Based on Extreme Value Theory	150
6.3	Integrity Analysis	155
6.3.1	Application in Real-world Data	158
6.4	Conclusions from Integrity Analysis	170
7	Conclusions and Recommendations for Future Research	171
7.1	Conclusions	171
7.1.1	Conclusions on the Prototype Triple-Sensor Solution	172
7.1.2	Conclusions on the Low-Cost Triple-Sensor Solution in Real Urban Canyons	173
7.1.3	Conclusions on Position-Domain Integrity Analysis	174
7.1.4	Summary of Research Contribution	175
7.1.5	Recommendations for Future Research	176

7.2	Recommendations for Future Research	177
7.2.1	Recommendations on the Sensor Fusion	177
7.2.2	Recommendations on Building and Testing the Triple-Sensor Solution	178
7.2.3	Recommendations on Evaluating the Solution and Integrity Monitoring	179
References		181

List of Tables

3.1	Relations of phase noise spectral density with Allan variance	63
4.1	Horizontal rms statistics of the PPP baseline solution in Run 1	97
4.2	Horizontal rms statistics of PPP-only, PPP/TC and PPP/TC/RCM solutions in Run 1	98
4.3	Horizontal rms statistics of PPP-only, PPP/TC, and PPP/TC/RCM solutions in Run 2	100
4.4	Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM and PPP/TC/RCC solutions in Run 1 with the four-satellite mask	103
4.5	Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM and PPP/TC/RCC solutions in Run 1 with the three-satellite mask	106
5.1	Horizontal rms statistics of PPP-only, PPP/TC and PPP/TC/RCM solutions in Run 3	127
5.2	Description of the artificial elevation masks used in Run 3	129
5.3	Horizontal rms statistics of for PPP-only, PPP/TC, PPP/TC/RCM, and PPP/TC/RCC solutions with different number of visible satellites in Run 3	132
5.4	Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM, and PPP/TC/RCC solutions in Run 4	135
5.5	Horizontal rms statistics of PPP/TC and PPP/TC/RCM solutions in the remaining datasets	138
6.1	GPD fitting parameters computed for each processing mode in Run 4	164
6.2	Integrity statistics of PPP-only, PPP/TC, and PPP/TC/RCM solutions in Run 4 specifying the percentage of safe operations, system unavailabilities and misleading epochs	169

List of Figures

2.1	Kalman filter flowchart	29
2.2	Kalman filter processing in time	29
3.1	Typical mechanism of an accelerometer	37
3.2	Typical mechanism of a gyroscope	38
3.3	Axes of the body frame	42
3.4	Euler angle rotations (Groves, 2013)	43
3.5	Flowchart of IMU mechanisation	48
3.6	Tightly-coupled sensor fusion scheme	50
3.7	Survey of typical AVAR performance (Fernandez & Calero, 2015; Teunissen & Montenbruck, 2017) . .	64
3.8	Theoretical behaviour of receiver clock coasting (blue) and EKF-based receiver clock modelling (red) in an urban canyon compared to the open-sky estimation (black)	75
3.9	Typical triple-sensor hardware flow during a field test	78
3.10	Clock discipline process using GPS Time	80
4.1	CSAC phase difference from reference and steering value during discipline	89
4.2	Road test setup with the CSAC-aided triple-sensor fusion	92
4.3	Roving sensors integrated with the vehicle rooftop	93
4.4	Google Street views of the routes: Highway 7 (above) and Pond Road (below)	94
4.5	Time series of horizontal errors in the GNSS-only, PPP solution - Run 1 + 2	96

4.6	Cumulative distribution of horizontal errors in the GNSS-only, PPP solution - Run 1 (left) and Run 2 (right)	96
4.7	Results from the dual- and triple-sensor integration: cumulative distribution of horizontal errors (top) and time series of horizontal errors (down)	98
4.8	A snapshot of trajectory of all processing methods in Run 1	99
4.9	Results from the dual- and triple-sensor integration: cumulative distribution of horizontal errors (top) and time series of horizontal errors (down)	100
4.10	Number of visible satellites and GDOP computation for Run 2	101
4.11	Sky plot of the four-satellite mask implemented on Run 1	102
4.12	Time series of horizontal errors over the first masked period in Run 1	103
4.13	Sky plot of the three-satellite mask implemented on Run 1	104
4.14	Time series of horizontal errors over the second masked period in Run 1	105
5.1	GXClok-500 package containing a u-Blox low-cost GNSS receiver chip and patch antenna	110
5.2	Typical OCXO operation	112
5.3	EK-5 evaluation board for GXClok-500 OCXO	113
5.4	PPS signals during a disciplining session	114
5.5	Status versus time (left) and internal frequency vs time (right) of the OCXO start-up sequence with external reference	116
5.6	Status versus time of a sample OCXO disciplining process (left) and internal frequency vs time for a sample OCXO disciplining process (right)	117
5.7	Status versus time of an OCXO holdover test (left) and internal frequency vs time of the OCXO holdover Ttst (right)	119
5.8	Status versus time of an OCXO free-run test (left) and internal frequency vs time of the OCXO free-run test (right)	120
5.9	Operation flowchart for an adaptive OCXO-aided GNSS receiver	121
5.10	Experimental setup used for testing low-cost triple-sensor fusion	123
5.11	Clock status over time during a road test in urban canyons	124

5.12 Google View of the underpasses involved in Run 3	126
5.13 Results from Run 3 using low-cost triple-sensor fusion: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)	127
5.14 Trajectory of all solutions in Run 3 during the two outages: railway underpass (left) and highway underpass (right)	128
5.15 Results from Run 3 with a 5-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)	130
5.16 Results from Run 3 with a 4-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)	131
5.17 Results from Run 3 with a 3-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)	131
5.18 Comparison of the horizontal rms errors across all solutions with various mask strengths in Run 3 . .	133
5.19 Google Street views of the urban routes: Yonge Street and the surrounding streets	134
5.20 Results from Run 4 on Yonge Street: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)	135
5.21 Number of visible satellites and GDOP in Run 4	136
6.1 Pyramid of the four parameters used to evaluate GNSS performance	142
6.2 Stanford Diagram (Tossaint et al., 2007) for evaluating integrity performance	144
6.3 Normalised histogram of PPP Horizontal Errors in Run 4	153
6.4 Normalised histogram of PPP/TC Horizontal Errors in Run 4	153
6.5 MRL plot for Run 3 without simulated outages	159
6.6 MRL plot for Run 3 with a 4-satellite mask	160
6.7 MRL plot for Run 4	161
6.8 Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC) with different L_0 values	162
6.9 Component-wise HPE and tail fitting with GPD for Run 4 (PPP-only)	163
6.10 Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC)	163
6.11 Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC/RCM)	164

6.12 Time series of component-wise HPL and HPE for Run 4 (PPP-only)	165
6.13 Time series of component-wise HPL and HPE for Run 4 (PPP/TC)	166
6.14 Time series of component-wise HPL and HPE for Run 4 (PPP/TC/RCM)	166
6.15 Stanford Diagram for Run 4 (PPP-only)	167
6.16 Stanford Diagram for Run 4 (PPP/TC)	168
6.17 Stanford Diagram for Run 4 (PPP/TC/RCM)	168

Chapter 1

Introduction

Modern Global Satellite Navigation System (GNSS) enables many time- and location-based services (LBS) by allowing a ubiquitous outdoor Position, Navigation, and Timing (PNT) solution available at a range of cost requirements. The GNSS technology stands out as a universal position-fixing method through ranging signals from satellite constellations, and thus supports a theoretically unlimited number of ground users and no marginal costs for each additional user. It is now seen in consumer products such as smart wearables and highly demanding applications such as spacecraft and aircraft guidance. The GNSS technology suffers from key drawbacks in modern urban applications due to effects on the ranging signals from obstructions (Geng et al., 2024; Groves et al., 2013). The problem worsens as emerging applications demand low-cost solutions, and low-cost hardware suffers more greatly from urban environments. The advent of low-cost, multi-constellation, multi-frequency GNSS receivers, as well as the advances in low-cost inertial measurement units (IMUs) with micro-electromechanical sensor (MEMS) technology, has improved the GNSS PNT in the said scenarios in recent years

(Abd Rabbou & El-Rabbany, 2015; Li et al., 2017; Vana & Bisnath, 2023). However, a gap still exists to bring the performance of the technology to the level demanded by emerging safety-of-life or liability-critical applications in urban environments.

The dissertation proposes a novel and low-cost triple-sensor fusion method based on GNSS-PPP (Precise Point Positioning) technology (Kouba & Héroux, 2001; Zumberge et al., 1997), an IMU, and an external frequency source in the form of a precise clock. Triple-sensor fusion is a step forward in closing the gap, providing an improved and tested solution for real-world urban canyons within a reasonable budget. The dissertation begins with an overview of relevant technologies and theories, followed by studies of the clock-aided GNSS approach, methodologies for field experiments and processing where a more costly prototype and a low-cost evolution are tested, and then followed by analysis of the performance of the setups, including a position-domain integrity framework for real-life considerations. The result shows an improvement in accuracy and resilience in real urban canyon scenarios compared to traditional solutions with dual sensors.

1.1 The Challenge of GNSS Technology in Novel Applications

GNSS signals are broadcast from satellite constellations, typically in Medium Earth Orbit (MEO). The satellites are tracked by the user equipment, and the signals are measured to produce (1) pseudorange, (2) carrier-phase, and (3) Doppler measurements (Farrell & Barth, 1999; Groves, 2013). The GNSS is a form of space-based, one-way radio-ranging navigation system. In such a navigation system, the user at an unknown location needs knowledge of the absolute position of the transmitter, i.e., the GNSS satellites, the transmitter clocks need to be synchronised with each other, and the receiver clock's unknown needs to be addressed. With enough availability of the ranging signals, the GNSS becomes available to the general population as a public service, allowing wide-range consumer applications. In addition, technological advancement led to cost reduction and miniaturisation of previously bulky GNSS user equipment in the last two decades. PNT solutions have since been proliferated for use in a range of technologies from professional industries to daily activities (Aggrey, Bisnath, et al., 2019; Vana & Bisnath, 2023; Yi et al., 2022).

If the reception of the ranging signals is good, GNSS solutions without any precise techniques reach an accuracy at the metre level. In an urban environment, however, where the usage of GNSS has become prevalent in the last two decades, GNSS suffers from degradation. These factors include (Groves et al., 2013; Teunissen & Montenbruck, 2017):

1. Signal Blockage: Buildings and large vehicles are common obstructions, blocking signals of low-elevation satellites and potentially leaving a narrow cone for signal reception.

2. **Multipath Errors:** Signals may be received by the antenna indirectly through the reflection of various surfaces. The reflected signals may interfere with the directly received signals, leading to noisy measurements. Multipath errors are especially common in urban canyons.
3. **Non-Line-of-Sight (NLOS) Signals:** Signals not directly reachable to the user antenna may be received through reflection on urban surfaces. If not correctly identified, and are instead processed, these signals will mislead the position estimation and lead to an increase in inaccuracy.
4. **Attenuation:** Some GNSS signals in urban environments are not sufficiently blocked but get attenuated through obstructions such as foliage. The degree of attenuation relates to the object's geometry and material. Signal interference, deliberate or not, can also cause a degradation effect on the GNSS signals.

These negative effects cause degradation of GNSS performance metrics, including accuracy, continuity, integrity, and availability. Modern GNSS applications require a range of performance levels. Less demanding LBS applications such as location-based information or push-ads require accuracy as low as tens of metres (Rizos, 2005). This level of performance can be achieved using basic positioning techniques in most scenarios or alternatively by cellular signal positioning if available (Spinney, 2003). For land navigation, such as road guidance, a few metres of accuracy is desired. At this level of accuracy, service outages are expected with low-cost hardware in urban environments. The magnitude of the accuracy requirement increases when greater autonomy is desired by the system. Lane-level navigation typically requires one-metre-level accuracy. The performance requirement becomes more systematic and stringent when it comes

to safety-of-life or liability-critical applications. Modern V2X (Vehicle-to-Everything) applications often have specific and constantly evolving performance indicators being developed by authorities (Standardization Administration of the People's Republic of China, 2024; U.S. Department of Transportation, 2025). Furthermore, unlike existing LBS, these applications also require the availability, integrity, and latency of the solution. These performance metrics will be detailed in later chapters.

To address the degradation of GNSS solutions in challenging and GNSS-denied environments, the user equipment fuses GNSS with other sensors, such as Inertial Measurement Units (IMUs), visual sensors, or terrestrial ranging systems (Melendez-Pastor et al., 2017). GNSS/IMU integration is the most widely adopted among sensor fusion schemes (Elsheikh et al., 2019; Shin & El-Sheimy, 2002; Vana & Bisnath, 2023).

However, while offering a great deal of resilience by extending GNSS solutions into short periods of GNSS-denied environments, IMUs exhibit drifts and biases over time, in addition to the calibration requirement and proper physical alignment. Vibration and temperature introduce errors to IMU measurements more than electronic sensors, especially on consumer-grade or automotive-grade IMUs at lower costs. The monetary cost of the system becomes a constraint in bringing the fused PNT solution to the level of performance demanded by emerging applications.

1.2 Problem Statement

The technology of the day presents a conundrum: Services enabled by the GNSS require improvement of GNSS PNT solutions in multiple metrics in challenging environments, while their growing cost constraints are making such improvement even more

difficult with low-cost hardware.

1.2.1 The Crucialness of the GNSS PPP Solution

The GNSS technologies were developed for military use in the 1980s, respectively in the Western Bloc and the Eastern Bloc, as the Global Positioning System (GPS) and GLONASS (GLObalnaya NAVigatsionnaya Sputnikovaya Sistema in Russian). Followed by its gradual opening to civilians and the establishment of the Chinese BeiDou and European Galileo constellations in the early 2000s. Regional systems such as the Quasi-Zenith Satellite System (QZSS) from Japan and the Navigation Indian Constellation (NavIC) work in similar principles. Spanning over four decades, the technology ages without fundamental alterations to its principles, maintaining its universal presence of ranging signals outdoors.

The modern challenge, as depicted in the previous section, has been a major theme of research in the navigation community. Alternatives to the GNSS have been discussed, and stellar progress has been made with technologies such as the 5G navigation (Mogyorósi et al., 2022). However, the crucialness of the GNSS lies in its principle of passive-ranging from space infrastructures. The passive ranging ensures inherent PNT availability without the need to build and maintain any land-based infrastructures except for those required for the operation of the satellites. Since availability is crucial in any safety-of-life applications, maintaining and improving a GNSS solution that the safety-of-life system may rely on and fall back on is necessary.

A similar argument may be presented for the continuous improvement and development of the PPP (Precise Point Positioning) over relative positioning techniques and hybridisation with other signals. The PPP is a precise GNSS technique that can be viewed

as the basic, stand-alone positioning technique of Single Point Positioning (SPP). A global network of stations provides key products such as precise satellite orbit and clock information that can be used to improve the positioning of stand-alone receivers. It reaches decimetre-level accuracy in an ideal environment. In contrast with Network Real-time Kinematic (NRTK) or hybridisation, the PPP requires no local land infrastructure. Therefore, for safety-of-life applications, a GNSS-PPP solution is preferred because not only is its solution availability outdoors guaranteed by the GNSS principles, but its accuracy is also widely maintained without concerns for land infrastructures (Bisnath & Gao, 2009; Choy et al., 2017).

1.2.2 Improving PPP in Challenging Environment through Sensor Fusion

IMUs are a type of dead-reckoning sensor, meaning they measure the user's motion with respect to the inertial frame. An IMU incorporates multiple accelerometers and gyroscopes (gyros). An Inertial Navigation System (INS) consists of IMUs and a navigation processor to produce a navigation solution through a set of navigation equations known as mechanisation equations. IMU hardware is now easier to mass-produce at a lower cost. Microelectromechanical systems (MEMS) technology allowed consumer-level IMUs to be used on smartphones and personal vehicles by providing more affordable options than traditional IMUs (Yang et al., 2023; Yi et al., 2022). GNSS/IMU integration has been widely studied and is the go-to option for most PNT applications in urban environments with demanding performance.

Sensor fusion with IMUs presents the same conundrum as the GNSS-only solutions. The low-cost requirement limits the performance of the GNSS/IMU sensor fusion while the industry continues to look for better performance that allows safety-of-life appli-

cations and enables more possibilities. Studies on improving the low-cost GNSS/IMU fusion have been a major topic in GNSS research, from the then-novel multi-GNSS, multi-frequency fusion with MEMS-IMUs to ultra-low-cost possibilities with smartphones. However, the challenge is still present, and the sensor fusion solution needs to be further improved.

1.3 Research Statement

To address the problem of improving the performance of modern GNSS solutions in urban areas, this research aims to develop a novel sensor fusion technology based on sensor fusion with an external frequency source. The problem of improving GNSS performance in challenging environments has been a regular and popular topic of research. The research community and GNSS solution providers have been addressing the problem by improving the GNSS base solution and by sensor fusion. Oftentimes, the solution involves an IMU due to its dead-reckoning ability. Other sensors, such as odometers, cameras or other visual or radar ranging sensors, have been investigated with varying levels of success and cost-effectiveness. The novel triple-sensor fusion scheme will include three low-cost sensors: a multi-constellation and multi-frequency GNSS receiver with a patch antenna, an automotive-grade IMU, and an external clock. The solution is produced through PPP for greater accuracy and applicability. For clarification, the 'low-cost' used in the context of this research refers to a consumer- or automotive-grade receiver with a patch antenna and IMUs. The triple-sensor fusion is designed to have a few qualitative requirements in mind:

1. Availability: the solution must be achieved without land infrastructures and thus be

able to serve as a backup solution in a safety-of-life or liability-critical navigation system.

2. Mass-market viability: the solution must be achieved with mass-market hardware components within a reasonable cost range and can be scaled down during mass production.
3. Performance: the solution must show improved performance over existing sensor fusions. The performance should further, if not completely, close the performance gap of low-cost sensor fusions from real-world requirements.
4. Resilience: the solution must show improved trustworthiness to withstand the adverse signal environment in urban canyons and similar challenging environments. It should be a step forward from the widely applied GNSS/IMU integration.

Therefore, the triple-sensor fusion is devised with the GNSS receiver as its core and is aided by fusing an IMU and an external clock. The IMU forms a tightly-coupled fusion with the GNSS measurements. The external clock, novel in the research, aids the system further by providing a stable frequency to the GNSS receiver, improving the measurement, and allowing clock algorithms in processing. The final block is the PPP techniques fused with the two aiding sensors. The PPP preserves the performance of the solution in the absence of infrastructures, provides decimetre-level kinematic solutions in more open-sky environments, and allows metre-level or better solutions in urban canyons. The Kalman filter is further enhanced by the clock algorithms specific for the external clock hardware, including receiver clock coasting (RCC) and receiver clock modelling (RCM) (Knabe & Kalafus, 1984; Krawinkel, 2018; Yang & Bisnath, 2020).

The goals of the triple-sensor fusion are:

1. To improve solution performance under poor geometry due to the presence of obstructions
2. To reduce error peaks/spikes caused by short outages common in urban environments
3. To achieve a meaningful three-satellite solution that helps bridge the gap between GNSS-denied environments and an obstructed environment.

The latter chapters of the dissertation will provide surveys on related aspects of triple-sensor fusion, details on implementation and methodology, and an analysis of various results to demonstrate the completion of the research goals. The research is heavily oriented towards automotive applications in defining low-cost hardware, the selection of sensors, and the demonstration of navigation results.

1.4 Novelty and Significance

The triple-sensor fusion is based on the well-studied dual-sensor fusion of GNSS/IMU, with the novelty of using the additional timing sensor as an aiding sensor for the GNSS. The research involves experimentally selecting, assembling, and configuring the hardware and, through software programming, building the algorithms into the Extended Kalman Filter (EKF) in York-PPP codes. Therefore, the novelty of the research can be approached from two perspectives.

1.4.1 Hardware Integration – External Oscillator

The frequency stability of a clock is commonly expressed using the Allan Deviation (ADEV) over a period of time τ , known as the averaging time. GNSS receivers typi-

cally adopt lower-end crystal oscillators (XOs), e.g. Temperature-Compensated Crystal Oscillators (TCXOs) for timing purposes. The clock aiding algorithm depends on a high-precision clock with good frequency stability, or a lower ADEV than the TCXOs.

There are two possible replacements within a reasonable cost range for the clock-aided GNSS: Chip-scale Atomic Clocks (CSACs) and Oven-Controlled Crystal Oscillators (OCXOs). Both oscillators will be introduced in detail in later chapters. Existing research has extensively used CSACs of various models, or their larger equivalent, the Miniturised Atomic Clocks (MACS) for RCC and RCM research (Fernández et al., 2017; Krawinkel et al., 2014; Meng et al., 2024) to improve the GNSS solution. However, the existing literature is limited as there are very few studies in further integration with inertial sensors, little emphasis on tackling challenging environments, and no attempts to reduce the cost of the clock-aided receiver scheme further.

This study addresses the blank by re-introducing the CSAC as the external oscillator to verify a triple-sensor fusion scheme with an IMU. Then the hardware is replaced by lower-cost alternatives, including the novel approach of using an OCXO in place of the precise external clock. Finally, the integrated setup is taken to real urban canyons for data collection and analysis. The project adopted Microsemi SA.45s as the higher-cost CSAC for building the initial prototype. Safran GXClok-500 was chosen as the low-cost OCXO timing module.

1.4.2 Algorithmic Improvement

Receiver clock errors are accustomed to being treated as an unknown variable along with the three spatial coordinates due to their significant effect on solution accuracy and inconsistency from various clock setups. GNSS receivers are typically equipped with

quartz clocks such as TCXOs. TCXOs are cheap but offer inferior frequency stability.

Chip-scale atomic clocks are the civilian descendant of U.S. DARPA projects to provide robust GNSS solutions at the personnel level in the event of war. They are still costly at a few thousand U.S. dollars, but compared to traditional atomic clocks such as hydrogen-maser clocks, CSACs are more portable and much cheaper. The latest variant of Microsemi SA.45s requires around 120 mW of power, not much larger than a U.S. dollar coin. A precise, external clock like CSAC improves GNSS solutions in a number of ways. For timing applications, the CSAC acts as a holdover clock to continuously provide timing signals as accurate as $1 \mu s$ over 24 hours despite frequency drifts without extra signals. The positioning potential of such an external clock has been investigated in the last decade. One way to utilise the clock is clock coasting. Similarly to holdover clocks in timing applications, the clock can be "coasted" to replace receiver clock error estimation in GNSS filters. Clock coasting leads to one less required satellite range measurement to produce an exact solution, and increased availability of GNSS solutions. Similarly, while coasting the clock, the clock bias can be estimated independently using receiver clock modelling in Kalman filters. Ideally, RCM will further reduce clock errors from the holdover clock if the modelling is correct. However, precise clock modelling is too difficult to achieve for mass-market applications due to cost and time constraints. The study employs RCM and clock coasting at the same time to investigate the performance of both methods. It is also in the interest of the study to investigate how a non-ideal parameter in the modelling affects RCM performance.

While there has been continuous interest in clock-aiding algorithms to improve the GNSS solution (Ma et al., 2015; Ramlall et al., 2011; Sturza, 1983), studies are limited in either applying the sensor fusion scheme in urban canyons (Meng et al., 2024; Yang &

Bisnath, 2020) or with IMU integration (Ma et al., 2016). The study presents the relevant theories, applications and results from the RCC and RCM algorithms in a triple-sensor setting with PPP as the GNSS processing method. Datasets were collected in various urban environments, including simulated GNSS outages and real busy urban canyons in Toronto for analysis.

1.5 Dissertation Overview

Following the introduction to the research, the subsequent chapters will gradually review the fundamentals, methodology, experimental procedures and results, and real-life implications of the research. These chapters are:

1. GNSS Positioning Techniques: a review of the fundamental algorithms and theories behind the research, including the precise techniques that provide the baseline solution.
2. PPP Sensor Fusion Techniques: the two aiding sensors are explained in detail through an overview of the IMU fusion techniques and studies in the external clock used in the research.
3. Prototype Analysis with PPP/IMU/CSAC integration using Geodetic-grade hardware: a prototype of the triple-sensor fusion is built for verification purposes using high-end hardware. The experimental results from the prototype are then analysed in detail. Research goals of performance improvement in urban canyons and three-satellite solutions are verified through simulated real-time post-processing and simulated outages.

4. A Mass-Market Navigation System with PPP/IMU/OCXO integration using low-cost hardware: the triple-sensor fusion has evolved by adopting mass-market hardware. The change leads to challenges that need to be tackled. Datasets from variable environments are collected. Validation from the previous, CSAC-aided setup is re-evaluated. Finally, real-world urban datasets in challenging environments are analysed.
5. Interpretation in Real Life through Position-Domain Integrity Analysis: the results from the triple-sensor fusion are further studied through the means of position-domain integrity analysis. The analysis completes the integrity metric evaluation of the triple-sensor fusion for real-life considerations.
6. Conclusions and Recommendations for Future Research: the research achieves the goals and leaves open a few possible ways for continuation.

Chapter 2

Estimation Techniques in of York-PPP GNSS Processing

GNSS signals are transmitted globally from MEO (Medium Earth Orbit) constellations, each with a network of at least 24 satellites. The signals are on the radio spectrum as part of the L-band with a wavelength between about 19 - 25 cm. The selection was specific for the PNT purposes allows an appropriate precision. The radio signals consist of at least two parts: a pseudorandom noise (PRN) code and a navigation message. The PRN code repeats at intervals from milliseconds to seconds and is used to produce observations and provide a unique identification to distinguish among satellites within the constellation. The navigation message stream, also known as the broadcast navigation message, transmits key information about the satellite itself, such as clock corrections and ephemeris information. The message also provides the timing components by syncing the receiver's local time with the GNSS system time. The receiver can thus compute the location of the satellite at the time of receiving and produce observations from the

PRN codes. With enough observations, the unknown position can be estimated through trilateration of the ranging signals (Kaplan & Hegarty, 2017; Teunissen & Montenbruck, 2017).

To make a measurement, the receiver locally generates a copy of the PRN sequence for comparison and alignment with the received GNSS signals. By the subtraction of the local time of receiving and signal generation, the receiver is able to measure the time t_r^s taken for the signal to propagate from the satellite to the receiver antenna. Since the GNSS signals are EM waves, they travel at the speed of light. Thus, the range of the transmission can be computed as,

$$\rho_r^s = ct_r^s \quad (2.1)$$

However, the product of the speed of light c and the time computed from the receiver t_r^s will vary from the true theoretical range ρ_r^s as the clock errors for the satellite Δt^s and the receiver itself Δt_r must also be accounted for. For this reason, the measurement is referred to as pseudorange. Since the satellite clock error can be computed through broadcast ephemeris or precise products, the term can be resolved individually alongside other error sources, leaving only a variable of receiver clock offsets. A common definition of the pseudorange can then be given.

$$p_r^s = \rho_r^s + c\Delta t_r \quad (2.2)$$

There are two other types of measurements typically made by the receiver, including the carrier phase and Doppler. The carrier phase measurement of the signals provides centimetre-level or higher precision and is essential for precise positioning. Measuring the carrier phase leads to a one-wavelength ambiguity that is at the centre of the

development of precise point positioning and RTK.

The Doppler measurement is caused by the relative motion of the receiver and the satellite, and is a measure of the line-of-sight velocity. It is more commonly used in applications where a high velocity relative to Earth is considered. The Doppler measurements can also be converted into the pseudorange rate (PRR) measurement (Groves, 2013).

2.1 GNSS-PPP Observation Model and Error Sources

The PPP technique provides superior positioning performance compared to SPP due to its use of multi-frequency pseudorange as well as carrier-phase measurements, precise satellite orbit and clock products, and additional products such as differential code biases (DCBs) (Liu et al., 2019; Montenbruck et al., 2014). While PPP requires station networks to estimate the above-mentioned corrections, it is considered a more financially viable option to maintain than the RTK or NRTK techniques. With differential techniques, the user receives corrections corresponding to observations (Naciri, 2022). Given the number of constellations and satellites available, an ideal NRTK operation becomes computationally and data-wise intensive. The differential techniques also require a large number of local stations to maintain. In addition, the communication between the user and base stations is two-way instead of one-way communication for PPP, as there are intermediary control centres to regulate the communications based on the user's rough location. In such regards, the PPP processing may be viewed as a natural evolution of the standard GNSS positioning as it preserves the core characteristics of universality and close-to-negligible marginal burdens per user.

The PPP observation equation used in this research assumes an uncombined model,

and the equations for the pseudorange and carrier-phase are (Aggrey & Bisnath, 2019; Katsigianni et al., 2019; Kouba & Héroux, 2001):

$$\begin{aligned} p_r^s &= \rho_r^s + c(\Delta t_r - \Delta t^s) + T + \gamma_i I_i^s + (b_{r,i} - b_i^s) + \epsilon_{p_i} \\ \phi_r^s &= \rho_r^s + c(\Delta t_r - \Delta t^s) + T - \gamma_i I_i^s + \lambda_i (N_i^s + B_{r,i} - B_i^s) + \epsilon_{\phi_i} \end{aligned} \quad (2.3)$$

The new terms introduced in the equations are:

T : Tropospheric delay;

$\gamma_i = \frac{f_1^2}{f_i^2}$: Ratio of frequencies applied to the first frequency ionosphere delay I_i at frequency i ;

$b_{r,i}$: Pseudorange receiver bias;

b_i^s : Pseudorange satellite bias;

$\lambda_i = \frac{c}{f_i}$: Wavelength of the GNSS signal;

N_i^s : Integer ambiguity on frequency i ;

$B_{r,i}$: Phase receiver bias;

B_i^s : Phase satellite bias;

ϵ_{p_i} : Unmodelled errors in pseudorange such as multipath

ϵ_{ϕ_i} : Unmodelled errors in phase such as multipath

Depending on the method to deal with some errors, some of the terms are to be estimated instead of computed from observables or accessed from data product providers. This presents additional errors when not enough observations are available, leading to a rank deficiency. The following list examines the methods to mitigate the errors and how they differ from PPP and SPP processing.

Satellite Orbit and Clock Errors

Most commonly, the broadcast orbit that is either acquired from GNSS stations or received by the receivers themselves is used. The broadcast orbit usually has orbital accuracy errors in the range of metres. The broadcast information also includes clock correction data.

The research adopts the precise .sp3 orbit product commonly used in PPP processing, and the file name extension refers to Extended Standard Product- 3, a product originally proposed in 1989 (Elsheikh et al., 2023; Hilla, 2016). The SP3 products provide orbital accuracy of less than 5 centimetres, allowing precise techniques and providing a more accurate position estimate. A separate clock file is used alongside the orbit correction product, and the clock correction file similarly provides a more precise and accurate clock error for the post-processing (Kouba, 2009).

Tropospheric Error

The refractive index of the atmosphere increases as altitude decreases. GNSS Signals slow down and get refracted while entering the troposphere, a lower part of the atmosphere. Therefore, it is necessary to model the troposphere delay as the signal shows a different behaviour from travelling in a complete or near vacuum. The error modelling also includes the stratosphere as part of it due to its similarity in modelling. The tropospheric delay is separated into "dry" and "wet" components. The "dry" component makes up 90% of the tropospheric delay and is modelled relatively easily (Hopfield, 1969). The "wet" component will also have a contribution from water vapour and is subject to improvement in research.

Ionospheric Error

GNSS satellites are typically located at MEO, meaning that the signal will travel through the entirety of the ionosphere. The free electrons from the ionised particles in the ionosphere affect the EM signal propagation. The effect exhibits weather-like behaviour from solar ultraviolet radiation, solar wind, and electrodynamics in the thermosphere, among other factors. Traditional models such as the Klobuchar model are commonly used in SPP (Klobuchar, 1975). To bring the solution accuracy to centimetre-level, however, it is crucial to improve the ionospheric model, or to estimate the ionospheric error as a variable in the EKF. In this research, the ionospheric errors are estimated as variables in the EKF. The ionospheric error is consistent in any localised area, which means the user can receive a better estimation from nearby, higher-quality, stationary GNSS stations. GIM (global ionospheric maps) is a new tool exploiting this fact. The maps allow a mapping of the total electron content (TEC) globally. GIM provides a useful tool to improve the ionospheric error modelling for PPP (Ho et al., 1997; Jee et al., 2010; Reddybattula et al., 2019).

Hardware Biases

The receiver or satellite hardware biases result from internal processing that delays the transmission. On the satellite end, the different frequencies and modulations have an effect, and therefore corrections such as Differential Code Bias (DCB) may be processed for GNSS processing (Montenbruck et al., 2014; Xiang et al., 2020). These code and phase corrections are additive to the user's observations. On the user-end, the hardware bias is seldom studied due to the various nature of the user's equipment and the errors are

absorbed into the estimation process.

Antenna Offsets

The signals are broadcast from the phase centres of the satellite antennas rather than the equipment's centre of mass. Therefore, the PPP user should take the displacement into consideration. The research utilises the ANTEX (Antenna Exchange) product that provides information for common satellites and user receivers (Prange et al., 2023; Schmid et al., 2016). Low-cost and ultra-low-cost receiver users commonly disregard the user-end antenna offset and absorb the errors into the estimation.

Relativistic Effect

The relative motion of the satellites and the user under different gravitational fields leads to a relativistic effect in GNSS processing. The relativistic effect leads to both a propagation delay due to space-time curvature and deviations in satellite clock frequency. The error is mitigated through correction based on computing the instantaneous satellite position and velocity. The relativistic clock effect on the receiver clock is neglected as it is present in all measurements and becomes a component in the receiver clock offset (Ashby, 2003; Hegarty et al., 2017).

Site Displacement

These refer to a series of error sources often at mm to cm level due to Earth movement (Aggrey, 2014; Naciri, 2022; Witchayangkoon, 2000).

1. Plate motion: movement of tectonic plate (Altamimi et al., 2012)

2. Solid Earth tides: a function of the gravitational motion of the Earth and other masses in the solar system such as the Sun. The error is due to the fact that the Earth is not rigid and thus affected by these gravitational sources. The errors affect the vertical positioning more than the horizontal and can be modelled (McCarthy, Petit, et al., 2004).
3. Ocean loading: the ocean tides cause displacement of the Earth's crust. The impact can usually be seen more obviously along the coastal area at a few centimetres (Agnew, 1997).
4. Other site displacement effects: these include minor error sources, including pole tides (Niu et al., 2021) and atmospheric loading (Tunalı & Özlüdemir, 2019). They typically contribute at a millimetre level at most and therefore are not particularly addressed.

Phase wind-up

The GNSS signals are right-hand circularly polarised. A rotating antenna relative to the user leads to a change in carrier-phase measurement, and due to the satellite motion relative to the user, the phase wind-up is not negligible, as it could lead to a few decimetres of errors. The error source can be computed with the carrier-phase measurement, and it does not affect pseudorange measurements.

The observation equations of the two observables p_i^s and ϕ_i^s and the error sources modelling make up the basis of the PPP processing. The PPP operates on similar principles to SPP, but with additional corrections that are made available by the use of precise products (Wu et al., 1993).

Furthermore, as constellations such as BeiDou-3 and Galileo become available and widely supported in a range of mass-market receivers, the low-cost PPP benefits from multi-constellation and multi-frequency processing. The availability leads to improved geometry and availability of signals, and accelerates convergence time, helping PPP to approach RTK techniques in terms of convergence time. The caveat of low-cost multi-constellation PPP is that the PPP solution is greatly affected by the quality of measurements, and low-cost receivers do not always present the same level of measurement quality over all constellations or frequencies. As in the case with smartphones, and less prevalent with consumer-grade receivers, the user often finds gaps in the measurement for certain frequencies. The phenomenon leads to systematic errors in the processing and therefore leads to a worse solution if some signals are included.

Utilising the available measurements with corrections applied, a PNT solution can be optimally estimated using one of the positioning techniques. The post-processing is called *filtering*, referring to the algorithms estimating the state variables based upon all past measurements. It is called *optimal* estimation as it is statistically defined to minimise the estimation error. In the next section, the filtering technique used in the research is introduced.

2.2 Kalman Filter

There are two primary optimal estimation techniques used in GNSS processing, namely the Kalman Filter (KF) and Least Squares estimation methods, such as Sequential Least Squares (SLS). The research uses KF exclusively, as the algorithm is able to allow the fusion of information from various sensors.

The Kalman Filter requires knowledge of a system model - a deterministic property of the dynamic system, statistical properties of the system, and measurements to achieve an optimal estimation based on the knowledge. These components describe the art of making, tuning, and running any KF (Brown & Hwang, 1997; Gelb, 1974). The physical property of the system must be correctly modelled. In the case of GNSS processing, the dynamics are widely studied and determined. The tuning process requires knowledge of the measurement. This component presents itself more notably with low-cost sensors as they suffer from great and less predictable random errors. Finally, the measurements themselves are part of the estimation. Systematic errors not properly modelled will lead to greater estimation errors.

2.2.1 Key Definitions

The variables subject to estimation form the *state vector*, denoted by $\mathbf{x}$. The estimated state vector is denoted by $\hat{\mathbf{x}}$. In GNSS processing, the number of states to be estimated depends on the technique. Four fundamental variables are the x , y , z coordinates in ECEF (Earth-centred, Earth-fixed coordinate system) and the receiver time offset. Given the availability of modelling options and the requirement for precision, other error sources, such as ionospheric errors, may be estimated through the KF as well. The error covariance matrix presents the uncertainties of the states. Through the error covariance matrix, if there is not enough information to independently estimate the state vector, the correlation information from the matrix allows an estimate before the next available measurement.

The Kalman Filter runs recursively with time evolution. GNSS epochs are often measured at 1 - 5 Hz. For a discrete system such as GNSS measurements by epochs, a discrete

variation of the KF is used. Conversely, a continuous dynamic system requires the KF to run continuously. A *measurement vector*, denoted by $\mathbf{z}$, is expected at each epoch, representing the latest measurements coming from the sensors. The measurement vector is a different set of variables from the state vector, where their relation is described by the *measurement model*. The measurement vector is associated with a *measurement noise* matrix. The measurement noise, as its name suggests, refers to the instrumental noise generated by the sensors. For example, a very-low-cost cellphone IMU is expected to show greater accelerometer noise than that of a tactical-grade IMU (Yang et al., 2023).

2.2.2 Kalman Filter in Steps

After initiation, the iterative process of a Kalman Filter begins with *system propagation*. The system propagation phase produces a *prediction*, an estimate of the state vector at a future time, of the state vector and error covariance matrix. The prediction will utilise the known properties of the system with the most recently available measurement. Following the system propagation, the *measurement update* will take into account the new measurement vector as it comes in from the sensors. The measurement updates correct the prediction by updating the deterministic and stochastic components of the measurement model. *Kalman Gain* is computed to provide an optimal estimation of the states using the information of the uncertainties of current states and measurement noise. The covariance matrix is then updated afterwards.

To derive the KF equations, the problem is abstracted into the following statement: let $\mathbf{x}$ be a vector of non-random constants based on k sets of noisy measurements $\mathbf{z}_i$ sequenced $i = 1, 2, \dots, k$, and $\mathbf{z}_i = \mathbf{x} + \mathbf{v}_i$. The subscripts i indicate discrete epochs in the time sequence. The estimation $\hat{\mathbf{x}}$ is supposed to be unbiased, meaning its expected value

is identical to the estimated value; of minimum variance, meaning that its error variance is less than or equal to any other estimation. Then, the system model has the following form:

$$\mathbf{x}_k = \Phi_{k-1}\mathbf{x}_{k-1} + \mathbf{w}_{k-1} \quad (2.4)$$

Here Φ_k denotes the transition matrix at time k , and $\mathbf{w}_k$ denotes the system process noise with covariance Q_k .

It is assumed that there are m measurements available, that is, the measurement vector $\mathbf{z}$ is an $m \times 1$ vector. And the state vector $\mathbf{x}$ is an $n \times 1$ vector. Denoting the $m \times n$ measurement matrix H and the measurement model follows:

$$\mathbf{z}_k = H_k\mathbf{x}_k + \mathbf{v}_k \quad (2.5)$$

The additional term $\mathbf{v}_k$ is a vector with a dimension of $m \times 1$. $\mathbf{v}_k$ denotes an additive random measurement noise with covariance R_k . In GNSS processing and other non-linear applications, an extended Kalman filter (EKF) provides a linearised estimation for the measurement matrix.

The idea is that at time k , a state estimate $\hat{\mathbf{x}}_k(-)$, where the minus sign indicates a prediction, is computed and then updated with the current measurement $\mathbf{z}_k$. The state estimate after the measurement update, denoted $\hat{\mathbf{x}}_k(+)$ is additive from two gain vectors, that is

$$\hat{\mathbf{x}}_k(+) = K'_k\hat{\mathbf{x}}_k(-) + K_k\mathbf{z}_k \quad (2.6)$$

The two gain matrices K' and K shall be derived. First, the estimated state vectors are

the sum of a 'true' state and an error component denoted with a tilde:

$$\begin{aligned}\hat{\mathbf{x}}_k(+) &= \mathbf{x}_k + \tilde{\mathbf{x}}_k(+) \\ \hat{\mathbf{x}}_k(-) &= \mathbf{x}_k + \tilde{\mathbf{x}}_k(-)\end{aligned}\tag{2.7}$$

Using the equation 2.7 and 2.5 to expand 2.6, the error estimation after an update is

$$\tilde{\mathbf{x}}_k(+) = [K'_k + K_k H_k - I] \mathbf{x}_k + K'_k \tilde{\mathbf{x}}_k(-) + K_k \mathbf{v}_k\tag{2.8}$$

The estimator is unbiased, giving $E[\mathbf{x}_k(\pm)] = 0$, and by definition $E(\mathbf{v}_k)] = 0$, so the following must be true:

$$\begin{aligned}K'_k + K_k H_k - I &= 0 \\ K'_k &= I - K_k H_k\end{aligned}\tag{2.9}$$

Updating 2.7, the state vector after measurement update can be expressed as

$$\begin{aligned}\hat{\mathbf{x}}_k(+) &= (I - K_k H_k) \hat{\mathbf{x}}_k(-) + K_k \mathbf{z}_k \\ &= \hat{\mathbf{x}}_k(-) + K_k [\mathbf{z}_k - H_k \hat{\mathbf{x}}_k(-)]\end{aligned}\tag{2.10}$$

To find the error covariance, 2.10 can be used to expand $P_k(+)$, and by eliminating the non-correlated terms, we get

$$\begin{aligned}P_k(+) &= E[\hat{\mathbf{x}}_k(+) \hat{\mathbf{x}}_k(+)^T] \\ &= (I - K_k H_k) P_k(-) (I - K_k H_k)^T + K_k R_k K_k^T\end{aligned}\tag{2.11}$$

The question of finding the optimal K_k , the Kalman gain matrix, remains. Therefore,

the minimum of the cost function $J = E[\hat{\mathbf{x}}_k(+)^T S \hat{\mathbf{x}}_k(+)] \forall S$ must be found. Let $S = I$, and take the partial derivative with respect to K_k , then it follows

$$\frac{\partial J_k}{\partial K_k} = \frac{\partial \text{tr}[P_k(+)]}{\partial K_k} \quad (2.12)$$

Setting 2.12 to zero yields

$$K_k = P_k(-) H_k^T [H_k P_k(-) H_k^T + R_k]^{-1} \quad (2.13)$$

and therefore,

$$P_k(+) = (I - K_k H_k) P_k(-) \quad (2.14)$$

The state estimate in 2.10 can finally be updated. Here, the concept of measurement innovation $\delta \mathbf{z}_k^-$ at time $t = k$ can be introduced. The measurement innovation signifies the difference between the true measurement vector and what is computed from the predicted state vector before the measurement update.

$$\begin{aligned} \hat{\mathbf{x}}_k(+) &= \hat{\mathbf{x}}_k(-) + K_k [\mathbf{z}_k - H_k \hat{\mathbf{x}}_k(-)] \\ &= \hat{\mathbf{x}}_k(-) + K_k \delta \mathbf{z}_k^- \end{aligned} \quad (2.15)$$

The concept of measurement residual is similarly defined as the difference between the actual measurement vector and the computed measurement vector from the state vector, after the measurement update:

$$\delta \mathbf{z}_k^+ = \mathbf{z}_k - H_k \hat{\mathbf{x}}_k(+) \quad (2.16)$$

Figure 2.1 shows a flowchart of the Kalman Filter spanning from the final step at time $t = k - 1$ to the final step at time $t = k$.




Figure 2.1: Kalman filter flowchart

The Kalman filter builds up predictions based on past estimates, and the predicted state becomes updated by the latest measurements as they come in. Such procedural steps allow Kalman filters to be customised according to sensors of varying rates and availability. To illustrate a kinematic position estimation process, Figure 2.2 shows an example trajectory and its KF prediction and estimation as it propagates and is corrected by measurements. The green trajectory is the theoretical truth. The ellipse shows the error covariance of the estimated position.




Figure 2.2: Kalman filter processing in time

In GNSS or other fields where the measurement model is simply assumed with the linear transition matrix H , the EKF is used. The EKF is a non-linear variation of the standard KF. In an EKF, the state transition function F (a continuous, locally linear estimate of Φ), and measurement function H are linearised through the first-order partial derivative with respect to $\mathbf{x}$. The linearization may cause instability in certain cases and require more work on tuning the EKF to improve the stability and performance. For low-cost and ultra-low-cost receivers and IMUs, the sensitivity becomes more prominent due to greater and less predictable errors in measurement that invalidate theoretical assumptions.

Tuning with a very-low-cost receiver system, such as smartphones will thus often require a larger initial P , to compensate the greater state uncertainties before running EKF; an inflated Q , as the issues due to cheap hardware lead to unknown system errors invalidating the EKF's system model; and a moderately increased R to reflect the worse quality of measurements. The need for tuning is typically greatest with lower-cost hardware in more challenging environments, as the system model is most compromised.

In the least-squares estimation, the optimality comes from choosing an $\mathbf{x}$ to achieve the minimization of the sum of squares of $\mathbf{z} - \hat{\mathbf{z}}$. The optimality is achieved by minimising the following value, weighted by the matrix R ,

$$J = (\mathbf{z} - H\hat{\mathbf{x}})^T R^{-1} (\mathbf{z} - H\hat{\mathbf{x}}) \quad (2.17)$$

and taking the derivative $\frac{\partial J}{\partial \mathbf{x}} = 0$, giving

$$\hat{\mathbf{x}} = (H^T R^{-1} H)^{-1} H^T R^{-1} \mathbf{z} \quad (2.18)$$

The research will primarily adopt low-cost industrial/automotive-grade sensors. Following the earlier master's work on a prototype clock-aided GNSS-PPP solution using clock coasting (Yang & Bisnath, 2020), the Ph.D. research will bring the prototype into a functional localisation subsystem that shows performance improvements in resilience in urban canyons. The previous work on clock-aided GNSS-PPP solution utilised a geodetic GNSS receiver and antenna, and a Chip-Scale Atomic Clock (CSAC) to show that a three-satellite solution can be achieved with an external holdover clock, and a solution under poor geometry of 4-5 satellites can be improved.

Chapter 3

PPP Sensor Fusion Techniques

GNSS provides a ranging positioning solution outdoors. In many cases, however, aiding sensors are required to either improve the solution in general, to provide a backup, or to allow solutions in deep urban environments or GNSS-denied environments. Such solutions require fusion techniques in combining the inputs from all sensors, oftentimes coming at different rates with various quality, together in such a way that they estimate the best possible solution.

Examples of additional sensors include IMUs, barometers, magnetic compasses, and odometers (Zhu et al., 2020). The aiding sensor is often a type of dead-reckoning sensor, meaning that it operates continually without external information. GNSS/IMU navigation is the most prevalent form of such sensor fusion. In a highly integrated system like smartphones, the navigation engine may opt to utilise information from cellular signals. Barometers provide information regarding altitude; magnetometers or magnetic compasses assist pedestrian navigation by quickly initialising the heading angle without continuous velocity computation (Domínguez et al., 2015; Perul & Renaudin, n.d.). In

vehicular systems, odometers are a native sensor that helps correct dynamic variables such as speed with the sensor on wheels. Vehicles may also use map matching using the images from cameras. Autonomous vehicles typically have a wider range of perception sensors available for localised road guidance and collision avoidance.

While GNSS suffers from mostly modelable error sources, each aiding sensor presents a unique set of challenges of systematic errors. Practical analysis such as integrity monitoring (IM) become much more complex with multiple sensors involved. Some external sensors suffer from environmental factors such as temperature or atmospheric pressure that typically have no effect on GNSS processing. Identifying common and distinct error sources and designing the filter such that the errors are properly modelled or estimated is essential for a successful sensor fusion. Similar to the realm of GNSS receivers, low-cost external sensors, such as a smartphone-grade IMU, present noisier measurement behaviour and suffer from greater bias growth over time. Unlike in ideal situations where, e.g., an IMU solution perfectly replaces a GNSS solution when a vehicle enters a tunnel, engineering a navigation system with low-cost hardware often risks feeding noisy or faulty information from one sensor to another. Understanding and sometimes conforming to the challenges becomes part of the filter design.

Another practical challenge is the integration of hardware from different grades of manufacturers. The interaction exists in three levels: (1) raw communication, (2) user-controlled communication, and (3) output integration. Raw communication refers to the hardware connection, sending and receiving signals from the other sensor, and utilising it internally, outside of user control. This realm of interaction is essentially a black-box in which individual users can only rely on information from manuals to investigate or deduce from outputs. In higher-grade and often more customizable hardware, the users

are given more information on the hardware interaction and may have the ability to control it with console commands or software. Finally, the interaction may only exist at the post-processing level, such as the case in loosely- and tightly-coupled GNSS/IMU integration, where the separate outputs of the two sensors are fused in the Kalman filter to produce a solution, and the two sensors do not need to communicate at all.

This chapter introduces the two aiding sensors used in the dissertation. The first section introduces the GNSS/IMU integration techniques. The second section provides a detailed introduction to external clocks, where the research presents a novel aiding sensor possibility. Then, the theoretical and practical aspects of building a GNSS/IMU/clock triple-sensor fusion navigation system are discussed.

3.1 Inertial Measurement Unit

The IMUs are dead-reckoning sensors, and thus, the user's motion is computed within the body frame. They alone do not provide an absolute positioning solution, as the body-frame axes must be resolved with external axes, such as in an ECEF axes using GNSS information. IMUs consist of accelerometers and gyroscopes, also known as gyros. Accelerometers measure the specific force the body endures, and the gyros measure the angular velocity at which the body is turning. Each of the sensor are made up of 3 sub-sensors for each axis, leading to 6 sensors in total. The research adopts strapdown IMUs, meaning that the IMU is fixed with respect to the user or the body (Elsheikh et al., 2020; Groves, 2013).

The development of microelectromechanical systems (MEMS) technology allowed small and low-cost IMU sensors to be produced. Moreover, their performance suffers

from greater bias growth and noisier measurements. The IMUs are categorically divided into a few grades. The highest grade of IMUs is used by the military, such as naval ships, submarines, and intercontinental ballistic missiles. Their cost goes up to a million USD per unit, and offers a daily drift of no more than a couple of kilometres per day. Aviation-grade IMUs cost less than 100,000 USD typically, and offer similar solution drift per hour, rather than per day. For users, a more common IMU grade is the tactical-grade IMUs used in expensive civil purposes or unmanned drones (Petovello, 2003). They sacrifice long-term performance but can rely on other positioning systems such as GNSS. In the ultra-low-cost category are the automotive-grade and consumer-grade IMUs. Automotive-grade IMUs cost around a hundred dollars and have the size of a SIM card (Ahmed & Tahir, 2016). Smartphones adopt consumer-grade IMUs with the lowest cost and smallest size at the cost of performance. The literature among IMU studies varies in its terminology. Typically, tactical-grade or lower IMUs are considered low-cost. To avoid confusion, the dissertation refers to the consumer-grade IMUs seen on wearable devices and smartphones as ultra-low-cost.

3.1.1 IMU Measurements and Errors

An IMU produces two separate types of measurement: the specific force and the angular rate. A 6-axis IMU measures each with three axes for a complete three-dimensional motion measurement. The internal IMU processor converts the internal measurement to output units such as ms^{-2} for accelerometers and $rads^{-1}$. It also validates and performs a range check of the output. The processor can also compensate for known errors in the output. For example, some sensors are sensitive to temperature, and a relationship between the temperature and output has been found (Ali, 2016). In this case, the IMU

may incorporate a temperature sensor to compensate for the processor's errors.

Accelerometers

There are a number of ways to construct accelerometers. They can be mechanical, piezoelectric, or capacitive. Figure 3.1 shows an example of such a mechanism. It shows a simplistic capacitive accelerometer commonly found on MEMS-IMUs (El-Sheimy & Niu, 2007; Zwahlen et al., 2012). Without external forces, the movable mass remains still in its neutral position. When an acceleration along its sensitive axis is experienced by the sensor, that is, there is an acceleration acting on the body it is strapped to, the movable mass moves along either the + or - direction. The movement of the mass causes a small but detectable change in the capacitance between the electrodes. The displacement of the movable mass can then be computed from the change. The physical or electrical limit of the mechanism leads to the detection range of the accelerometer. Accelerometers typically measure their range and measurement in terms of gravity g . Damping is often necessary to limit the responses and to protect the sensor in the event of shock. Sensors shown in Figure 3.1 can be connected in serial so that multiple measurements can be taken at once.




Figure 3.1: Typical mechanism of an accelerometer

Gyroscopes

Gyroscopes measure the rotation experienced with respect to the inertial frame. They output the angular velocity in radians per second. Like accelerometers, gyroscopes operate with a variety of principles, such as vibratory or optical. Figure 3.2 shows the basic principle of an optical gyroscope with four waveguides forming a closed loop when a light beam passes. If the waveguides rotate clockwise, in the inertial frame, the beam will travel a longer distance due to the Sagnac effect. If the waveguides rotate counter-clockwise, the beam will travel a shorter distance. The difference in travelled distance is used to measure how the body has rotated between each signal.




Figure 3.2: Typical mechanism of a gyroscope

IMU Measurement Errors

There are four main categories of IMU errors: biases, scale factors, cross-coupling errors and random noises. Errors in GNSS observables are primarily from the natural environment, while the IMU errors are closely related to the hardware itself and user dynamics.

Run-to-run errors are the fixed error contributions that remain consistent during each run but vary between IMU start-ups. Such errors could be a result of initialisation or incorrect calibration. The errors cannot be corrected by the sensor itself and should be corrected through integration or realignment.

In-run errors are the errors that develop and vary during the run. It is another example of errors that cannot be corrected by the IMU itself, such as accumulated bias. They may not be obvious with a short data collection and may only be noticed in the long run. The duration of short or long is a relative term, as the grade of the IMU greatly impacts the drift.

Some systematic errors are calibrated by the IMU itself, such as the size effect on

accelerometers or temperature. Many IMUs allow the user to manually perform an alignment of the IMU or perform calibrations with GNSS integration. In the former option, the user fixes the IMU on a flat and static surface. Knowing the body is static, the existing biases and drifts over time can be recorded by the IMU processor and compensated during actual runs. Some IMUs have the ability to integrate GNSS outputs or signals. In many cases of low-cost sensors, a cheap GNSS chip and patch antenna are provided in the same package, or integrated into the sensor. The sensors are then expected to calibrate themselves with the GNSS solution.

Bias is the offset between the actual physical measurement and the sensor readings. It is independent of the specific force or angular velocity at the time. Biases exist in all accelerometers, and gyros are considered to be the most prominent form of IMU errors (Abd Rabbou & El-Rabbany, 2015; Groves, 2013). The accelerometer bias $\mathbf{b}_a$ and gyro bias $\mathbf{b}_g$ can be split into the sum of static and dynamic components. The static component represents the fixed systematic biases that are instrumental, remain between each run, or persist after calibration. The dynamic component is the in-run biases that cannot be calculated before start-up, but typically contributes less than the static component. Bias in low-cost IMUs is typically in the order of $10^{-2} \sim 10^{-1} \text{ ms}^{-2}$. For low-cost gyroscopes, the bias is typically in the range of $10^{-6} \sim 10^{-4} \text{ rads}^{-1}$. For comparison, military-grade IMUs for marine purposes typically have $\mathbf{b}_a \sim 10^{-4}$ and $\mathbf{b}_g \sim 10^{-9}$ (Matthews, 1990; Yuksel et al., 2010).

Scale factors exist where the actual input-to-output ratio differs from the theoretical truth, causing a disproportionate output from expectation. Cross-coupling errors exist where a misalignment affects the orthogonality of the axes, leading to a cross-axis sensitivity. Although these two types of errors have different physical natures, they are

typically considered mathematically together.

Random noises are from various, uncategorized sources, e.g., mechanical instabilities and electrical noises. MEMS-IMUs are exceptionally vulnerable to random noises and may exhibit high-frequency noises. These noises present themselves as larger white noises over seconds than higher-grade IMUs, such as in the case of smartphone IMUs. They are preferred to be resolved with higher measurement noise tuning in sensor fusion than low-pass filters, as filters may erase essential dynamics that lead to instability of the EKF (El-Rabbany & El-Diasty, 2004; Skaloud et al., 1999).

3.1.2 IMU Navigation Equations

The measurements from an IMU are processed through a series of navigation equations to produce position, velocity, attitude, and sometimes acceleration information. A position fix in an Earth coordinate system requires external information given from other sensors, typically from GNSS. Oftentimes, off-the-shelf inertial systems have built-in navigation processors to directly give the user the IMU states, and these combined systems are called inertial navigation systems (INS). INS and IMU are seen interchangeably in much of the literature. However, in the GNSS fusion processing for this research, the raw output from accelerometers and gyroscopes is processed using the navigation equations programmed in York-PPP, and then fed into the EKF. Therefore, the dissertation adopts the GNSS/IMU fusion terminology.

Frame of Motion and Coordinates Transformation

Defining a coordinate system requires the definition of the three non-coplanar axes and an origin. The fusion of ranging navigation systems and dead-reckoning navigation

systems requires a transformation of different coordinate frames.

GNSS processing is done in the ECEF frame, where the z axis points along the Earth's axis of rotation, and the x axis points to the intersection of the equator with the IERS (International Earth Rotation System Service)'s Reference Meridian in a right-hand coordinate system.

GNSS solutions are also presented in the local frame by custom. In a local navigation frame, an origin is defined to provide the East, North and Up (ENU) position of the user. In kinematic processing, the origin is usually selected by the starting point. The local frame is chosen small enough as a surface, so that the geometry can be considered as 'flat'. Local frames do not work around poles as the ENU directions become ill-defined.

Dead-reckoning sensors such as IMUs resolve motion within the body frame. With a strap-down IMU, the body frame is equivalent to the IMU frame, as the IMU is considered fixed relative to the body. In a practical GNSS/IMU navigation system, the antenna phase centre does not coincide with the IMU centre. In such a case, the user should measure the physical distance of the two as lever arms and incorporate them into the processing.

Figure 3.3 shows a right-handed body frame. The convention is that the forward direction is labelled as x_b , and the rotation around the body's x_b axis would be the roll. The downwards direction, assumed along the direction of gravity, is labelled z_b , and rotation around the axis would be the yaw vector. Given the right-handedness, the remaining y_b vector points rightwards, and its rotation is the pitch vector.




Figure 3.3: Axes of the body frame

The velocity of motion is described by the vector $v_{\beta\alpha}^\gamma$, meaning that the velocity of frame α with respect to frame β is resolved about frame γ . To illustrate, the body's velocity v_{ib}^i in its own inertial frame is always zero.

Describing the body's attitude is more difficult, and there are a few different ways to accomplish it. Quaternions have been a popular mathematical tool for this purpose due to their computational efficiency. Another method, as used in the research, is the Euler angles. Euler angles describe the body's attitude with three sequential rotations for aligning frame β to α , and these are yaw $\psi_{\beta\alpha}$, pitch $\theta_{\beta\alpha}$, then roll $\phi_{\beta\alpha}$. The sequence is crucial for Euler angles as the rotational operators do not commute. Figure 3.4 shows an illustration (Groves, 2013) of Euler angles.




Figure 3.4: Euler angle rotations (Groves, 2013)

The Euler angles are tools to build up coordinate transformation matrices (CTMs), which will be used to describe attitude changes with IMU. For the transformation of vector $\mathbf{x}^\alpha$ to $\mathbf{x}^\beta$, the convention is

$$\mathbf{x}^\beta = C_\alpha^\beta \mathbf{x}^\alpha \quad (3.1)$$

The CTMs are also referred to as the direction cosine matrix (DCM) in some texts. Using Euler angles defined earlier, the DCM can be expanded into its complete rotational form:

$$C_\beta^\alpha = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \phi_{\beta\alpha} & \sin \phi_{\beta\alpha} \\ 0 & -\sin \phi_{\beta\alpha} & \cos \phi_{\beta\alpha} \end{pmatrix} \begin{pmatrix} \cos \theta_{\beta\alpha} & 0 & -\sin \theta_{\beta\alpha} \\ 0 & 1 & 0 \\ \sin \theta_{\beta\alpha} & 0 & \cos \theta_{\beta\alpha} \end{pmatrix} \begin{pmatrix} \cos \psi_{\beta\alpha} & \sin \psi_{\beta\alpha} & 0 \\ -\sin \psi_{\beta\alpha} & \cos \psi_{\beta\alpha} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad (3.2)$$

Small angle perturbation is a common approximation technique; by applying small

angle perturbation, 3.2 becomes

$$C_{\beta}^{\alpha} \approx \begin{pmatrix} 1 & \psi_{\beta\alpha} & -\theta_{\beta\alpha} \\ -\psi_{\beta\alpha} & 1 & \phi_{\beta\alpha} \\ \theta_{\beta\alpha} & -\phi_{\beta\alpha} & 1 \end{pmatrix} \equiv I_3 - [\Psi_{\beta\alpha} \wedge] \quad (3.3)$$

Here, $[\Psi_{\beta\alpha} \wedge]$ represents the skew-symmetric matrix of Euler angles.

IMU State Update

Dynamic processes are continuous; the body accelerates over time and turns over time. The IMU outputs the best measurement to describe the user's motion from time t to $t + \tau$, where τ is typically a short time increment. For reference, the IMU used in this research has a data output of 100 Hz. Similar to how EKF assumes a linearization for GNSS observations, the IMU navigation equations assume an infinitesimal change between measurements (Groves, 2013). The assumption plays a critical role in the derivation of equations. A concise version used in the York-PPP GNSS measurement processing software is explained in the subsection.

The attitude solution is updated with the instantaneous angular velocity measurement ω_{ib}^b , and the DCM of transforming the body attitude from t to $t + \tau$ must be found. Using the time derivative of the DCM, the instantaneous rotational transformation from the body axes to the ECEF axes is

$$\dot{C}_b^e = \lim_{\delta t \rightarrow 0} \left[\frac{C_b^e(t + \delta t) - C_b^e(t)}{\delta t} \right] \quad (3.4)$$

It can be assumed that the ECEF reference frame is stationary through the rotation. Thus,

the term $C_b^e(t + \delta t)$ can be decomposed as a successive rotation where the object frame b changes over time:

$$\begin{aligned}
C_b^e(t + \delta t) &= C_{e(t)}^{e(t+\delta t)} C_b^e(t) \\
&= (I_3 - [\Psi_{e(t)e(t+\delta t)} \wedge]) C_b^e(t) && \text{Using (3.3)} \\
&= (I_3 - \delta t [\omega_{be}^e \wedge]) C_b^e(t) \\
&= (I_3 - \delta t \Omega_{be}^e) C_b^e(t)
\end{aligned} \tag{3.5}$$

Thus, Equation 3.4 can be rewritten as, utilising the invertibility of DCMs,

$$\begin{aligned}
\dot{C}_b^e &= -C_b^e \Omega_{be}^e = C_b^e \Omega_{eb}^b \\
&= C_b^e \Omega_{ib}^b - \Omega_{ie}^e C_b^e
\end{aligned} \tag{3.6}$$

The two skew-symmetric matrices Ω_{ib}^b and Ω_{ie}^e correspond to the measured angular rate and the Earth rotation, respectively. By integrating 3.6 from t to $t + \tau$, and using the (+) and (−) notation to represent states before and after measurement, the following can be derived:

$$\begin{aligned}
C_b^e(+) &= C_b^e(-) \exp(\Omega_{eb}^b \tau) \\
&= C_b^e(-) \exp(\Omega_{ib}^b \tau) - C_b^e(-) [\exp(\Omega_{ie}^b \tau) - I_3] \\
&\simeq C_b^e(-) (I_3 + \Omega_{ib}^b \tau) - \Omega_{ie}^e C_b^e(-) \tau
\end{aligned} \tag{3.7}$$

where the final step assumes a constant angular velocity during the interval. The final expression requires knowledge of previous attitude, the angular measurement from gyros, and the measurement interval. The term Ω_{ie}^e is a known constant.

The acceleration experienced in the inertial frame is the sum of the specific force,

that is, the force on unit mass except for gravity, and the gravity at the body's position. At time t , it can be expressed as:

$$f_{ib}^e(t) = C_b^e(t) f_{ib}^b(t) \quad (3.8)$$

where $f_{ib}^b(t)$ is the measurement from the accelerometer, which is the average of the specific force measurement from time t to $t + \tau$ given by the IMU processor. The program takes another assumption here that the average of the attitudes before and after should also be taken, then

$$f_{ib}^e = [C_b^e(-) + C_b^e(+)] f_{ib}^b / 2 \quad (3.9)$$

To find the position update, the velocity should be updated. The time derivative of velocity is

$$\dot{v}_{eb}^e = \dot{C}_b^e \dot{r}_{eb}^e + a_{eb}^e \quad (3.10)$$

Here a denotes acceleration, and r denotes the body's position. One cannot simply assume that the time derivative of velocity is equivalent to acceleration, as the resolving frame may not necessarily be the same as the reference frame. However, in this case of transformation to the ECEF frame, the former term equals zero, then Equation 3.10 becomes,

$$\begin{aligned} \dot{v}_{eb}^e &= a_{eb}^e \\ &= f_{ib}^e + g_b^e(r_{eb}^e) - 2\Omega_{ie}^e v_{eb}^e \end{aligned} \quad (3.11)$$

The acceleration that a unit mass of the body experiences is the sum of the specific force, gravity at its location, and the Coriolis effect. An approximate velocity update can be

derived:

$$v_{eb}^e(+) = v_{eb}^e(-) + [f_{ib}^e + g_b^e(r_{eb}^e(-)) - 2\Omega_{ie}^e v_{eb}^e(-)]\tau \quad (3.12)$$

The Coriolis term $-2\Omega_{ie}^e v_{eb}^e(-)$, always perpendicular to the velocity, is usually neglected outside of space applications or high-end inertial sensors, as their magnitude is small compared to the specific force term f_{ib}^e or the gravity term $g_b^e(r_{eb}^e(-))$. Sometimes the applications require a precise update and the sensors provides enough precision, in which case the term will remain.

The position update follows another assumption that the velocity changes linearly over the interval.

$$r_{eb}^e(+) = r_{eb}^e(-) + v_{eb}^e(-)\frac{\tau}{2} + v_{eb}^e(+)\frac{\tau}{2} \quad (3.13)$$

Equation 3.13 concludes the final step of the mechanisation process for IMU navigation. The steps are illustrated in Figure 3.5. To begin the IMU mechanisation process, an approximate initial condition must be given. In the GNSS/IMU fusion method used in the study, the first 100 GNSS epochs are ignored, as these converging solutions may not provide a good approximation for the IMU states. Then, the position state is simply initiated from the GNSS epoch; the initial velocity vector can be initiated by estimation or linear position difference. The difficulty lies in the attitude. While the engine can use the velocity to approximate an attitude, at least one of the Euler angles cannot be initialised. With land vehicles, however, the user can assume that the roll angle is zero and the pitch angle is within a reasonable range.

In the study, the IMU states for GNSS/IMU fusion include the bias terms for the gyroscope and accelerometer. The estimated biases are added to the sensor input.




Figure 3.5: Flowchart of IMU mechanisation

The derivation of the IMU navigation equations leaves room for inaccuracy when the linearity assumption fails. However, the study does not involve dynamics requiring high precision. Moreover, given the measurement noise with low-cost IMUs, expanding the Taylor series further is unlikely to produce practical benefits.

3.2 GNSS/IMU Fusion

There are three primary types of GNSS/IMU fusion, categorised by the domain in which the IMU information is used by the system. The loose coupling presents a cascaded architecture where the position solutions are computed separately for the two sensors and fed through an EKF. The EKF needs input from both navigation filters to run,

and that means the possible rank deficiency in challenging environments will disable the fusion solution. Tight coupling presents a range-domain method where the GNSS measurements become the input over an expanded state vector that includes the IMU-exclusive terms. It is an example of a centralised integration where sensor measurements are the input. Finally, a deep coupling involves hardware-level interaction, where the GNSS tracking process combines IMU involvement (Farrell & Barth, 1999; Groves, 2013).

The research has adopted the tightly-coupled (TC) GNSS/IMU fusion. The TC method is chosen because, firstly, even if the user enters a challenging environment where the satellite number is too low to keep the GNSS solution without rank deficiency, the GNSS measurements can still aid in correcting IMU states along the way. Secondly, the IMU mechanisation solution is used as the prediction for the integrated EKF, giving a more timely prediction of the user's current state. The mechanisation solution is further utilised in the ranging processing of the GNSS measurements. It exploits the advantages of IMUs as dead-reckoning sensors and avoids the problems with loose coupling. Deeply coupled fusion requires hardware-level interaction that requires compatible sensors, where the researcher's role becomes limited. Therefore, the TC fusion serves as an ideal method to proceed.

Figure 3.6 shows the flow of a tightly-coupled sensor fusion scheme. The dashed rectangles depict two separate loops: the EKF loop and the IMU loop. The EKF loop runs at the rate of GNSS observations and incorporates the latest IMU solution from the IMU loop. The IMU loop runs at a different frequency than IMU measurements come in, which is typically a much higher frequency than GNSS observations. The IMU mechanisation uses the solution from the last fused epoch to initialise, and uses the IMU errors from EKF estimation to correct the input readings. In the case of missing

GNSS epochs, the IMU loop keeps on running and updating the user's position based on IMU readings. If the IMU fails to provide a meaningful solution, usually blocked by the program for sanity check purposes, the EKF becomes a conventional GNSS EKF and uses the correction data to gradually improve the IMU solution.




Figure 3.6: Tightly-coupled sensor fusion scheme

3.2.1 GNSS-PPP/IMU Fusion Models

Fusing GNSS and IMU outputs in a tightly coupled configuration requires the design of a state vector that incorporates the GNSS-PPP states as well as the necessary terms for IMU corrections. Despite the fact that PPP is used in the research, ambiguities are not resolved and instead estimated as float values. The reason for the design choice is that the focus is a metre- to decimetre- level solution using low-cost sensors in very challenging environments. The IMU states of position δr , velocity δv , and attitude $\delta \psi$ are to be corrected by the EKF. Since the GNSS already estimates position and velocity, the problem remains in attitude and IMU biases.

IMU State Propagation

There are five IMU states that are considered in the research, the position error δr_{eb}^e , velocity error δv_{eb}^e , attitude error $\delta \Psi_{eb}^e$, the biases for accelerometers and gyroscopes b_a and b_g . Each of the vectors is three-dimensional, so the IMU state vector is a 15×1 dimensional vector (Groves, 2013; Vana & Bisnath, 2023).

$$x_{imu} = \begin{bmatrix} \delta r_{eb}^e \\ \delta v_{eb}^e \\ \delta \Psi_{eb}^e \\ b_a \\ b_g \end{bmatrix} \quad (3.14)$$

To find the transition matrix, the expectation value of the error states must be found.

For example of attitude:

$$\begin{aligned}
E(\delta \hat{\Psi}_{eb}^e) &\simeq E(\hat{C}_b^e(\dot{\omega}_{eb}^b - \omega_{eb}^b)) \\
&\simeq \hat{C}_b^e \delta \omega_{ib}^b - \hat{C}_b^e(\tilde{C}_e^b - C_e^b) \omega_{ie}^e \\
&= E(\hat{C}_b^e \delta \omega_{ib}^b - \Omega_{ie}^e \delta \Psi_{eb}^e) \\
&\simeq -\Omega_{ie}^e \delta \Psi_{eb}^e + \hat{C}_b^e b_g
\end{aligned} \tag{3.15}$$

The last step uses the fact that the expected rate of change for attitude error should come from the gyroscope biases. The expression can be rephrased into a row in the transition matrix F_{imu} , namely

$$\begin{bmatrix} 0_{1 \times 3} & 0_{1 \times 3} & -\Omega_{ie}^e & 0_{1 \times 3} & \hat{C}_b^e \end{bmatrix} \tag{3.16}$$

Similarly, the rest of the transition matrix can be derived:

$$F_{imu} = \begin{bmatrix} 0_{1 \times 3} & I_3 & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} \\ F_{21}^e & -2\Omega_{ie}^e & F_{23}^e & \hat{C}_b^e & 0_{1 \times 3} \\ 0_{1 \times 3} & 0_{1 \times 3} & -\Omega_{ie}^e & 0_{1 \times 3} & \hat{C}_b^e \\ 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} \\ 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} \end{bmatrix} \tag{3.17}$$

The simplified terms are:

$$F_{21}^e = -\frac{2\gamma_{ib}^e \mathbf{r}_{eb}^{\hat{e}^T}}{r_{eS}^e \hat{L}_b |\hat{\mathbf{r}}_{eb}^e|} \quad F_{23}^e = [-(\hat{C}_b^e \hat{\mathbf{f}}_{ib}^b) \wedge] \tag{3.18}$$

The newly introduced terms γ_{ib}^e and $\hat{L}_b$ are, respectively, the gravitational force and

the longitude of the body. Due to long derivations, the dissertation will refer to the details in (Groves, 2013) at equation (2.142).

The discrete transition matrix Φ_{IMU} is derived by taking linearisation followed by a multiplication by the time interval τ :

$$\Phi_{imu} = \begin{bmatrix} I_3 & I_3\tau & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} \\ F_2^e 1\tau & I_3 - 2\Omega_{ie}^e \tau & F_{23}^e \tau & \hat{C}_b^e \tau & 0_{1 \times 3} \\ 0_{1 \times 3} & 0_{1 \times 3} & I_3 - \Omega_{ie}^e \tau & 0_{1 \times 3} & \hat{C}_b^e \tau \\ 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & I_3 & 0_{1 \times 3} \\ 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & 0_{1 \times 3} & I_3 \end{bmatrix} \quad (3.19)$$

3.2.2 Measurement Model

The state vector becomes the following:

$$\delta \mathbf{x} = [\delta r \quad \delta v \quad \delta \Psi \quad \delta t_r \quad \dot{\delta t}_r \quad d_{trop} \quad b_a \quad b_g \quad N_1 \quad \dots \quad N_n \quad dI \quad r_{dcb}] \quad (3.20)$$

The terms are respectively: δr is the three-dimensional position error in ECEF;
 δv is the three-dimensional velocity error in ECEF;
 $\delta \Psi$ is the attitude error in roll, pitch and yaw radians;
 δt_r is the receiver clock offset;
 $\dot{\delta t}_r$ is the receiver clock drift;
 d_{trop} is the tropospheric wet delay b_a, b_g is the accelerometer and gyroscope biases;
 $N_1 \dots N_n$ is the carrier phase ambiguities;
 dI the estimated ionospheric errors;
 r_{dcb} is the differential code bias per constellation. It acts together with dI for ionosphere

constraining.

The measurement vector includes the code pseudorange and carrier-phase difference between GNSS and IMU prediction, then for satellite i , that is

$$\mathbf{z} = \begin{bmatrix} \rho^i - \hat{\rho}^i \\ \phi^i - \hat{\phi}^i \\ \dots \end{bmatrix} \quad (3.21)$$

The predicted measurements of $\hat{\rho}^i$ and $\hat{\phi}^i$ are produced from IMU mechanisation. The measurement model in the PPP can be extended to include additional measurements, such as GIM delays and receiver DCBs, optionally. Therefore, the design matrix for the measurement model is:

$$H = \begin{bmatrix} (\rho^i)_x & (\rho^i)_y & (\rho^i)_z & 0_{1 \times 3} & 0_{1 \times 3} & 1 & 0 & (\rho^i)_{zpd} & 0_{1 \times 3} & 0_{1 \times 3} & 0 & \frac{f_1^2}{f_2^2} & \dots \\ (\phi^i)_x & (\phi^i)_y & (\phi^i)_z & 0_{1 \times 3} & 0_{1 \times 3} & 1 & 0 & (\phi^i)_{zpd} & 0_{1 \times 3} & 0_{1 \times 3} & \lambda & -\frac{f_1^2}{f_2^2} & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \end{bmatrix} \quad (3.22)$$

The notation $(\rho^i)_x$ denotes the partial derivative of ρ^i with respect to x .

EKF Constraints

The York-PPP is programmed to incorporate constraints for land vehicle applications. These include the zero-velocity update (zupt) and zero-angular-rate update (ZARU) constraints (Li et al., 2012; Qin et al., 2014; Vana et al., 2020). These constraints are designed to exploit the physical nature of a land vehicle's motion.

In ZUPT, if the vehicle is considered 'stopped', then its position should be static. In urban canyons, the poor GNSS signal reception and quality often lead to the vehicle's position solution drifting within an area of a metre-level radius. The ZUPT requires detecting stationary scenarios. The study invokes horizontal velocity detection criteria, where a velocity less than 0.5 m/s is considered stationary. A better and more native option is to add the odometer of the vehicle to sensor fusion.

If stationary detection is triggered, the following measurement matrix kicks in for the IMU terms:

$$H_{ZUPT} = \begin{bmatrix} \mathbf{0}_{1 \times 3} & -I_3 & \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} \end{bmatrix} \quad (3.23)$$

and the measurement innovation becomes:

$$\delta z_{ZUPT}^- = -\hat{v}_{eb}^e \quad (3.24)$$

ZUPT also benefits the solution in that the IMU attitude and biases can be corrected for the duration. The selection of 0.5 m/s as the threshold is an educated guess to filter out position changes from estimation errors, and any user movement in the vehicle is unlikely to be affected.

The stationary detection scheme is further used in the zero-angular-rate update (ZARU). The method of ZARU follows ZUPT by assuming the car doesn't spin or rotate along any axes, if it has a low enough velocity. ZARU may also be performed when the vehicle travels at a constant velocity. However, such scenarios are rare in real driving scenarios and are not implemented. ZARU can be particularly useful to calibrate the attitude errors, particularly the heading angle, with low-cost IMUs. With ZARU, the

measurement matrix becomes

$$H_{ZARU} = \begin{bmatrix} 0_{1 \times 3} & 0 & 0_{1 \times 3} & 0_{1 \times 3} & -I_3 \end{bmatrix} \quad (3.25)$$

And the measurement innovation becomes:

$$\delta z_{ZARU}^- = -\omega_{ib}^b \quad (3.26)$$

3.3 External Clock as Aiding Sensors

The primary novelty of the research is the introduction of a novel aiding sensor, an external clock. Clocks have been historically crucial in the development of GNSS technology. Clocks are oscillators that provide a frequency reference for the GNSS receivers and satellites. While frequency sources are important in many systems, they are especially crucial in GNSS navigation. The reason is that the dynamics of GNSS are based on EM wave transmissions, which travel at the speed of light. There is no perfection in timing performance to measure the travel time of the fastest possible speed for today's science. The Universal Time Coordinated (UTC) used widely in the PNT community is an example of a time reference where the output is rigorously controlled by top-grade atomic clocks, such as hydrogen masers (Brown, 1991). Unlike the clocks for satellite or user equipment, these atomic clocks have less concern regarding their cost, size and wattage.

Space GNSS segments adopt space-qualified clocks, which are usually atomic clocks customised to be reliable, stable and with a reasonable wattage for space use. Time synchronicity is critical in the GNSS constellations and is an active research area. The

GPST (GPS Time) is a sophisticated time scale that does not come from any single master clock but rather is a timing system made up of the physical clocks on GPS satellites and ground segments (Leschiutta, 2005; Lewandowski & Arias, 2011). GPST is designed to be in the range of $1\mu\text{s}$ of UTC time modulo leap seconds. Clock corrections are transmitted from ground stations for satellites to broadcast. The stability of the satellite clocks determines how long the corrections are effective without expiring. Precise clock data is made available as data products from IGS servers and used for the research's PPP study.

User equipment requires clocks operating with less expensive principles. They also need to be smaller, consume less power, and be more portable. That said, a high-precision clock on GNSS receivers has always been a low priority for most applications. The receiver clock error, or offset, is always modelled as a variable and part of the estimation process. The process links the receiver clock offset stochastically with other spatial variables x, y, z in the resolving coordinates.

The study aims to change the perception that the clock hardware in GNSS receivers is not as significant. The development of lower-cost, lower-wattage, and high-performance oscillators allows a clock upgrade for GNSS receivers of most price ranges. This section will begin with introducing the hardware of GNSS receiver clocks, followed by a review of the frequency stability of the clocks, then the theoretical aspects of the two algorithms, the receiver clock coasting (RCC) and receiver clock modelling (RCM), and finally, the main theme of the study, the triple-sensor fusion, is presented.

3.3.1 Clocks in GNSS Receivers

Quartz crystal oscillators (QXOs) are the most common form of consumer clocks. In QXOs, the piezoelectric quartz crystals produce electric potentials from pressing. A

shape change from the mechanical forces leads to the generation of an electric field. Its elasticity returns the crystal to its original shape, and a signal is electrically emitted. Using the physical property in conjunction with its natural resonant frequency, a stable frequency source can be built. The output frequency is then adjusted to provide a stable frequency, such as 10 MHz or 1 pps (pulse-per-second) signals. The mechanism is relatively simple and easy to implement, but it is also vulnerable to slight environmental changes, such as *temperature* changes that affect the physical properties of the crystals. *Shock* and *mechanical stress* are other possible factors that disrupt the frequency stability of the oscillator. Furthermore, the crystal is subject to *ageing*, where the physical properties change over time. Ageing has a long-term effect that is observed in months or years. All of these factors further affect the *frequency stability* of the clock (Frerking, 1996; Lam, 2008).

Modern hardware solutions adopt the temperature-compensated crystal oscillator (TCXO) more commonly. In TCXOs, there are additional components that compensate for temperature changes and insulators to seal the quartz crystals from the outside temperature. The evolution of microelectronics allows TCXOs to be the most prevalent form of clocks in GNSS devices (Kusters & Vig, 1991; Vig, 1992). Oven-controlled Crystal Oscillators (OCXOs) are a step further in regulating the ambient temperature of quartz crystals. In OCXOs, an 'oven', or a thermally insulated chamber that actively has heating elements, encapsulates the crystal and sometimes the entire circuit. Unlike conventional QXOs, which operate around ambient temperature, OCXOs require warming up to an ideal operating temperature, typically higher than the room temperature. OCXOs used to cost more than TCXOs, have a higher wattage and are larger in dimensions. They also provide the best stability for the QXO family, only second to atomic clocks. Traditionally,

OCXOs are found in ground radio stations or cellular base stations.

In the most recent decade, the revolutionary technology of chip-scale atomic clocks (CSACs) became available to the general public as a by-product of military research in the 2000s (Knappe et al., 2005; Lutwak et al., 2002a). The CSAC uses MEMS technology to allow historically massive atomic clocks to fit into coin-sized chips by creating millimetre-sized caesium capsules. The rubidium atomic-vapour discharge lamp used in old atomic clocks is replaced by vertical cavity surface emitting lasers (VCSELs) of much smaller sizes (Kitching, 2018; Lutwak et al., 2004). The technology was developed by the U.S. to improve the army soldiers' positioning capability when GPS were not available. They became commercially available in the early 2010s as a high-performance but mobile oscillator for space, aeronautics and PNT applications. However, CSACs are nonetheless expensive for mass-market products, with price tags in the range of 3,000 U.S. dollars.

The study recognises that the CSACs provide a promising future for PNT timing equipment. However, it is yet feasible to mass-produce the CSACs for the consumer market. Since the study aims to build a new consumer-grade navigation system that is possible for the industry to adopt, the study will begin with CSACs as a higher-standard external clock in transition with existing literature, then move on to a much lower-cost OCXO with the latest technologies.

3.3.2 Clock Frequency and Stability

The main performance criterion of clocks in PNT is frequency stability. Most commonly, a clock's frequency stability can be described with the *Allan Deviation* (ADEV) or *Allan Variance* (AVAR). The name comes from the early work in categorising frequency characteristics by (Barnes & Allan, 1966) in the 1970s. The Allan methods are sample

time-average approaches — other measures of clock stability approach from time or frequency domain spectrum analysis. ADEV or AVAR are metrics that measure the deviation of running frequency with nominal frequency over a specified time interval τ . The value of τ varies in the order of 10s. For short-term performance, the stability between $\tau = 10^0 \sim 10^2$ is the main criterion. For longer-term performance, one looks at $\tau = 10^2 \sim 10^4$ depending on the periods for which the clock runs undisciplined.

Frequency and Oscillations

The frequency of an oscillator is denoted by ν , with a period of τ . They form the fundamental inversely proportional relationship of oscillations (Blair, 1974; Teunissen & Montenbruck, 2017):

$$\nu = \frac{1}{\tau} \quad (3.27)$$

At time t , the voltage signals can be written as

$$V(t) = (V_0 + \epsilon(t)) \sin(2\pi\nu_0 t + \phi(t)) \quad (3.28)$$

where V_0 is the nominal voltage amplitude and ν_0 the nominal frequency. $\epsilon_V(t)$ and $\phi(t)$, respectively, are time-dependent variations to the amplitude and frequency output. The instantaneous frequency can be found by taking the time derivative of the triangular term, that is

$$\begin{aligned} \nu(t) &= \frac{1}{2\pi} \frac{d}{dt} [2\pi\nu_0 t + \phi(t)] \\ &= \nu_0 + \frac{1}{2\pi} \frac{d\phi}{dt} \end{aligned} \quad (3.29)$$

Allan Variance

To study the frequency fluctuations, the initial frequency of the oscillator is irrelevant, and therefore, one may define the *fractional frequency* $y(t)$ as the following:

$$y(t) = \frac{\nu(t) - \nu_0}{\nu_0} = \frac{1}{2\pi\nu_0} \frac{d\phi}{dt} \quad (3.30)$$

Additionally, the time error of the clock $T_c(t)$ is modelled with the following:

$$T_c(t) = x_0 + y_0 t + \frac{1}{2} a t^2 + \int_0^t E(t') dt' + \epsilon_t(t) \quad (3.31)$$

where x_0 represents the initial time offset from reference; y_0 represents the frequency offset; D_0 represents a constant frequency drift; $\int_0^t E(t') dt'$ encapsulates the remaining variables such as environmental effects that cause a time offset, then finally followed by ϵ_t representing the random errors.

The Allan variance is derived to characterise random frequency fluctuations (Allan, 2005). Let the time-averaged fractional frequency be $\bar{y}$, and the averaging time be τ , the expression for any time-averaged fractional frequency over τ is:

$$\bar{y}_i = \frac{x_{i+1} - x_i}{\tau} = \frac{1}{\tau} \int_{t_i}^{t_i+\tau} y(t) dt \quad (3.32)$$

The Allan Variance for an averaging time τ is defined by (Beard & Senior, 2017):

$$\sigma_y^2(\tau) = \left\langle \frac{(\bar{y}_{k+1} - \bar{y}_k)^2}{2} \right\rangle \quad (3.33)$$

It is also known as a 2-sample variance. The Allan variance is related to another approach to clock stability in the frequency domain.

Spectral Density of Clocks

In the alternative frequency domain approach, one considers the one-sided spectral density of fractional frequency $S_y(f)$, phase spectral density $S_x(f)$, where f represents the Fourier frequency. A generalised model applicable to most clocks is derived as the following (Rutman & Walls, 1991; Walter, 2002):

$$S_x(f) = \begin{cases} \sum_{\beta=-6}^0 \frac{h_{\beta+2}}{4\pi^2} f^{\beta} & \text{for } 0 < f < f_h \\ 0 & \text{for } f > f_h \end{cases} \quad (3.34)$$

In the generalised expression, a finite sum of superposed 7 spectra represents the random fluctuation for clocks. The β terms are integers, with corresponding constants indicating the spectral level of noise g_{β} . A low-pass filter f_h provides a high-frequency cut-off and the frequency f_h is also referred to as the high corner frequency.

The pure power-law noises are named as follows, based on the power β :

- (1) f^0 : white phase (WHPH);
- (2) f^{-1} : flicker phase (FLPH);
- (3) f^{-2} : random walk phase (RWPH);
- (4) f^{-3} : flicker frequency (FLFR);
- (5) f^{-4} : random walk frequency (RWFR);
- (6) f^{-5} : flicker drift (FLDR);
- (7) f^{-6} : random walk drift (RWDR)

Some noise types may be more prevalent in some clock types, and some may be negligible depending on the technology being used. The Allan variance is related to the phase spectral densities in the following way shown in Table 3.1, where C is a positive constant from integration. A more generalised conversion model can be found from (Barnes & Allan, 1966). The table shows the most relevant power-law terms for clocks.

Table 3.1: Relations of phase noise spectral density with Allan variance

Phase Noise Slope	Power Coefficient	Phase Noise $S_x(f)$	Allan Variance $\sigma^2(\tau)$
f^0	h_2	$\frac{h_2}{4\pi^2}$	$h_2 \frac{3f_h}{4\pi^2\tau^2}$
f^{-1}	h_1	$\frac{h_1}{4\pi^2 f}$	$h_1 \frac{3(C+\ln(2\pi f_H\tau))-\ln 2}{4\pi^2\tau^2}$
f^{-2}	h_0	$\frac{h_0}{4\pi^2 f^2}$	$h_0 \frac{1}{2\tau}$
f^{-3}	h_{-1}	$\frac{h_{-1}}{4\pi^2 f^3}$	$h_{-1} 2\ln(2)$
f^{-4}	h_{-2}	$\frac{h_{-2}}{4\pi^2 f^4}$	$h_{-2} \frac{2\pi^2\tau}{3}$

Comparing Clocks in GNSS

The stability of clock models is characterised by their specific technologies. High-end TCXOs and OCXOs tend to have the best frequency stability over 1 to 100 seconds, whereas CSACs perform better around 10000 seconds of averaging intervals. GPS Time and other GNSS systems provide better long-term stability than most portable solutions and, therefore, are typically used by clock users as reference oscillators to improve their performance through the means of *disciplining*.

Figure 3.7 shows a collection of typical AVAR performance for a range of clocks. TCXOs perform with the worst stability among the clocks. Their stability has been consistently

measured below 10^{-10} for all averaging times. OCXOs provide superior short-term stability from 1 to 100 seconds, but the stability deteriorates in the long term. CSACs, the newcomer in the clocks, do well, especially around $\tau = 10^4$ seconds, but suffer from not as impressive performance for shorter τ . The hydrogen maser clocks represent a high standard for reference. GPS clocks show the on-orbit performance of atomic clocks, which is more stable over the long term.




Figure 3.7: Survey of typical AVAR performance (Fernandez & Calero, 2015; Teunissen & Montenbruck, 2017)

3.3.3 Clock Coasting and Clock Modelling

Today, researchers face poor geometry, compromised signals, and loss of visible satellites in urban canyons for contemporary and novel applications. It is a new problem, and yet, attempts to solve the problem may be traced back to the earliest days of civil GPS use.

The GPS became available for civilian use in the 1980s. Until the 1990s, the number of GPS satellites in orbit was less than 20, compared to the norm of more than 30 active GPS satellites today. The surveyors were often forced to face the lack of visible satellites. Planning became necessary, especially in locations with less GPS signal coverage. A topic of research at the time was to improve the solution when the number of satellites was low or the geometry was poor. Past research conducted by Sturza (1983) proposed the idea of creating three-satellite solutions with a precise clock to replace the receiver clock error estimation.

The idea was that once the initial receiver clock offset (Knable & Kalafus, 1984) is known through estimation, with a precise clock as the timing source, the change in receiver clock offset can be low enough to reduce the need for further estimation and, therefore, reduce the number of variables by one. A three-satellite solution can be achieved since the 3-D spatial coordinates require at least three independent observations.

The method, later named *receiver clock coasting (RCC)*, is only feasible if a precise clock is used. As clock errors are multiplied by the speed of light, a few nanoseconds of errors will result in range errors in metres. Therefore, the Rb clock was deemed required for meaningful application and was also too expensive and immobile to carry around. The study nonetheless presented theoretically the positive impact of using a precise

external clock on geometry through a simple proof. A key conclusion was that a precise clock can be more helpful in improving solutions in a three-satellite scenario than an extra satellite.

The research topic delved into silence as the number of GPS satellites grew in the next two decades. There was little motivation as the scenario with three satellites seemed obsolete. The revival of the research started in the early 2010s thanks to the invention of CSACs, a natural companion to GNSS, and the proliferation of GNSS solutions in modern urban applications. The study from (Ramlall et al., 2011). In Ramlall's research, CSAC's clock bias was predicted with EKF and least squares, and a static SPP simulation using a CSAC as an aiding sensor resulted in a solution in around ten metres compared to over a hundred metres with a QXO clock. The method of predicting clock bias using the clock's properties was an improvement over the RCC method, and the method became known as the *receiver clock modelling (RCM)*.

With the CSACs becoming commercially available, clock-aided GNSS became investigated by more groups. Fernandez et al. (2016) studied the experimental aspects in adapting a CSAC with GNSS receivers as well as disciplining a CSAC (Fernandez & Calero, 2015; Fernandez et al., 2016, 2017). Interfacing CSAC with off-the-shelf GNSS receivers is, in fact, a practical challenge. The GNSS receivers utilise the external oscillator's signal to replace the internal clock. The process affects signal acquisition and tracking and is an example of a 'black-box' scenario discussed in sensor integration. The study will detail the interfacing method and caveats in the research.

The RCM method has been extensively studied by Krawinkel et al. (2018), as the group modelled a number of CSACs available on the market. Their research also depicted that it is possible to physically predict the clock behaviour with an EKF over short periods to

achieve a three-satellite solution and an improved solution in general. Kinematic SPP experiments were also involved to show improvement in the solution.

As the precursor of the current research, the RCC method in urban environments was studied using high-end hardware with the York-PPP engine as the first to combine a precise software solution with a precise external clock, and real-world tests were conducted with two separate receivers. One receiver was connected with a CSAC, while the other operated normally. With RCC in post-processing, the CSAC-aided GNSS receiver improved solution accuracy in kinematic positioning. The horizontal accuracy under significant obstructions simulated by an elevation mask of 55° was improved by 0.3 m and maintained the same level of position accuracy with an elevation mask of 30°. A three-satellite within metre-level was also achieved. The study proposed a number of topics to further the research, including extending the methodology to receiver clock modelling, adapting low-cost clocks and GNSS receivers and data processing in actual urban canyons. This dissertation addresses all of these areas of future work.

Receiver Clock Coasting

The RCC method is commonly used in conjunction with SLS processing. In a simplified n-satellite model estimating only the pseudoranges, the linearised navigation equation can be represented as

$$A\tilde{\mathbf{x}} = \tilde{\rho} \quad (3.35)$$

where $\tilde{\mathbf{x}}$ is the correction state vector consisting of three spatial variables and receiver clock error; $\tilde{\rho}$ is the pseudorange residual vector; A is the design matrix.

The equation can be expanded into its full form:

$$\begin{bmatrix} \alpha_{1x} & \alpha_{1y} & \alpha_{1z} & 1 \\ \alpha_{2x} & \alpha_{2y} & \alpha_{2z} & 1 \\ \alpha_{3x} & \alpha_{3y} & \alpha_{3z} & 1 \\ \dots & \dots & \dots & \dots \\ \alpha_{nx} & \alpha_{ny} & \alpha_{nz} & 1 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta t_r \end{bmatrix} = \begin{bmatrix} \Delta \rho_1 \\ \Delta \rho_2 \\ \Delta \rho_3 \\ \dots \\ \Delta \rho_n \end{bmatrix} \quad (3.36)$$

If RCC is activated, the receiver clock error is removed from the equation, and Equation 3.36 becomes:

$$\begin{bmatrix} \alpha_{1x} & \alpha_{1y} & \alpha_{1z} & 0 \\ \alpha_{2x} & \alpha_{2y} & \alpha_{2z} & 0 \\ \alpha_{3x} & \alpha_{3y} & \alpha_{3z} & 0 \\ \dots & \dots & \dots & \dots \\ \alpha_{nx} & \alpha_{ny} & \alpha_{nz} & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta t_r \end{bmatrix} = \begin{bmatrix} \Delta \rho_1 \\ \Delta \rho_2 \\ \Delta \rho_3 \\ \dots \\ \Delta \rho_n \end{bmatrix} + \Delta t_{r0} \mathbf{1}_n \quad (3.37)$$

where $\mathbf{1}_n$ is the $n \times 1$ identity vector and Δt_{r0} is the coasted receiver clock error. Texts in RCC often truncate the matrix dimensions by 1, rather than setting a zero column in the design matrix. RCC is implemented as shown in equation 3.37.

The merit of RCC lies in the improvement in the dilution of precisions (DOPs), particularly the geometric dilution of precision (GDOP), that results from eliminating the temporal variable. In the regular case, if the pseudorange residuals are uncorrelated among measurements, the pseudorange covariance matrix can be reduced to (Sturza, 1983)

$$\text{cov}(\tilde{\rho}) = \sigma_{\rho}^2 I \quad (3.38)$$

The covariance matrix for the spatial and temporal states has the form of

$$cov(\tilde{\mathbf{x}}) = E[\epsilon_x \epsilon_x^T] = A^{-1} cov(\tilde{\rho}) A^T = [A^T cov(\tilde{\rho}) A]^{-1} = \sigma_\rho^2 (A^T A)^{-1} = \begin{bmatrix} \sigma_x & 0 & 0 & 0 \\ 0 & \sigma_y & 0 & 0 \\ 0 & 0 & \sigma_z & 0 \\ 0 & 0 & 0 & \sigma_t \end{bmatrix} \quad (3.39)$$

The zero terms are the result of non-correlation among measurements. Since there exists a transformation from the ECEF frame to the local navigation frame, the matrix can be rewritten with North, East, Up coordinates, that is,

$$cov(\tilde{\mathbf{x}}) = \begin{bmatrix} \sigma_e & 0 & 0 & 0 \\ 0 & \sigma_n & 0 & 0 \\ 0 & 0 & \sigma_u & 0 \\ 0 & 0 & 0 & \sigma_t \end{bmatrix} \quad (3.40)$$

The GDOP is defined by:

$$GDOP = \sqrt{Tr((A^T A)^{-1})} = \sqrt{\sigma_e^2 + \sigma_n^2 + \sigma_u^2 + \sigma_t^2} / \sigma_\rho \quad (3.41)$$

where Tr indicates the trace of a matrix.

The GDOP may also be represented by PDOP (Positioning Dilution of Precision) and TDOP (Timing Dilution of Precision), where

$$PDOP = (\sqrt{\sigma_e^2 + \sigma_n^2 + \sigma_u^2}) / \sigma_\rho \quad (3.42)$$

and

$$TDOP = \sqrt{\sigma_t^2 / \sigma_\rho} \quad (3.43)$$

With RCC, the equation 3.37 can be solved with

$$\tilde{\mathbf{x}} = A^{-1}(\tilde{\rho} - \Delta t_0 \mathbf{1}_n) \quad (3.44)$$

Again, solving for the covariance of the state vector corrections gives

$$cov(\tilde{\mathbf{x}}) = A^{-1}[\sigma_\rho^2 I + \sigma_{t_0} I] A^{-T} \quad (3.45)$$

This gives the three-state derivation of PDOP, that is,

$$PDOP_3 = \sqrt{Tr[A^T A]^{-1} + \frac{\sigma_{t_0}^2}{\sigma_\rho^2} A^{-1} A^{-T}} \quad (3.46)$$

In some texts, the RCC method is sometimes considered to provide better 'geometry' with three satellites and clock coasting than a regular, four-satellite solution. This derivation of $GDOP_3$ shows that such a scenario is indeed possible.

Receiver Clock Modelling

The RCM method adopts a clock model to describe the clock's behaviour over multiple epochs. The idea of clock modelling in a Kalman filter implementation is first seen in (Brown & Hwang, 1997; Misra, 1996). A common approach is to consider a two-state model consisting of x_p , the clock bias, and x_f , the frequency offset. The dynamics of the

two-state model is modelled as a random ramp and a random walk:

$$\begin{bmatrix} \dot{x}_p \\ \dot{x}_f \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_p \\ x_f \end{bmatrix} + \begin{bmatrix} w_p \\ w_f \end{bmatrix} = A_t \mathbf{x}_t + \mathbf{w}_t \quad (3.47)$$

with w_i being process noises. As the system is time-invariant, the state transition matrix Φ over time interval Δt is given by

$$\Phi = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \quad (3.48)$$

The covariance of the process noise is given by:

$$Q_{clk} = E([w_p w_f]^T [w_p w_f]) = \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix} \quad (3.49)$$

With RCM, the process noise matrix Q_{clk} shall be represented with characteristics of the clock hardware. To derive the specific process noise, the relation

$$\mathbf{w}_{t,k} = \int_{t_{k-1}}^{t_k} \Phi(t, \tau) \mathbf{w}_t(\tau) d\tau \quad (3.50)$$

is explored.

From 3.50, the process noise matrix $Q_{clk,k}$ at time k can be expanded as:

$$Q_{clk,k} = \int_{t_{k-1}}^{t_k} \Phi(\tau) Q(t) \Phi(\tau)^T d\tau \quad (3.51)$$

At this step, the derivation invokes the spectral densities of clock phase and frequency.

$$Q(t) = \begin{bmatrix} S_x & 0 \\ 0 & S_y \end{bmatrix} \quad (3.52)$$

where S_x and S_y are the phase and frequency spectral density defined in Chapter 3.3.2.

This leads to the expanded form of three terms:

$$\begin{aligned} E[x_p^2(\Delta t)] &= \int_0^{\Delta t} \int_0^{\Delta t} 1 \times 1 \times S_x \delta(\tau_1 - \tau_2) d\tau_1 d\tau_2 + \int_0^{\Delta t} \int_0^{\Delta t} \tau_1 \times \tau_2 \times S_y \delta(\tau_1 - \tau_2) d\tau_1 d\tau_2 \\ &= S_x \Delta t + \frac{S_y \Delta t^3}{3} \\ E[x_f(\Delta t)x_p(\Delta t)] &= \int_0^{\Delta t} \int_0^{\Delta t} \tau_2 \times 1 \times S_y \delta(\tau_1 - \tau_2) d\tau_1 d\tau_2 \\ &= \frac{S_y \Delta t^2}{2} \\ E[x_f^2(\Delta t)] &= \int_0^{\Delta t} \int_0^{\Delta t} 1 \times 1 \times S_y \delta(\tau_1 - \tau_2) d\tau_1 d\tau_2 \\ &= S_y \Delta t \end{aligned} \quad (3.53)$$

Then, Q_{clk} can now be expanded into the following form:

$$Q_{clk} = \begin{bmatrix} S_x \Delta t + S_y \frac{\Delta t^3}{3} & S_y \frac{\Delta t^2}{2} \\ S_y \frac{\Delta t^2}{2} & S_y \Delta t \end{bmatrix} \quad (3.54)$$

The terms S_x and S_y relate to the clock's stability performance, such as depicted in Table 3.1. A few models have been studied in the past in approximating the S_x and S_y using clock stability parameters h_i . The first model is given by (Dierendonck & McGraw, 1984) as a further development of the original work from (Barnes & Allan, 1966) using the

impulse response of different processes.

$$Q_{clk} = \begin{bmatrix} \frac{h_0}{2} + 2h_{-1}\Delta t^2 + \frac{2}{3}\pi^2 h_{-2}\Delta t^3 & 2h_{-1}\Delta t + \pi^2 h_{-2}\Delta t^2 \\ 2h_{-1}\Delta t + \pi^2 h_{-2}\Delta t^2 & \frac{h_0}{2\Delta t} + 2h_{-1} + \frac{8}{3}\pi^2 h_{-2}\Delta t \end{bmatrix} \quad (3.55)$$

The second model is given by (Brown & Hwang, 1997). The model ignores the flat segment, corresponding to the flicker frequency noise (FLFR). The omission is because the term gives rise to a variance expression of the order of Δt^2 , which is impossible to model exactly with a finite-order state model. A very simple expression can then be derived:

$$\begin{aligned} S_x &\approx \frac{h_0}{2} \\ S_y &\approx 2\pi^2 h_{-2} \end{aligned} \quad (3.56)$$

The dynamic model models the white phase noise and the flicker phase noise, leaving the lower-order power-law noises as part of the process noise modelling. Power coefficient below h_{-2} is typically ignored as their contribution is negligible in Equation 3.55.

Discussion

The two approaches differ in whether a clock model is designed to describe the clock's behaviour. A misconception is that both methods ignore receiver clock error Δt_r , while the RCC method requires estimation of Δt_r unless invoked, and the RCM method constantly estimates the Δt_r with a model better fit to the specific clock. To help distinguish the two methods, Figure 3.8 shows an illustration of how the two methods differ when the user enters an urban canyon and activates RCC at t_0 , then leaves the urban canyon at t_1 . The black solid curve shows an ideal estimation of receiver clock error

throughout the period, regardless of any obstructions. The blue curve representing RCC shows four segments: (1) before the urban canyon, an RCC solution remains identical to traditional processing; (2) after activation, the RCC method assumes a constant receiver clock error from the point of activation (t_0); (3) as the user leaves the urban canyon, the clock error estimation converges back to regular processing relatively more quickly than RCM; (4) the RCC solution converges back to be identical with traditional solutions. The red curve representing an EKF-based RCM starts out as a model of the black curve, and its solution is kept close to the ideal curve until the user reaches the urban canyon. As of t_0 , the EKF clock model continues to utilise the prior measurement to predict clock behaviour through the urban canyon. After t_1 , with better incoming measurements, the model adjusts itself to converge to the open-sky solution. The convergence finally completes at t_2 .




Figure 3.8: Theoretical behaviour of receiver clock coasting (blue) and EKF-based receiver clock modelling (red) in an urban canyon compared to the open-sky estimation (black)

The RCC method is an on-or-off approach where an algorithm is placed to determine if the RCC processing should be activated. The solution prior to activation remains identical to traditional processing. After activation, the prior receiver clock states are 'frozen', and the RCC model kicks in to assume the receiver clock error either remains unchanged or follows a predetermined pattern.

In reality, the RCM curve may either be closer to the open-sky curve or further away than the RCC curve, meaning that the receiver clock estimation may or may not be more accurate than the RCC method. The quality of the receiver clock model determines the real-world performance. Selecting clock hardware that is more easily modelled with less in-run variation is important.

3.3.4 Triple-Sensor Fusion

Fusing the three sensors of a GNSS receiver, an IMU, and an external clock is a natural evolution of the GNSS/clock integration scheme. However, it is also natural to question the benefit of adding a precise external clock to the traditional GNSS/IMU fusion. The external clock cannot improve the observables for IMU directly as it does with the GNSS receiver: the time precision provided by the hardware can be crucial in radio ranging, where the signal travels at the speed of light, but is negligible in vehicle dynamics. The GNSS receiver takes the timing signals from the external clock to assist in both the signal tracking and observation process. The same scheme makes little sense with an IMU.

The existing research in sensor fusion typically only employs a couple of the two sensors. Ma (2016) used all three sensors to produce a triple fusion solution with a CSAC as the external clock. The data was post-processed with an EKF, and software-induced masks were in place to stimulate harsh urban environments. The triple-sensor fusion presented the most accurate solution with a mean horizontal error of less than 10 metres. As a comparison, the GNSS/clock fusion presented a similar mean horizontal error, about 2 metres higher. The traditional GNSS/IMU solution was drifting too much, resulting in errors above 100 metres.

In this research, the IMU is considered an integral addition to the GNSS/Clock fusion scheme. Either one of the RCC or RCM algorithms is capable of improving the navigation solution with a low number of visible satellites, as well as providing a meaningful three-satellite solution without rank deficiency. Once the user enters a more challenging environment, further reducing visible satellites to the point where no visible satellites are available, an IMU would be an ideal dead reckoning sensor to keep the solution going. Furthermore, from a traditional GNSS/IMU integration perspective, improving

the solution obtained from the external clock will benefit the estimation of the IMU states. Therefore, the idea that the external clock provides no benefit to the GNSS/IMU solution can be rebuked. As for a quantitative estimation of the improvement, or whether there are specific cases where the clock will benefit existing, traditional GNSS/IMU solutions will be discussed and analysed in the following chapters, where data-collection campaigns took place.

3.4 Hardware Assembly and Integration

The hardware solution provides a framework for realising and testing RCC and RCM in real-world situations. First, a framework using mostly higher-cost, geodetic-grade sensors was built. This solution was built upon a CSAC-aided GNSS receiver, which has been traditionally used in the field of research. The research then transforms the solution into a low-cost configuration absent of any geodetic-grade sensors. The low-cost triple-sensor solution was built upon an OCXO. Either of the solutions consists of three core sensors: the GNSS receiver, the external oscillator and the inertial measurement unit.

The hardware assembly must also ensure that the timing signal from the external clock provides a live oscillation reference for the satellite acquisition and tracking procedures, as well as the data recording for the post-processing. For that, a 10 MHz timing signal is fed from the clock to the GNSS receiver.

Figure 3.9 shows an overview of the hardware flow of a typical triple-sensor integration during a field testing run. GNSS signals received by the GNSS receiver are not only used to generate GNSS observations but are also used to support the timing for other sensors. For example, a low-cost, consumer-grade IMU will require the GPS Time to produce

timestamps that will eventually be used to synchronise with GNSS observations during processing. The diagram shows only the operational flow of a triple-sensor integration. The external clock requires discipline from the GPS time to maintain its long-term performance. This process is omitted from the overview diagram and will be specified in detail in the coming section. The York-PPP adopts a time synchronisation algorithm, where the first 100 GNSS epochs are processed independently without IMU information, followed by matching the closest IMU epochs with the least timestamp differences. The mechanisation loop runs all IMU readings before the subsequent GNSS epochs arrive.




Figure 3.9: Typical triple-sensor hardware flow during a field test

The way that the GNSS receiver utilises the external frequency reference is very much open to the manufacturer's interpretation and consideration. In higher-end timing receivers, where a high-performance VCXO is often employed, an internal disciplining algorithm will be used to drive the internal frequency towards the external frequency. In

low-cost GNSS receivers, rather than disciplining the internal clock, it can be found that the external reference completely replaces the internal frequency (Lombardi, 2016).

3.4.1 Clock Discipline

Timing oscillators suffer from long-term deviation of frequency from effects such as ageing. To mitigate the issue, a process called discipline is used. Clock disciplining synchronises the internal frequency of the clock with an external reference frequency (Lombardi, 2008). The reference signals are assumed to be more stable. Through continuous tracking of the external frequency over specified intervals, the clock steers its frequency output towards the external frequency. The steering process commonly happens within the software, where an adjustment value is computed, memorised, and applied to the frequency output signals. The reference frequency source is assumed to be superior, at least in the time interval for the discipline iterations.

Clock discipline is used when the application in question requires relatively high-frequency stability from the clock, such as in the triple-sensor fusion in the research. High-performance oscillators are intuitively preferred for such an application, so it is unlikely that the user will have a better oscillator at hand. A convenient and effective way to discipline the clock in the triple-sensor integration is to discipline the external clock with the GPS Time.

Disciplining to a superior frequency source is crucial, as the reference time signal should be stable enough for the steering algorithm to converge to. Disciplining to an inferior frequency source lets the algorithm try to converge to a moving target and is thus unlikely to compute an adjustment value that compensates for the long-term drifting of the clock.

Figure 3.10 shows the typical discipline process of the external clock using GPS Time through the GNSS receiver within the system. Timing GNSS receivers commonly work both ways: they can utilise an external, more stable frequency source for their own PNT, or they can output their timing signals, synchronised to the GPS time in the long term as 10 MHz timing signals. The 10 MHz timing signals can be fed to the external clock, where the internal discipline process will take over to steer its frequency towards the signals from the GNSS receiver. The internal algorithms store the correction within the memory of the clock and will apply the correction next time the clock is booted up. The internal oscillator is typically a physical process that naturally generates oscillations when activated, and cannot be directly controlled. Rather, the correction is most commonly applied to the internal algorithms that interpret the detection of the oscillation and its conversion to an output frequency.




Figure 3.10: Clock discipline process using GPS Time

The process presents a few experimental or practical caveats.

- The time interval for each discipline iteration needs to be 'long'.

A range of recommended time intervals for a discipline by GNSS is suggested per the user manuals of clock models (Microsemi, 2017a). A longer time interval extends the time needed for each discipline. A shorter time interval reduces the time needed, but increases the risk for the algorithm to fail to steer the clock frequency towards converging with the reference frequency.

The best time intervals for a discipline by GNSS are ultimately determined by the time interval τ where the reference frequency presents superior stability. In shorter intervals, such as 10^0 to 10^2 seconds, the time reference from the GNSS receiver is unlikely to be more stable than the oscillator to be disciplined. Typically, a range of a few thousand seconds is recommended.

Since both the RCC and RCM require a more predictable clock behaviour, this research practically suggests choosing a more extended time interval from the recommended range to prioritise positioning performance.

- During a discipline, it is crucial to maintain a stable timing signal from the GNSS receiver. GNSS receiver produces a frequency replica more stable than the external clock in the long term if the PNT solution converges well and stays consistently converged. The stability is particularly more difficult with a low-cost GNSS receiver setting. Furthermore, the GNSS antenna should ideally be stationary, in an open-sky environment with the absence of any artificial or natural disturbance.

A disturbance in the timing signal provided by the GNSS receiver can cause the steering algorithm to 'oversteer', and take longer to effectively converge to a more stable frequency. Frequent disturbances have the potential to continuously mis-

direct the algorithm so that the discipline fails to converge to a stable adjustment value.

- Environmental factors play a critical role in the clock's stability. Disciplining helps to compensate for the long-term frequency drift of the clock. However, oftentimes, short-term factors such as environmental temperature, vibration and shock have a more prominent effect on the frequency stability (Microsemi, 2014, 2017b, 2018). The environmental impact is particularly crucial with temperature-sensitive clock technologies.

As a result, keeping the clock in a lab environment prevents environmental factors from affecting the discipline process. It is also worth noting that in an operational environment, such factors may cause inaccuracies with the RCC or RCM algorithm.

These factors contribute to potential errors for the RCC and RCM algorithms. In the worst-case scenario, the contribution from these factors could lead to a worse PNT solution than traditional processing. This research, therefore, prepared the hardware, including the clock discipline, in a lab environment whilst the field tests were conducted in real-world situations.

The specific technologies and design of clock models differentiate their performance, ageing, and environmental vulnerability. In the following sections, the two types of clocks used in the research will be briefly introduced.

Chapter 4

Building the Prototype Triple-Sensor Fusion Solution

The hands-on portion of the research begins with building a prototype triple-sensor solution. The goal of the prototype solution is to verify the feasibility of triple-sensor fusion using current hardware, to examine and verify the performance of the triple-sensor solution in various real-world datasets.

Three sensors are chosen: the NovAtel OEM7 GNSS receiver, the Microchip SA.45s CSAC, and an XSENS MTi-7 IMU. The sensors are relatively high-end, except for the IMU. The NovAtel OEM7 receiver is a geodetic-grade receiver. The sensors are chosen based on their performance and customisability. Lower-end, consumer-grade hardware often provide fewer customisation options for users, and for research purposes, it is always preferable that the researcher has more control over the operation of the hardware.

The sensors, in particular the GNSS receiver and the external clock, are intentionally selected to be high-performance. As a prototype, the goal of this part of the research is

verification. In the next part of the research, lower-cost sensors will replace the geodetic receiver and the CSAC to build a lower-cost solution that allows the solution to be applied to mass markets.

The section will begin with an introduction to the essential technology of Chip-Scale Atomic Clocks, for which the clock-aided GNSS receiver concept is first derived. The introduction will begin

4.1 Chip-Scale Atomic Clock

Chip-Scale Atomic Clocks (CSACs) have been increasingly seen in various demanding timing solutions, such as resilient timing and relatively low-cost space timing for CubeSats.

The technology was first introduced to the public by the U.S. National Institute of Standards and Technology (NIST) in 2004 (Brannon et al., 2006). Commercial products from various manufacturers emerged shortly after, with serialised products getting mass-manufactured. CSACs provide a superior level of frequency stability over their peers of similar sizes and power consumption. The Microsemi SA.45s CSAC used for this research has a physical dimension of $4.06 \times 3.56 \times 1.14$ centimetres, comparable to a dollar coin. They also have a low power consumption of less than 125 mW. The combination of small size and low wattage has made the CSACs an ideal option for portable applications powered with consumer-grade batteries. The CSACs have seen increasing uses in diverse applications from the timing source of Low-Earth Orbit navigation projects, CubeSats of various uses, Ocean Bottom Seismographs (OBS), and now, the triple-sensor prototype investigated in this research.

The research adopts the Microsemi SA.45s as the CSAC of choice to build the prototype tri-sensor fusion with a geodetic GNSS receiver. The reason for this choice is that SA.45s provides a large pool of console commands for developers to study and experiment with, and a development board which supports all the necessary functions to build the triple-sensor receiver prototype can be acquired. The rest of this section is dedicated to briefly introducing the relevant aspects of the CSAC model, the experimental procedures, and the caveats during the assembly, integration and testing of the prototype triple sensor fusion.

4.1.1 Operational Principles

The CSAC is an oscillator system built around its Physics Package (PP), where the critical process of Coherent Population Trapping (CPT) takes place, and it consists of a number of supporting components installed alongside the PP to power, tune, and convert the physical phenomena in the PP to timing signals (Lutwak et al., 2002b, 2004; Microsemi, 2016).

4.1.2 Physical Principles and Environmental Factors

The advent of the CPT technology is the critical innovation that enabled the CSAC technology. The motivation of the technology was the synchronisation of time-ordered encryption keys and accelerated GPS acquisition for P(Y) signals in the absence of civilian signals in a warfare situation. The military application restricted the power requirement for the system, which in turn benefits the civilian end-product for consumer use. With a pair of finely tuned laser beams, a Caesium gas cell is simultaneously interrogated by the two laser beams tuned to excite the Cs atoms from a pair of ground hyperfine

states to the same excitation state. The invention of low-wattage Vertical-Cavity Surface Emitting Lasers (VCSELs) soon before the CSACs allowed two such coherent frequencies to be generated from a single laser source. The excited Cs atoms in the resonance cell return to the ground states by releasing photons that can be detected by the detectors built around the VCSEL source, thus creating an oscillation that can be detected.

The design allows the minimisation of the oscillator based on existing technology and a reasonable budget constraint. However, there are a few major engineering concerns that eventually lead to environmental and ageing factors that lead to short-term or long-term stability changes to the CSAC.

To support a sufficient density of gaseous Cs atoms, the Cs atoms in the cell must be heated to 75 °C. Therefore, it is crucial to delegate enough power budget to the heating mechanism while preventing possible heat loss to the ambient through any means. The CSAC model achieves the goal through enclosing the PP within a vacuum, eliminating convection or conduction through air. While the conduction through the inner circuit or radiation cannot be prevented, it can only be optimised through design. This leads to a requirement for the operational temperature to be in the range of -10 °C to +70 °C. The temperature control mechanism has a limited ability to heat up the gas chamber and is unable to cool down the chamber if the ambient temperature is too high. However, there are variant models that allow a wider range of operational temperatures. As such, the system should avoid rapid ambient temperature changes as the heating mechanism will need to cope with the change, and the actual temperature within the gas cell may deviate from the ideal temperature, leading to increased timing errors.

The design also contributes to an ageing factor where the vacuum sealing may deteriorate over time and limit the CSAC's ability to stabilise. Similarly, material degradation

over time will have an effect on the physical operation of the PP and will eventually cause timing errors. There are a number of possibilities when it comes to long-term stability and ageing. Some of the possibilities include changes to the gas mixture, dropped Cs content through diffusion, interaction of Cs with residual impurities over time, etc.

Vibration and mechanical shocks are also factors of short-term environmental errors. Vibrations may disrupt the optical geometry from which the resonance is achieved, and any shocks that exceed what the PP's suspension layer allows may cause heat loss that destabilises the clock. The physical principle of a CSAC relies on the optical interaction within the PP, and instabilities incurred by the optical interactions by any means, including instabilities to the exciting laser source, may be a factor in frequency stability.

The CSAC functions by translating the PP's physical phenomena into timing signals through a loop of supporting subsystems. These systems include a TCXO, a loop filter, a microwave synthesiser, and a microprocessor which controls the operation at the system level. The microprocessor oversees the CSAC operation, monitors the health and status of the CSAC, receives user commands and returns with the information polled by the user's command through an RS-232 cable.

A feedback loop controlled by the microprocessors is used to provide a stable timing signal. The TCXO serves as the basic timing unit but receives constant correction from the loop filter that receives an oscillating reference from the atomic resonance of the PP. This process allows the frequency stability of the original TCXO to be improved by several orders of magnitude.

4.1.3 Clock Discipline

The discipline of the CSAC was done using the NovAtel OEM7 GNSS receiver using a geodetic antenna installed on the rooftop of the Lassonde Research Centre (LRC) building. During the discipline process, the OEM7 receiver is placed indoors and connected to the rooftop antenna in an open-sky environment. The CSAC was connected to the OEM7 receiver through the 1 pps signal from the OEM7 receiver to the 1 pps disciplining interface onboard the CSAC. The GNSS receiver and the CSAC are connected through serial ports to a controlling laptop. The OEM7 receiver is controlled through the default NovAtel Application Suite. The CSAC is controlled through a combination of the Microsemi CSAC GUI software and serial command console software. The latter is used in conjunction with the official GUI (Graphical User Interface) software to allow user input of custom commands.

The CSAC was disciplined through a telemetry command '!D6000', where the parameter '1PPS Disciplining Time Constant' is set to 6000 seconds. Once activated, the internal algorithm will first begin a 1 pps synchronisation with the reference 1 pps signal from the GNSS receiver. During the synchronisation mode, an internal phase meter begins to measure the relative phase difference between the CSAC's 1 pps signal and the GNSS receiver's. The phase meter has a resolution of 450 ps. The time constant is set according to the user manual's suggestion of a 5000-second disciplining interval with GNSS receivers. The manufacturer's recommendation is based on empirical timing stability from a GPS receiver, where it becomes more stable than the CSAC at an averaging time higher than 5000 seconds. If disciplined to a superior atomic clock, the CSAC's time constant for disciplining can be set at a much lower value at which the superior clock demonstrates a more stable frequency. In this research, the CSAC was disciplined with a

time constant of 6000 seconds,

Figure 4.1 shows a sample clock disciplining process over five days. In the disciplining process, the phase difference between each 1 pps tick from the clock to the external reference of an OEM7 GNSS receiver was initially large, prompting a high steering value by the algorithm in response. Gradually, the phase difference converges to the threshold of 50 ns set by the user at around 10 hours. Once the convergence was reached, the clock maintained a relatively stable 'STEER' value, the internal measurement in parts per trillion that steers the output frequency. It can be observed that there were two instances where the phase differences exceeded the threshold after the first convergence. Such phenomena are sometimes observed. In both instances, the steering value remained around -200, suggesting that the initial convergence was accurate, and some unknown factors originating from either the CSAC or the reference receiver came into effect.




Figure 4.1: CSAC phase difference from reference and steering value during discipline

The convergence threshold is typically set around 20 ns. A threshold too low will cause the discipline algorithm to constantly search for a 'STEER' value that is beyond the clock's inherent stability. In such a case, the user cannot reach a converging steering value. A threshold too high will lead the algorithm to stabilise too soon, reducing the precision of the final steering value.

Operationally, the final steering value can be stored in the clock's memory using the 'LATCH' function. Once the steering value is stored, the clock will automatically apply the steering value. The process is irreversible, so it is crucial to latch only a converged steering value using a superior external oscillator.

4.2 Road Tests Set-up and Methodology

In this section, the prototype triple-sensor fusion is realised by assembling a disciplined Microchip SA.45s CSAC, a NovAtel OEM729 GNSS receiver, and an XSens MTi-7 IMU. Among the three sensors, the GNSS receiver is geodetic-grade, and the IMU is MEMS-based and low-cost. The setup was taken to public streets in Ontario, close to York University, Toronto, for data collection. All data was collected in real-time and processed later through post-processing. The post-processing was done in simulated real-time by feeding observations to the engine epoch by epoch, and no smoothing or data manipulation was involved. The GNSS receiver's timestamps are in sync with the external frequency, if available. The IMU timestamps are also synchronised with the GPS time through its built-in GNSS receiver.

4.2.1 Experimental Set-up

For real-world road tests, the three sensors were assembled together with a series of additional equipment. The setup consisted of the rover part, which was an integrated box to be attached above the moving vehicle, and the stationary part located in the lab. Figure 4.2 shows an illustration of the experimental setup for the road tests, where GNSS observation data and IMU data were collected. The antenna for the rover was placed on the roof of the vehicle, and the base station antenna was stationary as it was in the disciplining tests. The NovAtel SPAN receivers were chosen to produce the reference solution for both the rover and the base station. It is worth noting that the base station's position has been verified to centimetre-level through separate stationary positioning across days using PPP. The IMU was a low-cost sensor using the patch antenna provided with the package. A NovAtel OEM729 timing receiver aided by the Microchip SA.45s CSAC is connected to the same antenna as the reference receiver.




Figure 4.2: Road test setup with the CSAC-aided triple-sensor fusion

The rover section of the setup was integrated into a box where the antennas were placed on the roof of the box. The box hosted the entirety of the roving sensors, tightly strapped over the car roof, with the necessary cables connected to the laptop within the vehicle. Figure 4.2.1 shows photos of the experimental setup. The antennas were fixed onto the top surface, eliminating the possible signal obstructions from the vehicle itself, but were still susceptible to environmental obstructions, as is required by the real-world conditions.



Figure 4.3: Roving sensors integrated with the vehicle rooftop

4.2.2 Road Test Conditions

The road tests were taken in dynamic real-world driving situations. The vehicle was driven within around six kilometres of the base station. The trips consisted mainly of suburban buildings as obstructions. However, recent developments in the region have brought a few high-rises along the roads. Figure 4.2.2 shows the environment in which road tests were conducted. The area around York University is a busy one, and arterial roads such as Highway 7 are often travelled by commuters and cargo trucks alike. The occasionally heavy traffic on the roads has added uncertainties to the environmental obstructions. Oftentimes, the vehicle passes trucks on one or both sides, and the vehicle may be stopped by a traffic light along with a trailer truck obstructing GNSS reception. Each data collection trip is designed to last 10 - 30 minutes in duration, and the routes are designed to be in close-ended loops.



Figure 4.4: Google Street views of the routes: Highway 7 (above) and Pond Road (below)

The reference solution is provided by the RTK solution computed by NovAtel's Inertial Explorer software. The software is commercially licensed and can fully exploit the two SPAN receivers' potential, so that the solution produced becomes a trustworthy reference. Smoothing has been enabled for the processing of the reference solution.

4.3 Results and Analysis

In this section, the datasets collected through the road trips with the CSAC-aided GNSS/IMU triple fusion are processed and analysed. First, the kinematic PPP solution is produced as a baseline solution for the analysis to come. The purpose of the baseline solution is to provide a reference for the vanilla, raw performance of the PPP-only solution in the tested environment. Afterwards, the datasets are studied with tightly-

coupled integration enabled, and then with the triple-sensor solution enabled with both RCM and RCC.

4.3.1 Baseline PPP Solution

The PPP solution was produced using the York-PPP GNSS engine. GPS, GLONASS, Galileo and BeiDou were all used for processing. Specifically, the engine processes GPS L1, L2 and L5; GLONASS G1 and G2; Galileo E1, E5a/E5b and E6; BeiDou2 B1-2, B2b and B3; BeiDou3 B1 and B2a/B2b. The processing was completed with the rapid precise orbit and clock products from GFZ (German research centre for Geosciences, GeoForschungsZentrum) (Uhlemann et al., 2015). Additionally, the differential code bias (DCB) products from the Chinese Academy of Sciences (CAS) are used (Liu et al., 2019). IGS14 ANTEX file was used for antenna correction (Schmid et al., 2016). The observations are in RINEX format with a data frequency of 5 Hz (Gurtner & Estey, 2007). GMF model and mapping function (Kouba, 2009) was used for dry tropospheric delays, and the wet delays were modelled as a random walk. Ionospheric delays were estimated as white noise. Finally, an elevation angle cut-off of 10° was used.

Figure 4.5 shows the navigation solution of the PPP baseline solution. A time series of the horizontal errors is shown over roughly 40 minutes. The plot is divided into two disjoint runs with a short break in the middle. Following the time series, Figure 4.6 shows the cumulative distribution function of all the positioning solutions drawn in blue. The horizontal error is then divided into the along-track and lateral components using the velocity vectors of the reference solution. A lightly drawn, bright-yellow dashed line shows the 1.5 m reference for lane-level positioning.




Figure 4.5: Time series of horizontal errors in the GNSS-only, PPP solution - Run 1 + 2




Figure 4.6: Cumulative distribution of horizontal errors in the GNSS-only, PPP solution - Run 1 (left) and Run 2 (right)

Table 4.1 shows the rms errors, the 95th percentile rms, and the 67th percentile rms of the horizontal errors. The latter two rms are computed using only the 2σ and 1σ errors.

They reflect the positioning accuracy in the absence of 'error peaks' typically caused by environmental obstructions. Comparing horizontally within the same solution, the different columns show the impact of obstructing environments on the solution. Comparing vertically across solutions, the last two columns show the change in rms errors in the absence of extreme horizontal errors.

Table 4.1: Horizontal rms statistics of the PPP baseline solution in Run 1

Solution	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP Run 1	7.36	0.75	0.68
PPP Run 2	4.32	0.35	0.33

4.3.2 Results with Tightly-Coupled Integration and Clock Aiding

The GNSS-only solution in the first dataset showed a large number of instances of environmental obstructions, particularly due to the complex traffic situation on the particular day and time. The performance is expected to improve with the TC integration, with IMU data involved in the processing.

Suburban Datasets - Run 1

The triple-sensor fusion with RCM is tested. Figure 4.7 shows the CDF and time series of the three solutions. Table 4.2 shows the statistical data of the three sets of solutions. Both the tightly-coupled, dual-sensor solution and the triple-sensor solution managed to eliminate the error peaks found in the GNSS-only PPP solution. The clock-aided solution using RCM showed a slightly improved accuracy than that of the dual-sensor solution using only the IMU. The RCC was not tested as the dataset did not show any continuous

periods of obstructions, as one would expect in an urban canyon. The RCC will be tested in the upcoming section, where simulation outages are implemented. Note that the PPP/IMU fusion is conducted in a tightly-coupled fashion, and to avoid conflicts with the loosely-coupled fusion, which is also present in the York-PPP engine, the solutions are labelled as "PPP/TC".




Figure 4.7: Results from the dual- and triple-sensor integration: cumulative distribution of horizontal errors (top) and time series of horizontal errors (down)

Table 4.2: Horizontal rms statistics of PPP-only, PPP/TC and PPP/TC/RCM solutions in Run 1

Solution	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP	7.36	0.75	0.68
PPP/TC	0.75	0.71	0.68
PPP/TC/RCM	0.74	0.70	0.67

The improvement of the solution can be observed more intuitively with the trajectory plotted. As shown in Figure 4.8, the vehicle experienced a brief period of signal obstruc-

tions with visible satellites dropping below five and fluctuating for several epochs. The PPP solution experienced clear 'jumps' away from the reference solution in bright green. However, the solutions with either sensor fusion managed to avoid the position jumps.




Figure 4.8: A snapshot of trajectory of all processing methods in Run 1

Suburban Datasets - Run 2

The navigation solution is analysed and plotted in a similar manner in Figure 4.9 and in Table 4.3. It can be observed that the identical pattern found in Run 1 is again found in Run 2. The dual-sensor fusion with a tightly-coupled IMU removes some of the error peaks in the PPP-only processing, and the use of the CSAC with RCM algorithm has further improved the solution in these circumstances, leading to a further 24% reduction of horizontal errors.

Table 4.3: Horizontal rms statistics of PPP-only, PPP/TC, and PPP/TC/RCM solutions in Run 2

Solution	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP	4.31	0.35	0.33
PPP/TC	1.49	0.35	0.31
PPP/TC/RCM	1.14	0.31	0.29




Figure 4.9: Results from the dual- and triple-sensor integration: cumulative distribution of horizontal errors (top) and time series of horizontal errors (down)

Noticeably, the second dataset involves a few brief outages caused not by nearby buildings, but by overpasses. The overpasses have caused several occasions where the visible satellites drop dramatically to less than three. This explains the cases where the algorithm failed to eliminate error peaks. Clock-aiding algorithms are expected to work with three satellites in view, but when the number of satellites drops below three, the processing engine will rely on the IMU states. Figure 4.10 shows a time series of the number of visible satellites during Run 2 and the GDOPs computed accordingly. The plot, in conjunction with Figure 4.9, shows that the extended outages at 1.96 and 2.33 hours led to the two error peaks in the time series.




Figure 4.10: Number of visible satellites and GDOP computation for Run 2

These results show the performance of the three processing methods, adopting firstly the GNSS-only, dual-sensor, and finally triple-sensor integration in regular driving scenarios with moderate urban obstructions. These obstructions lead to error spikes in Run 1, as well as brief GNSS outages from overpasses in Run 2. It remains to demonstrate the algorithm in a consistently challenging environment.

4.3.3 Results with Simulated Outages

To demonstrate the ability of the algorithms to handle challenging urban environments, the engine is tuned to implement artificial obstructions that consistently block satellite signals, allowing only a few satellites of high-elevation angles to be visible.

The software mask function blocks an arbitrary range of azimuth and elevation angles over a user-defined period of epochs. All visible satellite matches the blocking rules and gets bypassed in the engine, simulating completely blocked signals. The function does

not simulate multipath errors.

Run 1 was used to test the simulated outages. Two separate masks were individually implemented.

Four-satellite Scenario over 30 Seconds

The first mask was implemented over a 30-second period, blocking 17 out of 21 visible satellites. Figure 4.11 shows the sky plot for the 30 seconds. An all-around mask blocks all satellites below an elevation angle of 60° , leaving only four satellites of G16, G04, R19 and E16 visible.




Figure 4.11: Sky plot of the four-satellite mask implemented on Run 1

The masked data was processed in the same way as the original datasets. On top of the three processing methods, RCC was also tested, as the continuous outage warrants the use of RCC. Figure 4.12 shows the time series of horizontal errors near the masked

Table 4.4: Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM and PPP/TC/RCC solutions in Run 1 with the four-satellite mask

Solution	PPP	PPP/TC	PPP/TC/RCM	PPP/TC/RCC
Horizontal rms (all data, m)	>10	1.76	1.01	1.00

period. All integrated solutions managed to overcome the surge in errors found in the PPP solution. The traditional TC solution saw a drifting pattern of errors until the mask was taken off. The RCM method produced a more robust solution. The same conclusion can be drawn for the RCC solution, with a steadier and accurate estimation in the first 15 seconds before seeing an error growth. Table 4.4 shows the rms errors for the displayed period. The clock-aided solutions produced a 43% improved accuracy, reducing the horizontal errors over the period to nearly a single metre level.




Figure 4.12: Time series of horizontal errors over the first masked period in Run 1

The result has verified the performance advantage of sensor fusion in a challenging GNSS environment, as expected. Furthermore, the results have demonstrated the improved ability of the clock-aided, triple-sensor receiver with clock-aiding algorithms in a

four-satellite environment.

Three-satellite Scenario over 20 Seconds

The critical advantage of clock-aiding algorithms is the ability to compute three-satellite solutions without rank deficiencies. In this analysis, a 20-second segment of the dataset is masked with a more challenging mask to allow only three satellites of the highest elevations to be visible. In such a scenario, the processing becomes more challenging not only because of the low number of observations, but also because of the poor geometry. As shown in Figure 4.13, the number of visible satellites was reduced to three: G4, G16 and E33. All of the visible satellites were above 70° of elevation.




Figure 4.13: Sky plot of the three-satellite mask implemented on Run 1

Figures 4.14 and Table 4.5 respectively show the time series of horizontal errors and

the error statistics for the period shown. Both clock-aiding solutions have demonstrated a superior performance over the traditional TC solution. The PPP solution diverged quickly beyond a 10-metre threshold when entering the simulated urban canyon, as expected. The solution appears to suggest that clock coasting provided a more accurate solution in both masked periods. However, it should be cautioned that the clock coasting method relies on the 'good' epochs right before the masked period to provide a trustworthy clock estimation. In actual urban canyons, the environmental obstructions tend to be more continuous, and the epochs before the three-satellite period may not be able to provide a good clock estimation. A less accurate clock estimation from earlier epochs will contribute to a higher clock coasting positioning error.




Figure 4.14: Time series of horizontal errors over the second masked period in Run 1

Table 4.5: Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM and PPP/TC/RCC solutions in Run 1 with the three-satellite mask

Solution	PPP	PPP/TC	PPP/TC/RCM	PPP/TC/RCC
Horizontal rms (all data, m)	>10	2.07	1.36	0.55

4.3.4 Summary

In this chapter, the research has presented a prototype triple-sensor fusion by assembling a CSAC-aided geodetic GNSS receiver. The prototype sensor fusion was tested in kinematic, suburban driving scenarios with urban obstructions. All sensor fusion solutions presented horizontal errors lower than those of the PPP-only solution, suggesting the necessity of fusing inertial sensors in the said scenarios. The goal of investigating the performance of the triple-sensor fusion was achieved as the PPP/TC/RCM solution gave more accurate solutions than the traditional solutions. The obstructions were further made more challenging by implementing artificial elevation masks to simulate deep urban canyons. In the new scenarios, the number of visible satellites was purposely reduced to four and then three over tens of seconds. Both receiver clock coasting and clock modelling methods were implemented, and they both produced a superior accuracy than the dual-sensor fusion of PPP/IMU. Particularly, the availability of a three-satellite solution, one of the primary advantages of clock-aided GNSS processing since the earliest research, has been verified. These results suggested that a triple-sensor fusion with a GNSS receiver, an IMU, and an external clock complemented with RCC or RCM algorithms can indeed provide an improved navigation performance in urban canyons.

Chapter 5

A Mass-Market Navigation System: PPP/IMU/OCXO Integration using Low-cost Hardware

The triple-sensor fusion built with a geodetic receiver, IMU, and CSAC serves as a prototype for the verification of the hardware integration, clock disciplining, and the validation of performance for the two clock-aiding algorithms. In this chapter, a mass-market-friendly version is studied by using more affordable alternatives to the high-end sensors.

The low-cost solution adopts an oven-controlled crystal oscillator to replace the CSAC. The GNSS receiver and antenna have also been replaced by automotive-grade alternatives. Unlike the higher-end products, the consumer-grade sensors are less customisable, presenting additional challenges in assembling, integrating and testing. While the clock-aided GNSS solution has been studied in the literature of recent years,

there have been no attempts to build a truly low-cost solution using automotive- or consumer-grade sensors.

This chapter begins with a thorough introduction of the oscillator used for the low-cost solution, with an emphasis on the adaptive approach to building the hardware solution. Then, the road tests approach, more real-world-oriented than the CSAC-aided GNSS solution, is presented. The vehicle was taken to bustling urban canyons in Toronto on streets with high rises and newly built condos. The datasets are studied, concluding a comparable performance for the low-cost approach with the CSAC-aided approach, with only a fraction of the cost.

5.1 Oven-controlled Crystal Oscillator (OCXO)

While the CSAC offers superior frequency stability over 100 seconds, it is typically less stable in a somewhat averaging interval, i.e., 0.1 to 10 seconds, compared to an Oven-controlled Crystal Oscillator (OCXO). Moreover, the CSACS are nonetheless too expensive for the majority of mass-market applications, where their price tags - often around a few thousand dollars - are deemed too expensive for the industry. This fact drives the research to look into possibilities beyond CSACs.

OCXO offers an ideal alternative. Clock-aiding algorithms such as RCM and RCC are expected to operate over short periods of GNSS challenging environments, typically not exceeding a few minutes. IMU plays a more important part if the GNSS signals are completely lost or the challenging environments are longer than a few minutes. This means that the CSAC's typical superior frequency over the averaging time of 10^2 to 10^3 is an overkill for the application in question. OCXOs do not have an edge in stability

over the CSACs over longer periods, but for shorter periods, the OCXOs are typically more stable. Cost-wise, the OCXOs are often offered at a much lower price, ranging from hundreds of dollars to a bit over a thousand. Admittedly, the price may still be over the budget in consumer navigation systems, but in a mass-production scenario, the OCXO may provide a commercially viable option for premium systems. Another caveat for adopting OCXOs is their much higher wattage. Due to their physical nature, the power consumption of an OCXO, using the Safran/Orolia GXClok-500 as an example, reaches 3W. However, these drawbacks do not imply inferiority compared to the more expensive CSACs if the power constraints are not as harsh. While the CSACs advantage in size and power consumption makes them well-suited for space applications, the OCXOs see extensive use in the telecom industry (Lim et al., 2010; Wojciechowski & Olsson, 2015).

The research adopts a widely available OCXO model of Orolia GXClok-500. Figure 5.1 shows an illustration of the package. The Clock module comes with a built-in u-Blox M8-T timing GNSS receiver and is about the size of one and a half USB memory stick. The module is size-wise comparable to the CSACs. The package also comes with a patch antenna intended to be used along with the built-in, consumer-grade u-Blox GNSS receiver module to discipline the clock over an extended period of time. These functions and modules are redundant for the low-cost, triple-sensor solution presented in the research. In an actual product of the triple-sensor solution, there is no need for a separate antenna and GNSS receiver for the clock, and therefore, the cost can be reduced below the mass-market package. In the following subsections, the operational principles and experimental procedures specific to the OCXO will be explored.



Figure 5.1: GXClock-500 package containing a u-Blox low-cost GNSS receiver chip and patch antenna

Operational Principles

Crystal Oscillators (XOs) are the most commonly used type of oscillators in electronics. Their birth was borne out of the discovery of the piezoelectric and inverse piezoelectric effects of quartz crystals. The former refers to the electric field produced by the crystal when a mechanical deformation is applied. The latter refers to the observation of mechanical deformation when an external electric field is applied. This discovery made it possible to construct an oscillator mechanism that exploits the resonance of quartz crystals in a conservation chamber, leading to the development of XO clocks.

The physical principles of an XO oscillator determine that its performance is impacted by both electrical and mechanical environmental conditions (Frerking, 1996; Vig, 1992). Particularly, the 'crystal' component, that is, the quartz crystal chamber, is temperature-sensitive. The sensitivity leads to a number of techniques to improve the XO clocks' stability by addressing the temperature sensitivity. Temperature-compensated crystal Oscillators (TCXO) are most commonly seen in consumer electronics, as well as the

majority of GNSS receivers. The TCXO employs a temperature compensation circuit to calibrate the frequency drift when the environmental temperature changes. Voltage controls are sometimes used to build Voltage-Controlled Crystal Oscillators (VCXOs) where a control voltage is applied to the crystal chamber to adjust the resonance through the understanding of the crystal's properties, thus reducing the effect of temperature changes. This technique is also seen to be used in conjunction with the TCXO technology, as seen in the Voltage-Controlled TCXO (VCTCXO) equipped by the NovAtel OEM7 GNSS receivers (NovAtel, 2019). The OCXOs present a more aggressive approach to controlling the impact of ambient temperatures. A thermally insulated chamber with a heating circuit, i.e., the oven, is constructed to enclose the crystal. The oven may also enclose the control circuit within to reduce the effect of ambient temperature. The oven is equipped with temperature sensors to allow the controller to heat up the oven as required. This process is similar to the CSAC's operation, where the ideal temperature must be reached in the Cs chamber to allow the optimal hyperfine resonating frequency. Likewise, an OCXO requires a few minutes to warm up, and an ambient temperature that is too low may cause the oven to take longer to warm up or be unable to reach the optimal temperature at all.

The OCXOs are affected by ageing, mechanical vibrations and shocks as well (Safran, 2022). In the CSAC, ageing is less prominent and is primarily associated with the degradation of the vacuum chamber. Crystal oscillators face structural deformation over time, contamination, and outgassing. All these factors contribute to frequency changes over time. Mechanical environments affect the oven in a similar way to the CSAC. Moreover, the mechanical environment affects the air pressure within the enclosure, affecting the behaviour of the crystal and the oven.

However, as in the case with the CSAC, the environmental factors are negligible in a daily driving scenario, over hours or days. In the long term, the clock can be disciplined by the frequency of a superior frequency source to reduce the effects of ageing.

Finally, Figure 5.2 shows a typical OCXO operation where the built-in GNSS receiver chip provides the 1 pps signal as the external frequency reference. The microcontroller calibrates the OCXO's frequency output through a controlling voltage and receives user commands from the PC. A 10 MHz oscillating signal is generated as the output.




Figure 5.2: Typical OCXO operation

Disciplining the OCXO

The OCXO model used in the research is a mass-produced module purchasable from common online vendors of electronics. The module is found to be more customisable as a component in the navigation system than the CSAC. To fully utilise the potential of the GXClok-500, an EK-5 evaluation board shown in Figure 5.3 was purchased for this research. The main purpose of the evaluation board is to divert the various input and output pins into the SMA and MMCX interfaces to connect to other sensors in the

triple-sensor fusion.




Figure 5.3: EK-5 evaluation board for GXClock-500 OCXO

The software interface that is provided by the package was rudimentary, so the control of the module was primarily done through manual serial commands via an RS232 cable. The evaluation board allows an external 1 pps signal ("PPSREF") to replace the internal 1 pps signal ("PPSINT") as the reference oscillator, and the identical NovAtel OEM7 receiver used for disciplining the CSAC was used for the OCXO. The GNSS receiver was connected to a rooftop antenna of geodetic grade in an open-sky environment.

Figure 5.4 illustrates the three PPS signal channels during a disciplining operation, or "tracking", as is referred to in the user manual. All three PPS signal channels are the inputs and outputs that can be found at the pins on the evaluation board. The 'PPSREF' refers to the reference PPS signal. If the internal GNSS receiver is used, it is the timing signal sent from the uBlox M8-T onboard. If an external oscillator is used, the external signal replaces the internal 1 pps signal. If there are no GNSS signals and the internal GNSS chip is used, the clock enters 'holdover' mode, and there is no PPSREF signal present. The 'PPSINT' signals are the internal signals tracked to the reference signal.

It is aligned with PPSREF within 50 ns if PPSREF is present and tracking is enabled. Otherwise, the PPSINT follows the 'PPSOUT' signal. The PPSOUT signal is the output signal based on the OCXO's oscillator. The user can sync the output signal to the PPSINT signal. If the sync mode is enabled and the OCXO is receiving reference signals from the GNSS receiver, the PPSOUT signal will become a copy of the GNSS receiver's 1 PPS output signal with a phase error margin of no more than 50 ns. The figure shows the operation of the 'tracking' command. Once the 'tracking' command is applied to the OCXO, the PPSINT signal will immediately try to align with the reference signal, and the disciplining process will begin. The phase difference between the PPSOUT and PPSREF will gradually reduce towards convergence by the disciplining algorithm.




Figure 5.4: PPS signals during a disciplining session

The analysis of the OCXO behaviour is largely based on the polled output from the command console. The OCXO is able to transmit the polling results in NMEA format. The study uses the entries of the "\$PTNTS" field where the clock's status, internal frequency, holdover frequency and saved frequency are recorded over time. Figure 5.1 shows two plots of the polled data during a start-up sequence with the GNSS feeding timing signal. The clock status is presented as an integer from the NMEA output. The status codes are mutually exclusive, and the active status code is shown with blue markers. According to the result, the clock status reached 'Syncing with PPSREF' in about 9 minutes and remained stable. The right plot shows the three frequencies used by the clock. The frequencies are in internal units. The saved frequency is a manually set frequency by the user. The holdover frequency is a rolling average over time computed by the disciplining algorithm when the clock is tracked to an external reference for a period of time. The actual frequency is the internal frequency that the clock is currently running, and it constantly tracks the external 1 pps signal in discipline mode. It is worth emphasising that the actual frequency is an internal measurement without a reference. Therefore, a constant 'actual frequency' still suffers from instabilities not detectable by the clock itself.




Figure 5.5: Status versus time (left) and internal frequency vs time (right) of the OCXO start-up sequence with external reference

The operation of the low-cost triple-sensor fusion using the OCXO differs from the CSAC-aided solution in that the OCXO is integrated with a GNSS receiver. If the GNSS signal is lost, the OCXO will automatically enter the holdover mode, where the steering pauses and the PPSOUT frequency remains unchanged until the timing reference comes back. This inherent hardware integration means the OCXO is sensitive to a GNSS-challenging environment and will adapt to it in real-time.

Figure 5.1 shows the disciplining process of the OCXO model over 24 hours. The left plot shows the status of the clock over the tested period. The figure shows that during the disciplining process, the clock remained consistently in disciplining mode. The right plot shows the actual frequency of the clock and the holdover frequency - a rolling

average of the tracked frequency over time. The data shows a convergence pattern for the actual frequency, and if the trend were to continue, the blue line showing the holdover frequency will approach the convergence limit of the green line and remain steady. An operational decision can be made at this point to manually set the holdover frequency as the actual frequency, assuming that the actual frequency in green is accurate. This will move the 'saved frequency' higher up towards the convergence limit, and the OCXO will have a default frequency at the limit.




Figure 5.6: Status versus time of a sample OCXO disciplining process (left) and internal frequency vs time for a sample OCXO disciplining process (right)

The Adaptive Approach to OCXO-aided Sensor Fusion

The OCXO-aided GNSS receiver differs from the higher-cost CSAC-aided GNSS receiver in that the OCXO excels in stability over 1 to 100 seconds but less so over a longer period. During a driving session, the OCXO will likely accumulate higher timing errors than those of the CSAC if not disciplined, which is expected to be in the range of 10^3 seconds of magnitude. Therefore, the OCXO-aided GNSS receiver is programmed to follow an adaptive approach by utilising its mode-switching function when the external 1 pps signal is lost.

Figure 5.1 shows a working sequence of the OCXO where the OCXO's GNSS receiver was purposely disconnected and then reconnected several times to simulate GNSS outages. The OCXO was programmed to be in holdover mode, so if the PPSREF is lost or unstable, the internal algorithm will switch to the holdover frequency. The clock initially warmed up to the saved frequency and began steering towards the external reference provided by the GNSS receiver. The plot shows that, as the OCXO loses GNSS timing reference regularly, the actual frequency was kept consistent at the holdover frequency computed at the beginning of the experiment.




Figure 5.7: Status versus time of an OCXO holdover test (left) and internal frequency vs time of the OCXO holdover Ttst (right)

Another possibility is that a not-yet-disciplined clock is put to use in such environments, and this leads to the 'free-run' scenario. In the 'free-run' scenario, the OCXO will revert to the saved frequency instead of the holdover frequency that has yet to converge. Figure 5.1 shows the case for a free-run test. The clock warmed up and began tracking the PPSREF signal from the GNSS receiver as usual, and then was forced into free-run mode. The clock stopped updating the holdover frequency and the attempts to discipline the actual frequency, but rather reverted back to the 'safe' saved frequency from the previous sessions. Operationally, the clock is expected to go through a nominal discipline session and update the saved frequency whenever it gets an opportunity to converge its holdover frequency calculation so that the saved frequency always remains a good backup.




Figure 5.8: Status versus time of an OCXO free-run test (left) and internal frequency vs time of the OCXO free-run test (right)

The OCXO-aided, low-cost, triple-sensor fusion will run adaptively to the environment. Figure 5.9 shows the logical scheme for the adaptive OCXO-aided GNSS receiver. The GNSS receiver and the OCXO form a closed loop with the GPS-synchronised 1 pps signal from the receiver to the clock, and the best-available timing signal from the OCXO to the GNSS. Determined by the availability of GNSS signals and the current settings of the OCXO, one of the three frequencies will be fed back to the GNSS receiver. The timing signal feedback is then used internally by the GNSS receiver for the tracking and measurement of GNSS signals. The observations are timestamped with the precise timing provided by the OCXO's feedback and can thus be used for RCC or RCM processing.




Figure 5.9: Operation flowchart for an adaptive OCXO-aided GNSS receiver

The closed-loop approach with low-cost equipment has several practical considerations, as evident by the issues observed in real-world testing. For example, the design of the low-cost GNSS receiver used in the research, namely the Septentrio Mosaic-T, does not have the ability to switch between the external oscillator and its internal clock in-run. If the clock fails to consistently produce the 10 MHz signal to the GNSS receiver for whatever reason, the receiver will not be able to continue tracking GNSS signals. These issues may require a deeper level of hardware coupling between the two sensors beyond this research. Their practical presentation and solutions will be presented in the following section when the experiment procedures are discussed.

5.2 Setup, Road Tests, Results and Analysis

The goal of this chapter is to move beyond the prototyping stage to more practical, user-oriented scenarios. The setup was assembled and taken into real urban canyons and underpasses to test the real-world performance of the triple-sensor solution and the two algorithms.

5.2.1 Experimental Setup

The rover-based system for road testing used in the CSAC-aided sensor fusion was modified to build the OCXO-aided, low-cost fusion. Figure 5.10 shows the experimental setup for data collection. The majority of the hardware remains identical to the previous CSAC-aided road tests, and the methodology to produce reference solutions is unchanged. However, there are two noticeable differences. The OCXO is connected to both a GNSS antenna and the GNSS receiver, that is, the Septentrio Mosaic-T. The OCXO's behaviour is similar to the IMU module, as both sensors are mass-market products with integrated GNSS receivers and require the GNSS time from the internal receiver. For the OCXO, the internal GNSS module provides the disciplining signal in run. The GNSS observations from the internal receiver are not used in any processing. For the XSens MTi-7 IMU, the internal GNSS receiver was necessary for its initialisation; losing connection to the antenna during the run could cause a system reboot and affect the data collection. For exactly this reason, the IMU was connected with the geodetic antenna to drop the IMU's antenna connection as a potential variable.




Figure 5.10: Experimental setup used for testing low-cost triple-sensor fusion

The OCXO-aided, low-cost navigation system was assembled using the adaptive approach described in the earlier section. The OCXO was disciplined overnight in the laboratory, and the converging frequency was recorded in the OCXO's memory. Whenever the OCXO detects a GNSS outage, affecting its ability to track with the GPS time, it will revert to a 'safe' holdover frequency near the saved frequency determined by its algorithm. Figure 5.11 shows the time series of clock status during a road test along an urban canyon. In 51% of the total epochs of the sample dataset, the clock reported a holdover status, indicating a failure to track its internal frequency towards the GPS time. This underlines the importance of disciplining by exposure in open-sky environments prior to driving into urban canyons.




Figure 5.11: Clock status over time during a road test in urban canyons

Another problem encountered during the road test was the low environmental temperature. Since many of the earlier tests were conducted in winter, there have been multiple occasions where the temperature was below -10°C . Under the low temperature, it has been observed that the system may take a few minutes longer than in a lab environment. In the extreme cold weather below -20°C , the GNSS receiver failed to initialise as the OCXO was not able to provide a stable enough timing signal for satellite tracking.

5.2.2 Results and Analysis

To validate the performance of the low-cost, clock-aided sensor fusion, the road tests were conducted in two sets of routes. The first set of the route was in the same suburban area as Run 1 and Run 2. The second set was a collection of urban driving scenarios in Toronto's busy North York area.

Suburban Dataset - Run 3

The low-cost triple-sensor setup was first taken to a suburban area with largely open arterial roads outside of rush hour. The route was designed to reflect a typical North American suburban area under urbanisation. In particular, the route involved two underpasses, and each of the two underpasses could cause a period of GNSS outage. Referring to Figure 5.12, the two underpasses were a railway bridge and a major highway underpass, respectively. The first underpass was narrow and caused only a few epochs of GNSS outage. However, the highway underpass measured about 100 metres in width and was a collection of overhead bridges. The highway underpass causes around 10 seconds of GNSS unavailability, with intermittent epochs of limited GNSS signals between the gaps of the bridges. The outages were complete, and no GNSS satellites were available when the vehicle was underneath the bridge.



Figure 5.12: Google View of the underpasses involved in Run 3

The solution was first processed with PPP, traditional tightly-coupled PPP/IMU integration and triple integration with RCM. Figure 5.13 shows the cumulative distribution of horizontal errors on the left and the time series of horizontal errors on the right. Table 5.1 shows the rms statistics of the horizontal errors. The dual- and triple-sensor fusions produced similar 95th and 67th percentile rms performance, indicating the OCXO had effectively eliminated more error spikes or peaks and produced a more stable time series. Meanwhile, both statistics outperformed the PPP-only solution. The two underpasses can be found at 1.915 and 2.135 hours, respectively. The first error spike present in the PPP solution was eliminated by the PPP/TC and PPP/TC/RCM solutions. However, the second and longer outage presented a challenge for all solutions, where an error spike can be observed in all solutions. This observation indicates that the current algorithms

can be further improved to deal with situations where there are alternating epochs of complete GNSS outage and poor visibility.




Figure 5.13: Results from Run 3 using low-cost triple-sensor fusion: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)

Table 5.1: Horizontal rms statistics of PPP-only, PPP/TC and PPP/TC/RCM solutions in Run 3

Solution	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP	4.28	1.32	1.30
PPP/TC	2.55	1.12	1.09
PPP/TC/RCM	2.09	1.12	1.09

The two natural outages are of particular interest to this research. Novel applications such as intelligent vehicles have demanding performance criteria for the positioning solution of the vehicles. Therefore, the trajectory of the three solutions was plotted and magnified for reference. Figure 5.14 shows the trajectory of solutions during the

two outages. It can be observed that the first and shorter outage caused no deviation in the navigation trajectory for the multi-sensor solutions, but a deviation from the reference can be seen from the PPP solution. All three solutions suffered from a path deviation when travelling underneath the major highway. However, the triple-sensor fusion demonstrated the least affected trajectory among all solutions. The observation has prompted further and more rigorous investigation of the navigation solution, with a specific focus on proving that the triple-sensor fusion approach from this research may benefit novel applications by means of integrity.




Figure 5.14: Trajectory of all solutions in Run 3 during the two outages: railway underpass (left) and highway underpass (right)

Suburban Dataset with Masks - Run 3

Following the methodology employed earlier to validate the performance of the prototype triple-sensor solution, the dataset of Run 3 was modified with artificial software masks to produce continuous epochs with a specifically low number of visible satellites.

Table 5.2 shows the information on the masks implemented. The first mask covered 20 seconds, where the elevation cut-off was increased from 65° to 70° and then to 75° , each time reducing the number of visible satellites by one. Similarly, a shorter mask lasting 10 seconds was implemented later in the dataset with increasing masking intensity to reduce the number of visible satellites from 5 to 4 and then to 3. Existing research in RCC based on CSAC-aided receiver design has proved that when the number of visible satellites is low, the marginal performance improvement from the external clock increases with each missing satellite till a three-satellite solution is reached.

Table 5.2: Description of the artificial elevation masks used in Run 3

	Duration (s)	Elevation Cutoff ($^\circ$)	Azimuth Range ($^\circ$)	Visible Satellites
Mask A	20	65/70/75	0-360	5/4/3
Mask B	10	60/70/75	0-360	5/4/3

Navigation solutions were processed using all four processing modes, including receiver clock coasting. The time series and cumulative horizontal error distribution are plotted in Figure 5.15, Figure 5.16 and Figure 5.17. There are a few observations that can be made as the software mask increases in strength. Firstly, there are four error spikes that can be observed across all four solutions. The first and last spikes were from the underpass travelling below the railway bridge and the highway, and the two spikes in

the middle were from the artificial masks implemented to evaluate the performance of the fusion schemes. Secondly, the clock-aiding algorithms have shown improved efficacy in reducing the error peaks. However, both algorithms still suffered from error spiking as the vehicle was travelling underneath the highway, where an intermittent GNSS signal was received between the two bridges. The receiver clock coasting algorithm has shown a slow convergence back to the original solution, leading to a worse performance than expected. This observation presents a caveat in using the receiver clock coasting algorithm, and further work is needed to address the problem.




Figure 5.15: Results from Run 3 with a 5-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)




Figure 5.16: Results from Run 3 with a 4-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)




Figure 5.17: Results from Run 3 with a 3-satellite mask: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)

The statistics of the solutions are presented in Table 5.3. A general trend of worsening performance as the mask strengths increase from 5-satellite to 3-satellite can be observed. The RCC fusion solution had a higher horizontal rms error than the traditional TC

solution due to the extended convergence time taken after the highway underpass. Meanwhile, the triple-sensor solution with RCM presented a noticeable performance improvement of 9% for horizontal rms errors and a slight improvement when only the 95th or 67th percentile errors are considered. Finally, Figure 5.18 shows an intuitive plot showing the performance improvement with the dual- and triple-sensor solutions.

Table 5.3: Horizontal rms statistics of for PPP-only, PPP/TC, PPP/TC/RCM, and PPP/TC/RCC solutions with different number of visible satellites in Run 3

Solution	Mask Strength	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP	5-sat	7.57	1.31	1.29
	4-sat	9.02	1.32	1.30
	3-sat	9.07	1.33	1.30
PPP/TC	5-sat	2.63	1.20	1.16
	4-sat	4.27	1.18	1.15
	3-sat	4.20	1.19	1.15
PPP/TC/RCM	5-sat	2.39	1.16	1.13
	4-sat	2.41	1.15	1.11
	3-sat	2.46	1.15	1.12
PPP/TC/RCC	5-sat	4.57	1.17	1.14
	4-sat	3.25	1.04	1.00
	3-sat	3.10	1.03	0.99




Figure 5.18: Comparison of the horizontal rms errors across all solutions with various mask strengths in Run 3

The PPP/TC/RCC solution improved counter-intuitively as the number of satellites declines in this scenario. A likely explanation is that increasing the elevation mask eliminated the single Galileo and the single GLONASS satellite from processing. The clock bias estimation from the two eliminated constellations was likely less accurate due to the low number of visible satellites in the epochs before the masks, so they contributed more uncertainty in the estimation.

This road test concludes the effort to validate the performance of the novel sensor fusion schemes following the same strategy in building and testing the CSAC-aided GNSS receiver solution. The low-cost triple-sensor fusion scheme was then taken for a series of road tests in actual urban canyons.

Performance Validation in Real Urban Canyons

Following the road tests in suburban areas, the setup was taken to the urban canyon in Toronto, where high-density condominiums or office buildings occupied both sides of the streets. Each road test was shorter in duration, limited by the size of the urban

area. Figure 5.2.2 shows the Google Street View of the roads where the data sets were collected. Since the roads in the business area can be busy at times, the solution suffers from higher obstructions from the traffic in some data sets, while less so in others.



Figure 5.19: Google Street views of the urban routes: Yonge Street and the surrounding streets

Figure 5.20 shows the cumulative distribution of horizontal errors and the time series for Run 4, which took place on Yonge Street in the above-mentioned area. Comparing the navigation solutions, an expected pattern of dual- and triple-sensor solutions reducing error spikes can be observed. This can be further confirmed by the statistics shown in Table 5.4. The solution has shown a 25% and a 20% improvement with the RCM than the RCC solution over the traditional PPP/TC dual-sensor solution.




Figure 5.20: Results from Run 4 on Yonge Street: cumulative distribution of horizontal errors (left) and time series of horizontal errors (right)

Table 5.4: Horizontal rms statistics of PPP-only, PPP/TC, PPP/TC/RCM, and PPP/TC/RCC solutions in Run

4

Solution	Horizontal rms (all data, m)	Horizontal rms (95th percentile, m)	Horizontal rms (67th percentile, m)
PPP	2.95	1.70	1.58
PPP/TC	2.47	1.86	1.75
PPP/TC/RCM	1.86	1.65	1.59
PPP/TC/RCC	1.96	1.77	1.70

The real urban canyon dataset differs from the suburban datasets with artificial masks in that we see a more consistently low number of satellites throughout. Figure 5.21 shows the number of visible satellites during the road test. The plot forms a concave U-shape, indicating a declining number of visible satellites when the car travels deeper within the urban canyon, and a recovering number of visible satellites when the car moves beyond the border of the business area. There had been occasions when the number of visible

satellites reduced to four, and each of the occasions lasted no more than two seconds. Unlike the sharp and continuous drop in visible satellites with simulated urban canyons, the pattern resembles more of a series of V-shaped cones as the vehicle moved past high rises.




Figure 5.21: Number of visible satellites and GDOP in Run 4

In a similar observation as Run 3 with simulated obstructions, the RCC solution struggled when facing a continuous set of GNSS outages. This behaviour became more prevalent in more challenging solutions in the urban canyon, where busier traffic had caused even lower satellite visibilities. In fact, the RCC solution may present a worse accuracy than the dual-sensor PPP/TC solution, if the algorithm continuously decides that RCC should be enabled due to poor satellite visibility, leading to undesirable rises in positioning errors due to outdated or not adequately converged receiver clock bias estimations. In contrast, the RCM had shown a more consistent performance throughout.

The triple-sensor fusion setup was taken to more urban and suburban areas for more

runs in the coming weeks. The solution was processed in dual- and triple-sensor mode. The RCC solutions did not show consistent performance in some of the urban canyon datasets: slow convergence after consecutive GNSS outages occurred more regularly in real urban canyons, and in the worst case, will lead to worse performance than PPP-only solutions. In simulated urban canyons or underpasses, the RCC algorithm is typically initiated when the vehicle enters a challenging environment from an open-sky solution. A few plausible reasons include: (1) the engine's attempt to re-converge to a new clock bias estimation was interrupted by another outage; (2) a less optimal clock bias estimation was used in one or more outages. However, as in the case with the highway underpass, the algorithm enters GNSS outages swiftly from another outage, leaving little room for the clock estimation to re-converge. It is possible to carefully tune the algorithm to fit each dataset manually by adjusting the conditional process that controls whether the engine enters/exits RCC in certain epochs. However, this defeats the purpose of creating a simulated real-time processing, which can be easily adopted into mass-market applications.

The statistics of the remaining runs are shown in Table 5.5. The table also shows the improvement of the triple-sensor solution in PPP/TC/RCM mode versus the traditional PPP/TC processing. In most datasets, the OCXO-aided solution improved the performance by around 30 %. Outliers such as Run 9, the very last drive at the end of the day, which presented no performance improvement, sometimes appear in the GNSS data collections. Several aspects, including hardware interaction, environmental factors, and post-processing, could be responsible for forming an outlier. To determine if there is a systematic reason for the outlier, future work is required to study further or to re-create the scenario.

Table 5.5: Horizontal rms statistics of PPP/TC and PPP/TC/RCM solutions in the remaining datasets

Horizontal rms (all data, m)/ Solution	PPP/TC	PPP/TC/RCM	Improvement
Run 5, DOY 113, 2024	6.29	3.92	38%
Run 6, DOY 114, 2024	4.25	3.74	12%
Run 7, DOY 114, 2024	2.29	1.86	19%
Run 8, DOY 114, 2024	1.25	0.94	25%
Run 9, DOY 114, 2024	2.47	2.48	0%
Run 10, DOY 115, 2024	3.88	1.98	49%
Run 11, DOY 124, 2024	1.65	1.24	25%
Mean	3.15	2.31	26%

This concludes the results and analysis of the OCXO-aided triple-sensor solution in real-world urban canyons. Previously, there have been no studies to build a low-cost, clock-aided GNSS solution. While existing literature has demonstrated a three-satellite solution with RCC and RCM using CSAC or similar high-end external clocks, no studies have been attempted in real-world urban canyons where the composite effect of signal blockage and multipath exists concurrently. The OCXO-aided solution was successfully assembled and tested, showing effectiveness in reducing positioning errors caused by challenging GNSS environments. It also marks the lowest-cost clock-aide solution as of the date.

The PPP/TC/RCM solution presented a noticeable positioning improvement in either the simulated environment or the real-urban canyons. As the number of visible satellites declines with stronger masking in the simulated urban environment, the triple-sensor solution with RCM presents less deterioration in accuracy. The RCC solution proved to be less consistent due to a number of possible reasons: (1) the OCXO is less stable

over 10 seconds than a CSAC; (2) activation and deactivation of the algorithm during the outages affect the clock bias used to coast, and the re-convergence after the outage stochastically. The study was extended with a number of additional data collections around the same area over different days. The PPP/TC/RCM triple-sensor solution was proven to produce more consistently improved navigation solutions and has achieved the goal for this study.

The study so far has demonstrated the feasibility of a clock-aided GNSS hardware and software solution that can effectively improve navigation solutions, particularly in terms of accuracy. The hypothesis of improved performance of the solutions in challenging GNSS environments with obstructions, such as urban canyons, has been validated. Nonetheless, as illustrated earlier in the dissertation, the clock-aided GNSS solution has its limits. It does not replace an IMU, and must be used in conjunction with the GNSS observations. Its potential lies in filling the gap between a complete GNSS outage and an open-sky scenario, where a GNSS solution exists but may be compromised for liability-critical or safety-of-life applications. The effort in these applications is to not only reduce the error peaks, but also to solve corner cases to provide a trustworthy solution for the user. In such scenarios, the performance criteria of accuracy alone are insufficient to support the adoption of the novel method, and additional arguments and analysis are required. Therefore, the next research question is, can the use of clock-aided GNSS solutions be justified in liability-critical or safety-of-life applications by improving the corner cases where traditional solutions fail?

Chapter 6

Interpretation in Real Life:

Position-Domain Integrity Analysis

Studying GNSS positioning often requires metrics and analysis beyond accuracy. While studies keep pushing the solution accuracy towards centimetre-level using lower-cost equipment and providing the best possible navigation solution for mass-market users, it is essential to look back at the solution itself to ask the question: "Is the solution trustworthy? Can I trust the result for Intelligent and Connected Vehicles (ICVs), etc.?" This study has so far presented a novel approach to GNSS solutions with clock-aided triple-sensor solutions. The solution involves a hardware component of an external clock and a software component with PPP combined with RCM or RCC algorithms. In this penultimate chapter, the research shall address the question of quantifying some of the performance benefits that the solution provides for emerging applications.

The question of trustworthiness is usually addressed with the concept of integrity (Groves et al., 2013; Kaplan & Hegarty, 2017). Integrity describes the level of trust in

the information supplied by the navigation system, indicating whether the system can be trusted to be used (Groves, 2013; Wang et al., 2012). The correctness of information supplied by the navigation system is traditionally tied to signal spoofing, attacks, detection, and mitigation. Studies have been particularly concentrated on integrity in the aviation domain, as aerial applications represent key challenges for GNSS in safety-of-life applications (Gebre-Egziabher, 2018).

In recent years, the proliferation of low-cost GNSS solutions for road applications, for example, Intelligent Transport Systems (ITS), ICVS, or electronic toll systems, has brought new challenges and requirements to describe the trustworthiness of GNSS solutions. The challenge arises from the risk of the correctness of the GNSS solution not from being compromised by potential jamming and spoofing but rather from the urban environments where GNSS positioning performance is challenged by satellite obstruction, multipath errors, Non-line-of-sight (NLOS) signals, interferences, and many other adverse effects (Groves et al., 2012; Wang et al., 2019).

Overcoming the challenges and introducing viable integrity monitoring methods have become essential topics in recent GNSS research. However, applying an identical integrity scheme from existing applications to emerging applications is not always feasible. A crucial difference between building up an integrity scheme for urban and for aviation applications is the loss of measurements due to urban obstructions, while open-sky scenarios are well-expected in the civil aviation domain. Another noticeable difference is that low-cost receivers and other sensors are most commonly employed in urban applications, and system manufacturers are more cost-sensitive. Lower-cost sensors lead to greater negative impacts from environmental factors and worse measurement quality, and their sensor fusion creates a difficult and complex challenge for integrity

monitoring.

The chapter will review the relevant standard concepts that lead to GNSS integrity, existing methods in integrity monitoring specific to urban and vehicular applications, alternatives to traditional GNSS integrity models that describe the correctness of GNSS solutions, and finally, applied to the triple-sensor fusion for evaluation.

6.1 Overview of GNSS Evaluation Parameters

Evaluating a GNSS navigation system's performance typically calls for four parameters of accuracy, integrity, continuity and availability. The four parameters are sometimes represented as a pyramid, demonstrated in Figure 6.1.




Figure 6.1: Pyramid of the four parameters used to evaluate GNSS performance

6.1.1 Parameters for Evaluating GNSS Solutions

Availability stands at the top of the pyramid and describes the system's ability to provide a solution satisfying the application requirement, and is measured as the percentage of the time duration for which the system is active. As defined by ICAO (International

Civil Aviation Organisation), integrity is the percentage of time that the services of the system are usable by the navigator (Zhu et al., 2018). In the case of the proposed Ph.D. dissertation, where land vehicles on roads are considered, availability should measure the percentage of epochs where the system reaches a specified performance requirement. The requirement can be strict, as complying with recent industry standards or government regulations such as (SAE International, 2021; Standardization Administration of the People's Republic of China, 2024), or can be relaxed, as long as a meaningful output is provided. The latter is not always possible due to GNSS outages when the number of visible satellites is too low, and in such a case, the performance of a dead-reckoning sensor must be considered. In all cases, defining the availability of the system in question requires an evaluation of all three remaining parameters: accuracy, integrity, and continuity.

Accuracy measures the deviation of the estimated position solution from the ground truth. The accuracy of a GNSS solution may refer to the epoch-wise positional difference between the estimated states and the truth, as measured. In the analysis of navigation solutions from earlier chapters, the accuracy of each solution was presented as both a time series and a cumulative distribution function.

The next concept is integrity, measuring the level of trust the user should have on the navigation output. Integrity monitoring (IM) requires a model to evaluate integrity and provide timely warnings to the users should a threshold of trustworthiness be breached. Integrity introduces two additional quantities on top of the position error (PE), including the alert limit (AL) and the protection level (PL). The former provides a statistical upper bound for the horizontal errors, since the horizontal errors are not computable in real-time; the latter is defined according to the safety requirement of the system, where an

error beyond the AL is considered unsafe. Figure 6.2 shows the Stanford diagram of integrity (Tossaint et al., 2007). The safe operation of the navigation system requires that PL is no greater than AL, and the design of the integrity monitoring scheme should avoid situations where PE is greater than PL, indicating a modelling fault.




Figure 6.2: Stanford Diagram (Tossaint et al., 2007) for evaluating integrity performance

Breach of trustworthiness can be a result of intentional action, such as GNSS jamming in conflict zones, or a result of natural occurrences that affect GNSS signals. Studies into integrity often relate to the realm of spoofing and detection, but in the case of this research, the integrity is focused on the effect of urban environments.

Continuity is the third key performance parameter for building up availability. Continuity historically refers to the availability of GPS signals due to the lack of them in places or times around the world at a time when there were many fewer GNSS satellites in

orbit. In modern times, the concept is more commonly used to describe the continuity provided by the navigation system to provide a solution. An example of continuity loss is when a GNSS-only solution enters a complete GNSS outage. Without dead-reckoning sensors, the GNSS-only solution is expected to discontinue when such an outage is encountered.

Availability is built upon the three parameters and presents the system's reliability in supplying trustworthy information to the user to make safety-of-life decisions. The primary challenge to evaluate the performance of the triple-sensor fusion is to justify the system's practical performance in emerging applications. Previous work has been focused on the accuracy improvement in different scenarios, and on proving the concepts of RCM and three-satellite solutions, but to push the work into industry, more work is needed in the domain of integrity analysis based on Required Navigation Performance (RNP) criteria. Specifically for road users, the navigation performance is divided into either local navigation or absolute navigation. Local navigation is the positioning with respect to local terrain features as reference points for manoeuvring. Local sensors, such as cameras, are often used for the purpose. Absolute navigation refers to the ability to pinpoint the vehicle in a geographical map and is used more often in liability-critical applications. The study, as well as the report, will focus on absolute positioning.

For vehicle applications, the continuity requirement is often overshadowed by other parameters and is less important. The reason is that accuracy and integrity are more crucial in the safety-of-life applications. Some applications have contrary requirements where the accuracy becomes less important, and the requirement for AL can be much higher, as seen in Location-Based Services (LBS) apps (Rizos, 2005; Tiberius & Verbree, 2004). Examples include ride-hailing apps or vehicle-sharing apps. In these applications,

a timely yet rough estimation of position will satisfy the requirement.

6.1.2 Defining an IM Framework

A few definitions regarding the integrity monitoring framework shall be established.

1. Alert Limit (AL) AL is an externally defined parameter for allowable errors in either horizontal or vertical positions. An error beyond the AL is considered unsafe. If PL exceeds AL, a warning must be issued, and the system output is considered unavailable.

2. Protection Level (PL) PL is a statistical concept defined as the upper bound of the position difference between the system output and the truth. For road applications, horizontal protection levels (HPL) are more commonly used. Computing PL requires modelling of the navigation system known as IM (integrity monitoring) framework.

3. Time to Alert (TTA) TTA is the allowable time duration between when a positioning failure occurs and when an alert is announced by the IM framework. Should a failure become unannounced or announced later than the TTA, it becomes an integrity failure.

4. Misleading Information (MI) MI occurs when the actual position error exceeds the PL. MI could indicate a design flaw in the IM framework if the modelling of parameters is not accurate. When modelling navigation systems, a balance is required between being too conservative (PL is too high compared to PE) and being too tight (MI becomes more possible). Complex navigation systems introduce more error sources, such as the case of GNSS-PPP/IMU sensor fusion, where a very different sensor is introduced, and more error sources are introduced into GNSS processing. The complexity greatly increases the difficulty in modelling the IM framework, and more often than not, some error sources are ignored completely. By ignoring error sources or the additional sensor outputs themselves, the design and computational load are reduced, but MI becomes

more likely.

Balancing the complexity of the model with practicality is crucial, but a scenario that should be absolutely avoided is when $PE > AL > PL$, or the event of Hazardously Misleading Information (HMI). In this case, the system is declared available, and no alerts are issued as $PL < AL$. In the meantime, the user is unsafe because $PE > AL$.

6.2 Building an Integrity Monitoring Framework

There are several ways to build an IM framework. A more traditional approach is to introduce external information through augmentation systems. ABAS (Airborne Based Augmentation System), GBAS (Ground Based Augmentation System) and SBAS (Satellite Based Augmentation System) are the typical sources for such augmentation (Salós et al., 2010).

This section will begin with a brief review of the classic RAIM method and recent innovations in RAIM, followed by the argument that a position-domain IM framework is preferred in this research, and finally, the position-domain IM is introduced with historical background based on the concepts already established with the RAIM introduction.

6.2.1 Classic and Recent Innovations in RAIM

RAIM (Receiver-autonomous integrity monitoring) covers a range of methods to establish IM framework without augmentation systems or any external information. They are observation/measurement domain methods using either a snapshot scheme or a filtering scheme. Classic RAIM involves hypothesis testing where at least one test statistic is used for fault detection. Many RAIM schemes have been developed over the

years (Brown, 1992; Lee, 1986, 2004). The majority of the RAIM algorithms are based on checking the consistency of range measurements using the redundant measurements. RAIM algorithms may be categorised into a few types based on their domain of consistency check, the kind of measurement being used, algorithms or capability of Fault Detection and Exclusion (FDE).

Classic RAIM methods are observation/measurement domain methods using either a snapshot scheme or a filtering scheme. Classic RAIM involves hypothesis testing where at least one test statistic is used for fault detection. Snapshot schemes (Kaniuth et al., 2005) use the current epoch's satellite information and measurement to achieve an instantaneous position fix. Snapshot schemes are often based on Least Squares (LS) estimation, providing various computational efficiencies. Filtering schemes tend to perform better than snapshot schemes in PL estimations, but they can be computationally expensive.

Advanced RAIM (ARAIM) emerges as a multi-constellation solution compared to the GPS-based traditional RAIM method (Blanch et al., 2012; Nikiforov & Roturier, 2005; Simon et al., 2010). It is developed from a solution separation algorithm, where two solutions are compared: one with all visible satellites and another one excluding one satellite. While classical RAIM focuses on FD of code measurement and many novel RAIM methods are focused on combining code with other measurements, SS RAIM analyses the positioning solution of the navigation algorithm. The FD test statistic is determined by the maximum distance between the solutions obtained by removing each satellite. SS ARAIM requires multiple threshold tests for each of the coordinates and is, therefore, computationally more cumbersome. ARAIM has the advantage of being modifiable based on operational requirements, such as integrity parameters and

probabilities that are fixed in traditional RAIM. It is an example of a position-domain RAIM.

Beyond ARAIM, techniques such as Relative RAIM (RRAIM) and Carrier-based RAIM (CRAIM) (Feng et al., 2012; Gratton et al., 2010; Ochieng et al., 2010), where different measurements are used, provide an alternative to classical RAIM methods with various performance in PL computation and computational complexity.

6.2.2 An Alternative Position-Domain Approach

Whilst the RAIM and related approaches are commonly seen in IM literature, particularly concerning the spoofing and detections, the approach is impractical for this research to demonstrate the real-life implications of the clock-aided triple-sensor fusion through the lens of integrity. There are two major challenges: the lack of measurements and the overwhelming complexity of the three sensors.

Firstly, a major challenge in low-cost GNSS navigation in urban environments is the obstruction of signals. In the context of IM, this presents a unique difficulty as RAIM relies on the redundancy of signals. A low number of visible satellites with likely high multipath errors or NLOS signals will degrade the performance of most traditional RAIM methods or render them inoperable. The TC-based GNSS/IMU scheme may also face challenges as the low-cost nature of both sensors makes each of them an underperforming external reference for the other.

Furthermore, a range of methods have also been developed for integrating inertial sensors or other sensors. These methods usually require characterisation of the fault from additional sensors. In the example of a tightly-coupled (TC) scheme, the IMU can provide additional data for the fault detection of GNSS, and vice versa, for detecting

time-dependent error growth in IMU (El-Mowafy & Kubo, 2017). ARAIM, MSS (Multiple Solution Separation) RAIM and MHSS (Multiple Hypothesis Solution Separation) may also be used in GNSS/IMU IM schemes (Joerger et al., 2014). However, it should be noticed that the FD for the inertial sensors are often ignored. In the reality of low-cost and especially ultra-low-cost GNSS/IMU integration with handheld devices, the user may encounter faults in IMU outputs rather often. The inclusion of the external clock further introduces another factor of complexity that can be difficult to model and to verify through data collections.

These factors mean that any of the conventional IM frameworks in the observation domain will be impractical for the scope of the research. However, since the goal is to build a case showing the real-life implications of the novel sensor fusion scheme, there is no requirement for real-time integrity monitoring. In such a case, a convenient IM framework can be built upon the position-domain, utilising known position references, and is conducted in a similar manner as found in the aviation industry.

6.2.3 Applicability of the Position-domain IM Framework Based on Extreme Value Theory

Position-domain IM frameworks were seen in a very different application in the history of GNSS. In aviation, researchers developed an IM method based on the bounding of position errors without the need to collect years of data. The challenge faced by the researcher is that, in a similar situation, a rigorous argument of integrity improvement is required for a novel GNSS navigation solution.

Extreme Value Theory in GNSS Analysis

Classical IM methods presented a significant challenge when the demonstration of GNSS integrity for SBAS services was required for aerial purposes. The International Civil Aviation Organisation (ICAO) requires demonstration that the correctness of the position is guaranteed at a 1×10^{-7} order of magnitude (Braff, 2003). However, to show statistical significance through classical IM methods required several tens of years of data in the hunt for rare loss-of-integrity events. A statistical alternative tailored specifically to the tails of the error distribution was needed.

The tails of GNSS positioning errors may not be modelled as Gaussian, as many of the error sources, such as tropospheric and ionospheric errors, do not follow a Gaussian distribution. The result is a 'heavy tail' characteristic of GNSS positioning errors. Researchers were most interested in modelling the heavy tails as they represent the 'corner cases' where loss-of-integrity events reside. There were a few publications that recognised the non-Gaussian nature of the problem and began solving it through a position domain monitor or different bounding functions (Braff, 2003; Rife et al., 2004; Rife & Gebre-Egziabher, 2007).

To produce a statistical bound for the non-Gaussian, small-sample tails, a branch in statistics, the Extreme Value Theory (EVT) was applied, first for the European Geostationary Navigation Overlay Service (EGNOS), which has successfully evaluated and predicted the integrity of position errors (Azais et al., 2009; Kannemans, 2010; Ober et al., 2014). EVT deals with extreme deviation from the median of statistical distributions and is applied in various predictions of probability distributions, such as anomalous network traffic, road safety analysis, and equity risks. All of the applications require knowledge of rare but potentially devastating event probabilities, as in the case of Safety-of-Life

analysis in GNSS positioning.

Beyond GNSS, the EVT has seen extensive applications in dealing with similar peaks-over-threshold (POT) applications in fields including insurance, actuarial science and risk management (Haan & Ferreira, 2006).

Statistical Investigation on Horizontal Error Distribution

The applicability of the EVT depends on non-Gaussian, small-sample tails. This assumption matches intuitively with the horizontal errors seen in urban environments, where environmental blockages produce unbounded horizontal errors for a few consecutive epochs before the solution converges back. Sometimes the position change between epochs is limited by the user in anticipation of such position 'jumps' erroneously occurring in challenging environments. For the sake of statistical analysis, the study does not implement any arbitrary upper bounds to the horizontal errors or position changes between epochs.

To support the assumption, the horizontal errors found in Run 4 in real urban canyons are plotted as normalised histograms. Figure 6.3 shows the histogram for PPP-only solution, and Figure 6.4 shows the histogram for the dual-sensor fusion solution. The red solid line indicates a Gaussian fit of the errors. A heavy tail can be observed in either of the two plots, as illustrated with grey eclipses. The PPP-only had a longer tail that extended into larger horizontal errors not seen in the fused solution. While the sensor fusion effectively reduced the error spikes, the solution exhibits a non-Gaussian tail distinct from the majority of the error distribution.




Figure 6.3: Normalised histogram of PPP Horizontal Errors in Run 4




Figure 6.4: Normalised histogram of PPP/TC Horizontal Errors in Run 4

Additionally, the existence of the heavy tails highlights the 'corner cases' that this research is focused on tackling. Traditional sensor fusion and PPP improvement have

been bringing down the RMS errors over the years with low-cost equipment, but it remains to reduce these corner cases where safety-of-life applications may be impacted.

Applying EVT in GNSS Solution

EVT applications are typically approached by either deriving the block maxima/minima series or extracting peak values above/below a pre-defined threshold. For modelling the tails of positioning errors to evaluate integrity, the latter approach is most commonly chosen. Let $G_L(X)$ be the distribution of some random variable X over a given threshold L , we have

$$G_L(X) = P(X \leq x + L | X > L) \quad (6.1)$$

A few different parametric models exist in the EVT analysis of positioning errors, and the most commonly considered are the Generalised Extreme Value (GEV) distribution and the Generalised Pareto Distribution (GPD) (Ahmad et al., 2014; Larson et al., 2019; Li et al., 2023). GEV is a limiting distribution of the sample maximum of independent identically distributed (i.i.d.) random variables, and we assume a large sample size by taking the sample size $n \rightarrow \infty$. The two distributions are inherently related, and a detailed proof and explanation will be given. The choice of L in Equation 6.1 must seek a balance to be large enough to hold the asymptotic assumption while small enough to include sufficient data for an accurate parameter estimation. GPD is a univariate distribution long used to model heavy tails of any distribution. Its CDF is given by:

Let $X \sim GPD(\mu, \sigma, \xi)$, for $\mu, \sigma, \xi \in \mathbb{R}, \sigma > 0$

$$G(x) = \begin{cases} 1 - \left(1 + \frac{\xi(x-\mu)}{\sigma}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - \exp\left(-\frac{x-\mu}{\sigma}\right) & \text{if } \xi = 0 \end{cases} \quad (6.2)$$

The three parameters are respectively the shape ξ , the scale σ and the location μ . The shape parameter ξ ranges from $-1 < \xi < 1$. For $\xi < 0$, the tail has finite end points and errors beyond its endpoint cannot occur. For $\xi = 0$, the errors have an exponential tail. For $\xi > 0$, the errors have a heavy tail. Once the overbounding function is determined, the protection level can be computed by taking the inverse of the CDF function for the following integrity analysis.

6.3 Integrity Analysis

The goal is to develop a scheme to evaluate the integrity performance of the triple-sensor fusion in the dataset and prove that the integrity requirements are met. Hazardous Misleading Information (HMI) is where $PE > AL$ given HPL computed is not greater than HAL, and Misleading Information (MI) is where $PE > PL$. System availability is determined by the ratio of epochs where $PL < AL$ over all the epochs. That means what we are particularly interested in is the tails of the error distribution, where PES are the highest.

However, as seen in the positioning solution improvement from the sensor fusion with a precise external clock, it is intuitive to hypothesise that the novel solution provides a more ‘trustworthy’ performance under adverse environmental conditions such as urban canyons. Without access to conventional IM methods, proving the hypothesis requires studying the post-processed solution to transform the accuracy metrics into

statistical probabilities indicating generalised integrity metrics. The transformation is essential in the proposed sensor fusion, as the statistical significance is required to prove the safety-of-life performance of the fusion.

Translating the definition into probability, the probability of HMI is given by

$$Pr(HMI) = Pr(HPE > HAL | HPL \leq HAL) \quad (6.3)$$

To model the tails, the EVT is invoked. EVT seeks a limiting distribution of the sample maximum of independent and identically distributed (i.i.d.) random variables. Let a sequence of i.i.d. random variables be $Y_1, \dots, Y_n$ with distribution function G . Define $M_n = \max\{Y_1, Y_2, \dots, Y_n\}$ which follows the distribution of F^n . It has been proven that the only possible non-degenerate limiting distribution of the sample maximum is the GEV distribution, given by

$$G_\xi(x) = \begin{cases} \exp\left(- (1 + \xi x)_+^{-\frac{1}{\xi}}\right), & \xi \neq 0 \\ \exp(-\exp(-x)), & \xi = 0 \end{cases} \quad (6.4)$$

where the subscript indicates $x_+ = \max(x, 0)$.

Proof of Proposition

Proposition: Assume an error distribution follows $X \sim GEV(\xi, \sigma, \mu)$, with cdf $G(x)$. Then, for a given threshold L , the excess distribution $G(x|L)(x) = Pr(X - L \leq x | X > L)$ converges to a GPD as $L \rightarrow \infty$.

Proof: Consider the GEV distribution G in terms of extreme events where a random

variable X exceeds a given threshold L , and take the threshold to infinity,

$$\lim_{L \rightarrow \infty} \Pr(X > L + x \mid X > L) = \frac{1 - G(L + x)}{1 - G(L)}, \quad x > 0 \quad (6.5)$$

Following the conditions for sample maxima convergence (Pickands III, 1975), there are parameters ξ, σ, μ that for large n , the following holds (Dey & Yan, 2016):

$$G^n(z) \approx \exp \left\{ - \left[1 + \xi \left(\frac{z - \mu}{\sigma} \right) \right]^{-1/\xi} \right\} \quad (6.6)$$

$$\text{That is, } n \log G(z) \approx - \left[1 + \xi \left(\frac{z - \mu}{\sigma} \right) \right]^{-1/\xi} \quad (6.7)$$

$$\text{Since } \lim_{z \rightarrow \infty} G(z) \rightarrow 1, \quad \log G(z) \sim G(z) - 1, \text{ Eq. 6.7 is equivalent to } 1 - G(z) \sim \frac{1}{n} \left[1 + \xi \left(\frac{z - \mu}{\sigma} \right) \right]^{-1/\xi} \quad (6.8)$$

Evaluating z at $L + x$ and L , following Eq. 6.5, we have

$$\Pr(X > L + x \mid X > L) \sim \left[\frac{1 + \frac{\xi(L+x-\mu)}{\sigma}}{1 + \frac{\xi(L-\mu)}{\sigma}} \right]^{-1/\xi} = \left[1 + \frac{\xi x}{\sigma_L} \right]^{-1/\xi}, \quad \sigma_L = \sigma + \xi(x - \mu) \quad (6.9)$$

satisfies

$$\Pr(X > L + x \mid X > L) = 1 - \Pr(X - L \leq x \mid X > L) = 1 - G_{x|L}(x) \quad (6.10)$$

Consequently,

$$G_{x|L}(x) \sim 1 - \left[1 + \frac{\xi x}{\sigma_L} \right]^{-1/\xi} \quad \text{for } \sigma_L = \sigma + \xi(x - \mu) \text{ and } \xi \neq 0. \quad (6.11)$$

Therefore, a rigorous connection with GPD is established when the exceedance of a threshold L is considered. Let P_{IR} be the probability of integrity risk defined by $HPE > HPL$, then

$$P_B = 1 - P_{IR} = 1 - \left[1 + \frac{\xi(HPL - \mu)}{\sigma} \right]^{-1/\xi} \quad (6.12)$$

$$\text{Solving for } HPL, \quad HPL = \frac{\sigma}{\xi} \left[(1 - P_B)^{-\xi} - 1 \right] + \mu \quad (6.13)$$

6.3.1 Application in Real-world Data

With a computable HPL in the position domain, the analysis will be applied to the real-world data collected for the low-cost triple-sensor fusion.

Identifying the Tail Distribution

To apply the GPD modelling to the tail of the horizontal errors, it is crucial to first determine if the chosen threshold value L is valid, without either violating the asymptotic assumptions with a threshold too low or underestimating extreme events with a threshold too high. The threshold choice can be diagnosed by ME (Mean Excess) functions.

ME is defined as the expected value of the excesses of the sampled error over different potential thresholds L , that is $ME = E(X - L | X > L)$. An estimation of ME is given by:

$$\widehat{ME} = \frac{1}{k} \sum_{i=1}^k X_i - L \quad (6.14)$$

Applying Equation 6.14 to Run 3 collected with a low-cost triple-sensor setup, the Mean Residual Life (MRL) plot, where one can diagnose visually the suitability of the

method. First, the MRL plot for the raw data without simulated masks is shown in Figure 6.5. There are three major observations. First, the PPP-only solution presented a much longer tail, where the maximum threshold value was over 100 metres. The longer tail indicates the existence of a wider error distribution than the two fused solutions. Second, all of the solutions exhibit linearity beyond a certain inflexion point. The linearity is an expected phenomenon and suggests a Pareto-like, heavy-tail behaviour. Third, the OCXO-aided triple-sensor solution and the traditional PPP/IMU dual-sensor fusion show a relatively similar behaviour. However, the linear portion for the triple-sensor solution began earlier than the dual-sensor solution, suggesting a lower occurrence of extreme values.




Figure 6.5: MRL plot for Run 3 without simulated outages

Before moving on to the real urban-canyon dataset, it is reasonable to first test the method with simulated outages. In earlier analysis, this dataset was implemented with simulated outages of various strengths. The solution with the 4-satellite mask is plotted in Figure 6.6. The PPP-only solution suffered from an even longer tail. For the sake of

readability, the x-axis was cut off at 115 metres. The mask altered the tail behaviours for the fused solutions. While this alteration is difficult to make conclusions visually, the shape factor ξ will be affected mathematically. The MRL plots are qualitative tools to interpret the extreme value behaviours of the error statistics. Just as the determination of the alert level in integrity monitoring, the interpretation of MRL plots can be subjective to the degree demanded by the specific requirements of the application.




Figure 6.6: MRL plot for Run 3 with a 4-satellite mask

Plotting the MRL of Run 4 in Figure 6.7, which is the real-world urban-canyon dataset collected with the same setup, the error distribution shows a more irregular shape. Instead of a clear inflexion where the Gaussian and non-Gaussian distributions can be distinguished, the real-world data behaves less predictably with multiple 'humps'. This behaviour is reasonable since in a real-world urban canyon, the vehicle encounters a series of obstructions of various sizes, while in the simulated outages, only a few outages are implemented. At roughly $L_0 = 2.5m$, as indicated as a grey dashed line, the three solutions begin to show a non-Gaussian pattern. It matters not that the shape for each

distribution is different after this point, as they will be modelled individually.




Figure 6.7: MRL plot for Run 4

Applying GPDs to the tails

The modelling requires identifying the non-Gaussian 'heavy tail' that must be fitted with EVT. The selection of the L_0 value, where the tail begins, will affect the goodness of the fit for the tail. While the choice can be subjective, it is crucial to select a value that is reasonable according to the application in mind. A lower L_0 will produce a more stringent bounding function, which will eventually lead to a lower HPL. Conversely, a higher L_0 will lead to a more relaxed HPL, where HMI becomes less likely, but the solution becomes less meaningful. If the goal is to produce the best fit of the tail, then the problem is solved by identifying the inflexion point in the MRL plots.

By identifying an L_0 , the tails of along-track and cross-track errors can be fitted with GPD. To illustrate how the selection of L_0 impacts the fit, three L_0 values of 0.5 m, 2.5 m, and 5 m are illustrated and fitted to the tails of the PPP/TC solution in Figure 6.8. The figure is enlarged to show only the top 4.6% (i.e., two sigmas) of the error distribution.




Figure 6.8: Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC) with different L_0 values

The observation from the plots indicates that an L_0 too low fails to bound a large portion of the errors, and an L_0 too large fails to model some of the error distribution that is supposed to be part of the tail. The L_0 value identified earlier provides a reasonable fit of the cross-track errors, but the bounding is off in the along-track errors. Therefore, the subsequent analysis will treat the two horizontal error components separately due to their different statistical behaviour.

Vehicular applications commonly require the horizontal errors to be represented by cross-track and along-track errors. The former indicates a deviation of the trajectory of the vehicle laterally across car lanes and is thus significant in interpretation. The analysis to follow will therefore be dividing the cross-track and the along-track components of the horizontal errors. In some POT (Peak-over-Threshold) applications, the choice of the threshold L_0 is curated in such a way that optimally describes the tail behaviour. However, the study has identified a goal of automotive applications. To produce a universally comparable result, the value of L_0 is defined as 1.0 m for each of the two components. Existing literature has typically required an accuracy of 1.5 m for lane-

keeping and a number of manoeuvres. Therefore, a component-wise threshold of 1.0 m fits the requirement, and allows some buffer.

The method is now implemented with the three solutions of PPP-only, PPP/TC, and PPP/TC/RCM processing modes. Figures 6.9, 6.10, and 6.11, respectively, show the tail fitting for the three processing modes.




Figure 6.9: Component-wise HPE and tail fitting with GPD for Run 4 (PPP-only)




Figure 6.10: Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC)




Figure 6.11: Component-wise HPE and tail fitting with GPD for Run 4 (PPP/TC/RCM)

The fitted GPD parameters for each processing mode are listed in Table 6.1. The shape parameter ξ indicates whether the tail tends to be infinite or bounded. If $\xi < 0$, the tail is modelled as bounded. If $\xi > 0$ the tail is modelled as unbounded. The scale parameter σ indicates whether the exceedances concentrate or are widespread. A larger σ means the exceedances are wider and the error distribution has more variations.

Table 6.1: GPD fitting parameters computed for each processing mode in Run 4

Solution	Component	Shape ξ	Scale σ
PPP	Cross-track	0.19	0.63
	Along-track	0.81	0.52
PPP/TC	Cross-track	-0.13	0.80
	Along-track	0.41	0.79
PPP/TC/RCM	Cross-track	-0.27	0.86
	Along-track	0.14	0.69

Calculating Protection Levels and Solution Availability

With the parameters calculated, the corresponding HPL can be computed using Equation 6.13. Equation 6.13 requires defining the term of μ , known as the 'threshold'

value, beyond which data is considered extreme. The value has been universally chosen as 1 m for all processed datasets earlier. It is worth noticing that the value is additive to the HPL computation. Furthermore, the alert level is chosen to be 3 m to characterise the MI, HMI and other operational modes in the IM framework.

The time series for the computed HPL for along-track and cross-track components are plotted in Figure 6.12, 6.13 and 6.14, respectively, for the three processing modes.




Figure 6.12: Time series of component-wise HPL and HPE for Run 4 (PPP-only)




Figure 6.13: Time series of component-wise HPL and HPE for Run 4 (PPP/TC)




Figure 6.14: Time series of component-wise HPL and HPE for Run 4 (PPP/TC/RCM)

As expected by the theory and the selection of threshold values, the HPLs bounds the HPEs. The triple-sensor solution provided the lowest magnitudes of HPL values for bounding its errors, indicating a superior availability once the alert level is considered. The PPP-only solution, contrastingly, has HPL values well over 10 metres to cover the peaks of errors.

With the information available, the integrity performance can be derived for all three solutions. Epochs with comparisons of PE , AL , and PL values can be counted, leading to MI, HMI and system availabilities percentages. The epochs are also plotted in Stanford diagrams as shown in Figure 6.15, 6.16, 6.17 for the PPP-only, PPP/TC and PPP/TC/RCM, respectively.




Figure 6.15: Stanford Diagram for Run 4 (PPP-only)




Figure 6.16: Stanford Diagram for Run 4 (PPP/TC)




Figure 6.17: Stanford Diagram for Run 4 (PPP/TC/RCM)

The Stanford diagrams show that the IM framework behaves as expected, and all solutions were protected without misleading information or hazardously misleading information. The AL can be adjusted by moving along the diagonal, but only impacting the percentage of conservative warnings ($PE < AL < PL$), nominal operations ($PE < PL < AL$), or proper warnings ($AL < PE < PL$). At an identical AL, the triple-sensor

fusion outperformed the traditional solutions with the highest percentage of nominal operational epochs.

The exact statistics are shown in Table 6.2. In terms of cross-track errors, the PPP/TC increased the solution availability by 0.8%, and the triple-sensor solution gave a further 2.3% increase, leading to a total availability of 99.5%. The percentage of epochs where a proper warning is issued, that is, $PL > PE > AL$, is reduced from 1.5% to 0.2%. In terms of along-track errors, it is observed that there is a higher unavailability across the board. Although the two components are treated separately, the study employs a uniform methodology to approach them. Real-world applications often have different requirements for the along-track errors versus cross-track errors, as the former is more commonly seen in scenarios such as forward collision warning rather than lane-keeping. Nonetheless, the statistics show a 9.8% increase in safely operational epochs for the triple-sensor solution versus the PPP-only solution, equivalent to a roughly 11% performance improvement.

Table 6.2: Integrity statistics of PPP-only, PPP/TC, and PPP/TC/RCM solutions in Run 4 specifying the percentage of safe operations, system unavailabilities and misleading epochs

Solution	Component	Safe Operations	System Unavailable (Conservative Warning)	System Unavailable (Proper Warning)	MI/HMI (Misleading)
PPP	Cross-track	96.4%	2.0%	1.5%	0%
	Along-track	82.8%	5.6%	11.5%	0%
PPP/TC	Cross-track	97.2%	2.0%	0.8%	0%
	Along-track	82.3%	5.4%	12.2%	0%
PPP/TC/RCM	Cross-track	99.5%	0.4%	0.2%	0%
	Along-track	92.7%	2.2%	5.1%	0%

6.4 Conclusions from Integrity Analysis

The study adopted a direct-to-position-domain approach by modelling the heavy tails of horizontal position errors (HPEs) with an overbounding Generalised Pareto Distribution (GPD) function that is descriptive of the tail characteristics. The acquired probability density function (PDF) establishes a relation between position errors and probability. A protection radius of integrity, also known as the horizontal protection levels (HPLs) can then be derived. During the process, it was found that the cross-track component and the along-track component showed distinct tail distributions. Thus, the study treated the two components independently. The solution from real-world urban canyons was analysed with the methodology, producing integrity performances for the PPP-only, PPP/TC and PPP/TC/RCM solutions. The IM frameworks, built separately for each solution using identical criteria, correctly produced HPLs that avoided misleading operations. The novel triple-sensor fusion outperformed either of the two traditional methods by producing the highest proportion of nominal epochs, and reduced the epochs with 'proper warning', where a hazardous operation was correctly predicted.

The chapter on integrity monitoring elevated the analysis of the novel triple-sensor fusion with the external clock as the additional sensor. The position statistics from datasets collected in real urban canyons were transformed statistically to derive parameters that describe the solution's integrity performance. The performance of the triple-sensor fusion in the challenging environment has demonstrated an improved resilience under poor signal conditions by providing an improved integrity performance. The novel sensor fusion was able to maintain an accurate and reliable solution, and is shown to be a viable option in the emerging applications.

Chapter 7

Conclusions and Recommendations for Future Research

To address the challenge of enhancing the performance of modern GNSS solutions in urban environments, this research aims to develop a novel sensor fusion strategy that incorporates an external frequency source on top of the low-cost GNSS-PPP/IMU integration. The step-by-step research began with building the higher-end prototype, and concluded by proving solution improvement in accuracy and integrity with a more mass-market-friendly variation. This chapter provides a summary of conclusions from each step of the work, the results and reflections, and provides corresponding recommendations to future research projects.

7.1 Conclusions

The dissertation has proposed a novel and low-cost triple-sensor fusion method based on GNSS-PPP (Precise Point Positioning) technology, an IMU, and an external frequency

source in the form of a precise clock. The tri-sensor fusion is a step forward in closing the gap as it provides an improved and tested solution in real-world urban canyons within a reasonable budget. The dissertation begins with an overview of relevant technologies and theories, followed by studies of the clock-aided GNSS approach, methodologies for field experiments and processing where a more costly prototype and a low-cost evolution are tested, and then analysis of the performance metrics of the setups, including a position-domain integrity framework for real-life considerations.

7.1.1 Conclusions on the Prototype Triple-Sensor Solution

The primary work of the dissertation began with expanding existing CSAC/GNSS-PPP research with a tightly-coupled IMU integration. The two primary clock-aiding algorithms of RCM and RCC in the field were adopted to produce a navigation solution in a suburban setting. The environment provided some natural challenges with obstructions, such as a couple of overpasses, but to validate the performance of the solution, artificial satellite masks were implemented. The additional masks allowed the study to produce periods with a specific number of satellites. In particular, the three-satellite solution without rank deficiency was successfully tested, where both clock-aiding algorithms outperformed the traditional GNSS/IMU combination with less error growth over time.

The overall solution has seen the horizontal rms error reduced from 4.31 m in PPP solution to 1.49 m with tightly-coupled IMU integration, and to 1.14 m with the external CSAC implemented using RCM algorithm. With the three-satellite solution lasting 20 seconds, or 100 epochs, the traditional GNSS/IMU dual-sensor integration presented a horizontal rms of 2.07 m, where the RCM and RCC algorithms reduced the error value to 1.36 m and 0.55 m, respectively. However, the processing also cautioned the limits

of the RCC algorithm, where the performance of the solution is largely dependent on obtaining and identifying a good clock bias estimation prior to GNSS outages, and the ability to revert back to estimating the bias upon recognising a change in environment.

7.1.2 Conclusions on the Low-Cost Triple-Sensor Solution in Real Urban Canyons

The study evolved by adopting the Septentrio Mosaic-T timing GNSS receiver as a low-cost alternative to the NovAtel OEM7 geodetic receiver. Furthermore, the CSAC is replaced with a carefully chosen alternative of an OCXO, providing improved short-term timing stability over typical outage durations. CSACs typically cost a couple of thousand U.S. dollars, which is too expensive for most applications. There had been no previous literature that had attempted to bring a clock-aided GNSS into the low-cost domain. The OCXO is an ideal candidate performance-wise or cost-wise, as it costs usually around a few hundred dollars.

The characteristics of the OCXO allow the clock-aided receiver to be able to track GPS time anytime in its operational period, thanks to its internal GNSS capability. The GNSS capability allows the OCXO to automatically enter a hold-over mode when GNSS signals are lost, or if the GNSS chip fails to consistently calibrate its timing frequency over a specified duration, indicating intermittent or poor signal environments. The adaptivity means the low-cost triple-sensor fusion is most optimal for the RCM algorithm due to the automatic holdover mechanism.

The OCXO-aided GNSS receiver went through an identical process as the CSAC-aided, geodetic GNSS receiver in their respective triple-sensor fusion, and has proved a consistent performance when the number of visible satellites is reduced to 5, 4 and 3, respectively, through masking. In the most challenging 3-satellite scenario, the PP-

P/TC/RCM solution presented an overall rms of 2.46 m, significantly lower than that of the PPP/TC solution with 4.20 m, and 9.07 m without sensor integration. Another novelty regarding this chapter of the work is the test performance in real urban canyons at one of the most bustling streets with high-rises obstructing GNSS signals. The clock-aided solution presented an overall rms less than 2 metres, compared to 2.47 m from the dual-sensor fusion, or 2.95 m from the PPP-only solution. Additionally, this configuration was tested in similar environments for multiple runs, with a mean accuracy improvement of 26%. The solution improvement also justifies the potential engineering trade-off of spending the extra costs on a superior GNSS receiver or an external clock, as geodetic receivers tend to be an order of magnitude more expensive, while suffering from similar degrees of adverse effects from the urban canyons.

The proposed solution has achieved the requirement of ease for deployment: no additional land infrastructures are needed, and thus, it is able to serve as a backup solution in a safety-of-life or liability-critical navigation system. It offers a mass-market viability: the solution must be achieved with mass-market hardware components within a reasonable cost range and can be scaled down during mass production. Its performance has been verified through key metrics of accuracy and integrity, which are both essential for the emerging applications. The performance had further closed the performance gap of low-cost sensor fusions from real-world requirements.

7.1.3 Conclusions on Position-Domain Integrity Analysis

The motivation of the study is strongly associated with the real-life performance criteria demanded by safety-of-life or liability-critical automotive applications where low-cost GNSS receivers are widely used in urban environments. To build a stronger

case of proving the performance of the novel sensor fusion, demonstrating its capability beyond accuracy analysis is essential.

The study adopted a less common approach of integrity monitoring through the position-domain, exploiting the fact that the reference solution can be viewed as the 'truth' in the test runs, a similar approach used in the aviation navigation. Through justifying and evoking the extreme value theory, the tails of kinematic solutions, representing the corner cases where spikes of position error were detected, were modelled. Using 1.0 m as the component-wise threshold, a buffer from the typical 1.5 m accuracy for lane-level navigation, the protection levels for the cross-track errors and the along-track errors were computed for all solutions. Stanford diagrams were plotted for single-, dual-, and triple-sensor schemes in the real-world urban canyon dataset. The result shows the epochs with safe operations was increased by 2.3% in the cross-track direction, and 9.6% in the along-track direction when compared to traditional dual-sensor processing. The period of system unavailability due to warnings was reduced by more than half.

7.1.4 Summary of Research Contribution

In the context of existing and similar research, the proposed solution excels in several aspects. Firstly, the solution requires minimal investment in infrastructure. The PPP-based GNSS solution provides a cm-level baseline solution in ideal environments, and the sensor integration in the user equipment compensates for the loss of performance in challenging environments. The sensors do not require base stations, regularly updated high-resolution maps, or any external maintenance beyond those in existence to support the GNSS constellations. Secondly, existing research in clock-aided GNSS fails to provide a low-cost, practical solution to prove its worth. Moreover, the existing literature did not

explore the improvement of the solution in real urban environments, nor the impact on key performance metrics such as integrity. The research innovatively adopted OCXO in conjunction with other mass-market sensors as a step forward from existing CSAC-based clock-aiding solutions. The real-world urban canyon datasets were used with analysis in accuracy and integrity.

The study presents a comprehensive approach to the concept of a clock-aided GNSS receiver themed around automotive applications in urban canyons. The results and conclusion show there is a practical benefit in adopting an external frequency source as the third sensor in a traditional GNSS/IMU fusion. Such outcome can be used in mass-market V2X, intelligent vehicles, and similar applications that desire a higher accuracy and more trustworthiness in corner cases than existing sensor fusions.

7.1.5 Recommendations for Future Research

The pathway of this research, namely, beginning with building a clock-aided, triple-sensor solution using higher-end hardware, followed by switching to lower-cost alternatives, and finally, justification of the performance in real-life scenarios, has opened up plenty of opportunities for further research through the trade-off decisions along the way. The trade-off decisions are made based on the goal of each step in the research. During each step, new concerns and ideas pop up, but may not be addressed due to the scope of each step. For each of the primary aspects of the study, recommendations are made and detailed below.

7.2 Recommendations for Future Research

The dissertation is composed with theories, illustration of the algorithm, and more importantly, the methodology that developed through time and overcoming challenges to apply RCM and RCC algorithms with the triple-sensor fusions schemes. The analysis of various results presented questions that are reasonably answered, but worth fourth investigation and proof. This section will go through the key topics in the research, examine the founding, and make recommendations for further investigation.

7.2.1 Recommendations on the Sensor Fusion

The study has proposed and implemented a triple-sensor fusion method for the three sensors involved. The IMU acted as the dead-reckoning sensor, providing the information for continuous navigation in a complete GNSS outage. The external clock supports the GNSS by providing an extra source to reduce the required number of independent observations. The external clock enables three-satellite solutions without rank deficiency. There had been room for improvement for the RCC and RCM algorithms and operational improvement for the solution.

To address the frequency drift of either of the external clocks, clock discipline must be performed beforehand. This raises the concerns of (1) how often the clock discipline is required; (2) how much performance degradation is expected for any clock discipline due to; (3) whether the duration or frequency of clock disciplining changes over a course of years as the clock ages. These concerns can be answered in theory with the data at hand. However, it will require extra work, especially with the frequency behaviour monitoring for the clocks. As the clock behaviour becomes better understood and characterised, it is

possible to enhance the RCC/RCM algorithms accordingly.

Similarly, the RCM parameters used in the research are adopted from official datasheets and existing research. The clocks' performance may vary among batches or individuals, and may change with time. Due to the uncertainty, the RCC algorithm was presented alongside the RCM algorithm for earlier solutions. Theoretically, if the clock performance can be characterised precisely, the RCM should provide a better performance than the RCC solution.

The RCC solution is based on the assumption that the clock bias estimation in the initial epoch is low, as it will be carried along into the following epochs. The algorithm adopted in the code handles singular outages well, but encounters problems when there are consecutive outages and the navigation engine has no time to re-converge. Furthermore, as the clock drifts, there is a window of validity for the receiver time bias estimation. This window is longer for the CSAC and shorter for the OCXO. It is possible to improve the existing algorithm to handle consecutive outages better, for example, through 3D mapping or cameras for outage prediction, and balance with the window of validity in the meantime.

7.2.2 Recommendations on Building and Testing the Triple-Sensor Solution

The research began with the classic combination of NovAtel OEM7 GNSS receiver and the Microchip SA.45 series of CSAC. The IMU was an automotive-grade MEMS-IMU at a lower cost range. The GNSS receiver had a number of internal tweaks to utilise the external clock. However, the research did not have access to the internal algorithm used in the receiver to investigate or to improve its coupling with the CSAC. This left a black box, which future research may tap into, potentially creating a deeply coupled

integration approach with the receiver itself. This goes similarly with the Mosaic receiver and GXClock-500 OCXO coupling, where there are fewer configurable options for the researcher to capitalise on the external clock.

The IMUs in the research do not have access to the external clock directly, and their respective output streams are time-synchronised in the processing stage. This fact is due to the hardware limitation. As future research, it is recommended to investigate a deeply coupled solution. In a deeply coupled GNSS/IMU integration, the two sensors are fused at the signal level, meaning that the raw tracking loops of the GNSS receiver will benefit from the IMU, hence improving resilience with poor signals. It is possible to add the external clock in the GNSS/IMU deep coupling and improve the performance of a deeply coupled sensor integration. This also prevents the possible issues from time stamp synchronisation with a series of low-cost hardware.

Finally, the environmental factors may play a role in the operation of a clock-aided sensor. As some of the experiments were conducted in a Canadian winter with sub-zero temperatures, the clocks, regardless of CSAC or OCXO, would have been challenged physically by the environmental temperature. While the effect of cold weather is studied better with high-end clocks, particularly for the specialised model also intended for space or submarine uses, low-cost and mass-market models often do not have much information available. Studying the impact on the clock's behaviour, and hence the navigation performance, will be helpful in commercialising the clock-aided GNSS solution.

7.2.3 Recommendations on Evaluating the Solution and Integrity Monitoring

The penultimate chapter introduced a method to evaluate the solution's integrity performance by borrowing a case from proving SBAS in aviation. The position-domain

IM has the advantages that are particularly suitable for this study, as it requires fewer observations and less past data to work with. However, to rigorously reinforce the argument, it is plausible to take additional data in a variety of environments, grouping the data to build a more concrete snapshot of the integrity performance.

The study used a threshold value of 1 m for both along-track and cross-track errors. Real-world expectations vary from specific manoeuvre requirements and regulations. It is possible to extend the current study with specific safety-of-life scenarios identified by research and regulations. The along-track and cross-track errors will be customised accordingly.

Based on the current performance of the triple-sensor solution, its performance in real urban canyons, despite the improvement over traditional setups, has yet to reach a consistent performance of 0.5 m or less identified in some applications. Further studies can continue working on the GNSS-PPP aspects in making ambiguity resolution more possible and reliable in challenging environments with low-cost equipment, and continue to improve the baseline solution through further GNSS augmentations.

References

- Abd Rabbou, M., & El-Rabbany, A. (2015). Tightly-coupled integration of GPS precise point positioning and MEMS-based inertial systems. *GPS Solutions*, 19, 601–609.
- Aggrey, J., & Bisnath, S. (2019). Improving GNSS PPP convergence: The case of atmospheric-constrained, multi-GNSS PPP-AR. *Sensors*, 19(3), 587.
- Aggrey, J., Bisnath, S., Naciri, N., Shinghal, G., & Yang, S. (2019). Accuracy trend analysis of low-cost GNSS chips: The case of multi-constellation GNSS PPP. *Proceedings of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*, 3618–3635.
- Aggrey, J. E. (2014). *Multi-GNSS Precise Point Positioning Software Architecture And Analysis Of GLONASS Pseudorange Biases* [Doctoral dissertation, York University].
- Agnew, D. C. (1997). NLOADF: A program for computing ocean-tide loading. *Journal of geophysical research: solid earth*, 102(B3), 5109–5110.
- Ahmad, K., Sahmoudi, M., & Macabiau, C. (2014). Trusted overbounding position errors for gnss positioning in urban environments. *Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*, 3618–3628.
- Ahmed, H., & Tahir, M. (2016). Accurate attitude estimation of a moving land vehicle using low-cost MEMS IMU sensors. *IEEE Transactions on Intelligent Transportation Systems*, 18(7), 1723–1739.
- Ali, M. (2016). Compensation of temperature and acceleration effects on mems gyroscope. *2016 13th International Bhurban Conference on Applied Sciences and Technology (IBCAST)*, 274–279.
- Allan, D. W. (2005). Statistics of atomic frequency standards. *Proceedings of the IEEE*, 54(2), 221–230.

- Altamimi, Z., Métivier, L., & Collilieux, X. (2012). aITRF2008 plate motion model. *Journal of geophysical research: solid earth*, 117(B7).
- Ashby, N. (2003). Relativity in the global positioning system. *Living Reviews in relativity*, 6(1), 1–42.
- Azais, J.-M., Gadat, S., Levy, J.-C., Rols, B., Mercadier, C., Jordan, C., & Suard, N. (2009). GNSS integrity achievement by using extreme value theory. *Proceedings of the 22nd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2009)*, 1281–1287.
- Barnes, J. A., & Allan, D. W. (1966). A statistical model of flicker noise. *Proceedings of the IEEE*, 54(2), 176–178.
- Beard, R., & Senior, K. (2017). Clocks. In P. J. Teunissen & O. Montenbruck (Eds.), *Springer handbook of global navigation satellite systems* (pp. 121–164). Springer International Publishing.
- Bisnath, S., & Gao, Y. (2009). Current state of precise point positioning and future prospects and limitations. In M. G. Sideris (Ed.), *Observing our changing earth* (pp. 615–623). Springer Berlin Heidelberg.
- Blair, B. E. (1974). *Time and frequency: Theory and fundamentals* (Vol. 140). US National Bureau of Standards.
- Blanch, J., Walter, T., Enge, P., Lee, Y., Pervan, B., Rippl, M., & Spletter, A. (2012). Advanced RAIM user algorithm description: Integrity support message processing, fault detection, exclusion, and protection level calculation. *Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012)*, 2828–2849.
- Braff, R. (2003). A method for laas fault-free error overbounding using a position domain monitor. *Proceedings of the 2003 National Technical Meeting of The Institute of Navigation*, 326–338.
- Brannon, A., Jankovic, M., Breitbarth, J., Popovic, Z., Gerginov, V., Shah, V., Knappe, S., Hollberg, L., & Kitching, J. (2006). A local oscillator for chip-scale atomic clocks at nist. *2006 IEEE International Frequency Control Symposium and Exposition*, 443–447.
- Brown, K. R. (1991). The theory of the GPS composite clock.
- Brown, R. G. (1992). A baseline GPS RAIM scheme and a note on the equivalence of three RAIM methods. *Navigation*, 39(3), 301–316.
- Brown, R. G., & Hwang, P. Y. (1997). *Introduction to random signals and applied Kalman filtering: With MATLAB exercises and solutions*. John Wiley & Sons, 2012.

- Choy, S., Bisnath, S., & Rizos, C. (2017). Uncovering common misconceptions in GNSS precise point positioning and its future prospect. *GPS Solutions*, 21, 13–22.
- Dey, D. K., & Yan, J. (2016). *Extreme value modeling and risk analysis: Methods and applications*. CRC Press.
- Dierendonck, A. J. V., & McGraw, J. (1984). Relationship Between Allan Variances and Kalman Filter Parameters. *Proceedings of the 16th Annual Precise Time and Time Interval Systems and Applications Meeting*, 273–293.
- Domínguez, E., Moriana, C., Bonardi, L., Aguado, E., Lowe, D., Pattinson, M., Hutchinson, M., Seco-Granados, G., Lopez-Salcedo, J., Egea, D., et al. (2015). Vehicular and pedestrian GNSS integrity algorithms and results for urban and road environments developed after an extensive real data collection campaign. *Proceedings of the 28th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2015)*, 553–568.
- El-Mowafy, A., & Kubo, N. (2017). Integrity monitoring of vehicle positioning in urban environment using RTK-GNSS, IMU and speedometer. *Measurement Science and Technology*, 28(5), 055102.
- El-Rabbany, A., & El-Diasty, M. (2004). An efficient neural network model for de-noising of MEMS-based inertial data. *the Journal of Navigation*, 57(3), 407–415.
- Elsheikh, M., Abdelfatah, W., Noureldin, A., Iqbal, U., & Korenberg, M. (2019). Low-cost real-time PPP/INS integration for automated land vehicles. *Sensors*, 19(22), 4896.
- Elsheikh, M., Noureldin, A., & Korenberg, M. (2020). Integration of GNSS precise point positioning and reduced inertial sensor system for lane-level car navigation. *IEEE Transactions on Intelligent Transportation Systems*, 23(3), 2246–2261.
- Elsheikh, M., Iqbal, U., Noureldin, A., & Korenberg, M. (2023). The implementation of precise point positioning (PPP): A comprehensive review. *Sensors*, 23(21), 8874.
- El-Sheimy, N., & Niu, X. (2007). The promise of mems to the navigation community. *Inside GNSS*, 2(2), 46–56.
- Farrell, J., & Barth, M. (1999). *The global positioning system and inertial navigation*. McGraw-Hill.

- Feng, S., Ochieng, W., Samson, J., Tossaint, M., Hernandez-Pajares, M., Juan, J. M., Sanz, J., Aragón-Àngel, À., Ramos-Bosch, P., & Jofre, M. (2012). Integrity monitoring for carrier phase ambiguities. *The Journal of Navigation*, 65(1), 41–58.
- Fernandez, E., & Calero, D. (2015). Characterization of chip-scale atomic clock for gnss navigation solutions. *2015 International Association of Institutes of Navigation World Congress (IAIN)*, 1–8.
- Fernandez, E., Calero, D., & Eulia Pares, M. (2016). CSAC temperature calibration for improving GNSS positioning performance. *Proceedings of the 47th Annual Precise Time and Time Interval Systems and Applications Meeting*, 36–44.
- Fernandez, E., Calero, D., & Eulia Pares, M. (2017). CSAC characterization and its impact on GNSS clock augmentation performance. *Sensors*, 17, 370.
- Fernández, E., Calero, D., & Parés, M. (2017). CSAC characterization and its impact on GNSS clock augmentation performance. *Sensors*, 17(2), 370.
- Frerking, M. E. (1996). Fifty years of progress in quartz crystal frequency standards. *Proceedings of 1996 IEEE International Frequency Control Symposium*, 33–46.
- Gebre-Egziabher, D. (2018). Managing uncertainty in the design of safety-critical aviation systems-safety critical unmanned aerial systems, 2297–2320.
- Gelb, A. (1974). Applied optimal estimation.
- Geng, J., Zeng, R., & Guo, J. (2024). Assessing all-frequency GPS/Galileo/BDS PPP-RTK in GNSS challenging environments. *GPS Solutions*, 28(1), 5.
- Gratton, L., Joerger, M., & Pervan, B. (2010). Carrier phase relative raim algorithms and protection level derivation. *The Journal of Navigation*, 63(2), 215–231.
- Groves, P. (2013, March). *Principles of GNSS, inertial, and multisensor integrated navigation systems, second edition*.
- Groves, P., Jiang, Z., Rudi, M., & Strode, P. (2013). A portfolio approach to NLOS and multipath mitigation in dense urban areas. *26th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2013*, 4.
- Groves, P. D., Jiang, Z., Wang, L., & Ziebart, M. K. (2012). Intelligent urban positioning using multi-constellation GNSS with 3D mapping and NLOS signal detection. *Proceedings of the 25th interna-*

- tional technical meeting of the satellite division of the Institute of Navigation (ION GNSS 2012)*, 458–472.
- Gurtner, W., & Estey, L. (2007). RINEX-the receiver independent exchange format-version 3.00. *Astronomical Institute, University of Bern and UNAVCO, Boulder, Colorado*.
- Haan, L., & Ferreira, A. (2006). *Extreme value theory: An introduction* (Vol. 3). Springer.
- Hegarty, C., Leva, J., Dyke, K. V., & Walter, T. (2017). Performance of stand-alone GNSS. In E. Kaplan & C. J. Hegarty (Eds.), *Understanding GPS/GNSS: Principles and applications, third edition* (3rd). Artech House, Inc.
- Hilla, S. (2016). The extended standard product 3 orbit format (SP3-d) 21 february 2016.
- Ho, C., Wilson, B., Mannucci, A., Lindqwister, U., & Yuan, D. (1997). A comparative study of ionospheric total electron content measurements using global ionospheric maps of GPS, TOPEX radar, and the Bent model. *Radio Science*, 32(4), 1499–1512.
- Hopfield, H. (1969). Two-quartic tropospheric refractivity profile for correcting satellite data. *Journal of Geophysical research*, 74(18), 4487–4499.
- Jee, G., Lee, H.-B., Kim, Y., Chung, J.-K., & Cho, J. (2010). Assessment of GPS global ionosphere maps (GIM) by comparison between CODE GIM and TOPEX/Jason TEC data: Ionospheric perspective. *Journal of Geophysical Research: Space Physics*, 115(A10).
- Joerger, M., Chan, F.-C., & Pervan, B. (2014). Solution separation versus residual-based RAIM. *Navigation: Journal of the Institute of Navigation*, 61(4), 273–291.
- Kaniuth, R., Lemke, N., Heinrichs, G., IfEN, J. Ó. W., Ehm, H., Weigel, R., Schmid, A., Neubauer, A., & Nagel, P. (2005). A snap-shot positioning approach for a high integrated GPS/Galileo chipset, 7.
- Kannemans, H. (2010). The generalized extreme value statistical method to determine the GNSS integrity performance. *2010 5th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*, 1–8.
- Kaplan, E., & Hegarty, C. J. (2017). *Understanding GPS/GNSS: Principles and applications, third edition* (3rd). Artech House, Inc.
- Katsigianni, G., Loyer, S., & Perosanz, F. (2019). PPP and PPP-AR kinematic post-processed performance of GPS-only, Galileo-only and multi-GNSS. *Remote sensing*, 11(21), 2477.

- Kitching, J. (2018). Chip-scale atomic devices. *Applied Physics Reviews*, 5(3).
- Klobuchar, J. A. (1975). *A first-order, worldwide, ionospheric, time-delay algorithm* (Vol. 75). Air Force Cambridge Research Laboratories, Air Force Systems Command, United ...
- Knabe, N., & Kalafus, R. M. (1984). Clock coasting and altimeter error analysis for GPS. *Navigation*, 31(4), 289–302.
- Knappe, S., Schwindt, P., Shah, V., Hollberg, L., Kitching, J., Liew, L., & Moreland, J. (2005). A chip-scale atomic clock based on 87Rb with improved frequency stability. *Optics express*, 13(4), 1249–1253.
- Kouba, J. (2009). A guide to using International GNSS Service (IGS) products.
- Kouba, J., & Héroux, P. (2001). Precise point positioning using IGS orbit and clock products. *GPS Solutions*, 5, 12–28.
- Krawinkel, T. (2018). *Improved GNSS navigation with chip-scale atomic clock* [Doctoral dissertation, Gottfried Wilhelm Leibniz Universität Hannover].
- Krawinkel, T., Schon, S., & Bauch, A. (2014). Application of miniaturized atomic clocks in kinematic GNSS single point positioning. *2014 European Frequency and Time Forum (EFTF)*, 97–100.
- Kusters, J. A., & Vig, J. R. (1991). Hysteresis in quartz resonators - a review. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 38(3), 281–290.
- Lam, C. (2008). A review of the recent development of MEMS and crystal oscillators and their impacts on the frequency control products industry. *2008 IEEE Ultrasonics Symposium*, 694–704.
- Larson, J. D., Gebre-Egziabher, D., & Rife, J. H. (2019). Improving navigation system availability using gaussian-pareto overbounding. *Proceedings of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*, 3153–3161.
- Lee, Y. C. (1986). Analysis of range and position comparison methods as a means to provide GPS integrity in the user receiver. *Proceedings of the 42nd Annual Meeting of The Institute of Navigation (1986)*, 1–4.
- Lee, Y. C. (2004). Investigation of extending receiver autonomous integrity monitoring (RAIM) to combined use of Galileo and modernized GPS. *Proceedings of the 17th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2004)*, 1691–1698.
- Leschiutta, S. (2005). The definition of the ‘atomic’ second. *Metrologia*, 42(3), S10.

- Lewandowski, W., & Arias, E. (2011). GNSS times and UTC. *Metrologia*, 48(4), S219–S224.
- Li, L., Li, R., Wang, L., Wang, R., Li, M., & Li, M. (2023). GNSS integrity risk evaluation in the position domain based on the generalized Pareto distribution. *Measurement Science and Technology*, 34(9), 095010.
- Li, T., Zhang, H., Niu, X., & Gao, Z. (2017). Tightly-coupled integration of multi-GNSS single-frequency RTK and MEMS-IMU for enhanced positioning performance. *Sensors*, 17(11), 2462.
- Li, Y., Niu, X., Zhang, Q., Cheng, Y., & Shi, C. (2012). Observability analysis of non-holonomic constraints for land-vehicle navigation systems. *Proceedings of the 25th International technical meeting of the satellite division of the institute of navigation (ION GNSS 2012)*, 1521–1529.
- Lim, J., Kim, H., Jackson, T. N., Choi, K., & Kenny, D. (2010). An ultra-compact and low-power oven-controlled crystal oscillator design for precision timing applications. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 57(9), 1906–1914.
- Liu, T., Zhang, B., Yuan, Y., Li, Z., & Wang, N. (2019). Multi-GNSS triple-frequency differential code bias (DCB) determination with precise point positioning (PPP). *Journal of Geodesy*, 93, 765–784.
- Lombardi, M. A. (2008). The use of GPS disciplined oscillators as primary frequency standards for calibration and metrology laboratories. *NCSLI Measure*, 3(3), 56–65.
- Lombardi, M. A. (2016). Evaluating the frequency and time uncertainty of GPS disciplined oscillators and clocks. *NCSLI Measure*, 11(3-4), 30–44.
- Lutwak, R., Emmons, D., Riley, W., & Garvey, R. M. (2002a). The chip-scale atomic clock – coherent population trapping vs. conventional interrogation, 12.
- Lutwak, R., Emmons, D., Riley, w., & Garvey, R. (2002b). The chip-scale atomic clock – coherent population trapping vs. conventional interrogation. *Proceedings of the 34th Annual Precise Time and Time Interval Systems and Applications Meeting*.
- Lutwak, R., Emmons, D., Riley, w., & Garvey, R. (2004). The chip-scale atomic clock – low-power physics package. *Proceedings of the 36th Annual Precise Time and Time Interval Systems and Applications Meeting*.
- Ma, L., You, Z., Li, B., Zhou, B., & Han, R. (2015). Deep coupled integration of CSAC and GNSS for robust PNT. *Sensors*, 15(9), 23050–23070.

- Ma, L., You, Z., Liu, T., & Shi, S. (2016). Coupled integration of CSAC, MIMU, and GNSS for improved PNT performance. *Sensors*, 16(5), 682.
- Matthews, A. (1990). Utilization of fiber optic gyros in inertial measurement units. *Navigation*, 37(1), 17–38.
- McCarthy, D. D., Petit, G., et al. (2004). *IERS conventions (2003)* (Vol. 32). Bundesamt für Kartographie und Geodäsie Frankfurt am Main, Germany.
- Melendez-Pastor, C., Ruiz-Gonzalez, R., & Gomez-Gil, J. (2017). A data fusion system of gnss data and on-vehicle sensors data for improving car positioning precision in urban environments. *Expert systems with applications*, 80, 28–38.
- Meng, Y., Weng, D., Yang, C., Chen, W., & Hou, Z. (2024). Chip-scale atomic clock (CSAC) aided GNSS in urban canyons. *GPS Solutions*, 28(4), 195.
- Microsemi. (2014). White paper - shock and vibration testing of the SA.45s chip scale atomic clock (CSAC) validation units.
- Microsemi. (2016). White paper - leading-edge technology enables a chip scale atomic clock.
- Microsemi. (2017a). User guide - SA.45s chip-scale atomic clock.
- Microsemi. (2017b). White paper - thermally-improved microsemi chip scale atomic clock (CSAC).
- Microsemi. (2018). White paper - chip-scale atomic clock (csac) performance during rapid temperature change.
- Misra, P. (1996). The role of the clock in a GPS receiver. *GPS World*, 60–66.
- Mogyorósi, E., Revisnyei, P., Pašić, A., Papp, Z., Törös, I., Varga, P., & Pašić, A. (2022). Positioning in 5G and 6G networks — a survey. *Sensors*, 22(13), 4757.
- Montenbruck, O., Hauschild, A., & Steigenberger, P. (2014). Differential code bias estimation using multi-GNSS observations and global ionosphere maps. *Navigation*, 61(3), 191–201.
- Naciri, N. (2022). RTK-quality positioning with global PPP corrections. *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*, 2546–2562.

- Nikiforov, I., & Roturier, B. (2005). Advanced RAIM algorithms: First results. *Proceedings of the 18th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2005)*, 1789–1800.
- Niu, Y., Li, M., Wei, N., Shi, C., Chen, G., & Wang, L. (2021). The ocean pole tide loading and its effect on GPS position time-series. *Geophysical Journal International*, 227(1), 368–382.
- NovAtel. (2019). User guide - OEM7 GNSS receivers.
- Ober, P. B., Imparato, D., Verhagen, S., Tiberius, C., Veerman, H., Van Kleef, A., Wokke, F., Bos, A., & Mieremet, A. (2014). Empirical integrity verification of GNSS and SBAS based on the extreme value theory. *Navigation*, 61(1), 23–38.
- Ochieng, W. Y., Feng, S., Moore, T., Hill, C., & Hide, C. (2010). User level integrity monitoring and quality control for high accuracy positioning using GPS/INS measurements. *Positioning (POS) Journal Information*, 14.
- Perul, J., & Renaudin, V. (n.d.). Fusion of attitude and statistical walking direction estimations with time-difference carrier phase update for pedestrian dead reckoning, 23.
- Petovello, M. G. (2003). Real-time integration of a tactical-grade IMU and GPS for high-accuracy positioning and navigation.
- Pickands III, J. (1975). Statistical inference using extreme order statistics. *the Annals of Statistics*, 119–131.
- Prange, L., Arnold, D., Dach, R., Brockmann, E., Kalarus, M. S., Schär, S., Stebler, P., & Jäggi, A. (2023). CODE product series for the IGS-MGEX project.
- Qin, F., Zhan, X., & Zhan, L. (2014). Performance assessment of a low-cost inertial measurement unit based ultra-tight global navigation satellite system/inertial navigation system integration for high dynamic applications. *IET Radar, Sonar & Navigation*, 8(7), 828–836.
- Ramlall, R., Streter, J., & Schneck, J. F. (2011). Three satellite navigation in an urban canyon using a chip-scale atomic clock. *Proceedings of the 24th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2011)*, 4, 2937–2945.
- Reddybattula, K. D., Panda, S. K., Ansari, K., & Peddi, V. S. R. (2019). Analysis of ionospheric TEC from GPS, GIM and global ionosphere models during moderate, strong, and extreme geomagnetic storms over Indian region. *Acta Astronautica*, 161, 283–292.

- Rife, J., Walter, T., & Blanch, J. (2004). Overbounding SBAS and GBAS error distributions with excess-mass functions. *Proceedings of the International Symposium on GNSS/GPS, Sydney, Australia*, 6–8.
- Rife, J., & Gebre-Egziabher, D. (2007). Symmetric overbounding of correlated errors. *Navigation*, 54(2), 109–124.
- Rizos, C. (2005). Trends in geopositioning for lbs, navigation and mapping. *Proceedings of Int. Symp. & Exhibition on Geoinformation*, 27–29.
- Rutman, J., & Walls, F. L. (1991). Characterization of frequency stability in precision frequency sources. *Proceedings of the IEEE*, 79(7), 952–960.
- SAE International. (2021). Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles [SAE Standard J3016_202104, Revised April 2021].
- Safran. (2022). User manual - GXClok-500 auto-adaptive GPS/GNSS receiver and crystal oscillator module.
- Salós, D., Martineau, A., Macabiau, C., Kubrak, D., & Bonhoure, B. (2010). Groundwork for GNSS integrity monitoring in urban road applications. the road user charging case. *Proceedings of the 23rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2010)*, 1130–1144.
- Schmid, R., Dach, R., Collilieux, X., Jäggi, A., Schmitz, M., & Dilssner, F. (2016). Absolute IGS antenna phase center model igs08.atx: Status and potential improvements. *Journal of Geodesy*, 90, 343–364.
- Shin, E.-H., & El-Sheimy, N. (2002). Accuracy improvement of low cost INS/GPS for land applications. *Proceedings of the 2002 national technical meeting of the institute of navigation*, 146–157.
- Simon, J., Vazquez, J., Madrazo, A., Cezon, A., & Enderle, W. (2010). Advanced RAIM scheme and its evaluation for non-aeronautical users in different environments. *Proceedings of the 23rd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2010)*, 292–300.
- Skaloud, J., Bruton, A., & Schwarz, K.-P. (1999). Detection and filtering of short-term (1/f γ) noise in inertial sensors. *Navigation*, 46(2), 97–107.
- Spinney, J. E. (2003). Mobile positioning and lbs applications. *Geography*, 88(4), 256–265.
- Standardization Administration of the People's Republic of China. (2024). GB/T 45086.1-2024: Technical Requirements and Test Methods for On-Board Positioning System – Part 1: Satellite Positioning [Issued on November 28, 2024; Implemented on June 1, 2025].

- Sturza, M. A. (1983). GPS navigation using three satellites and a precise clock. *Navigation*, 30(2), 146–156.
- Teunissen, P., & Montenbruck, O. (2017). Handbook of global navigation satellite systems. Springer.
- Tiberius, C., & Verbree, E. (2004). GNSS positioning accuracy and availability within Location Based Services: The advantages of combined GPS-Galileo positioning. *2nd ESA/Estec workshop on Satellite Navigation User Equipment Technologies, GS Granados (Ed), ESA publications division, Noordwijk*, 1–12.
- Tossaint, M., Samson, J., Toran, F., Ventura-Traveset, J., Hernández-Pajares, M., Juan, J., Sanz, J., & Ramos-Bosch, P. (2007). The Stanford–ESA integrity diagram: A new tool for the user domain sbas integrity assessment. *Navigation*, 54(2), 153–162.
- Tunalı, E., & Özlüdemir, M. T. (2019). GNSS PPP with different troposphere models during severe weather conditions. *GPS Solutions*, 23(3), 82.
- Uhlemann, M., Gendt, G., Ramatschi, M., & Deng, Z. (2015). GFZ global multi-GNSS network and data processing results. In *Iag 150 years: Proceedings of the iag scientific assembly in postdam, germany, 2013* (pp. 673–679). Springer.
- U.S. Department of Transportation. (2025, January). *Positioning, navigation, and timing strategic plan* (tech. rep.).
- Vana, S., Naciri, N., & Bisnath, S. (2020). Benefits of motion constraining for robust, low-cost, dual-frequency GNSS PPP + MEMS IMU navigation. *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 1093–1103.
- Vana, S., & Bisnath, S. (2023). Low-cost, triple-frequency, multi-GNSS PPP and MEMS IMU integration for continuous navigation in simulated urban environments. *Navigation*, 70(2).
- Vig, J. R. (1992). Introduction to quartz frequency standards. *Army Research Laboratory Electronics and Power Sources Directorate, Tech. Rep. SLCETTR-92-1*.
- Walter, T. (2002). Characterizing frequency stability: A continuous power-law model with discrete sampling. *IEEE Transactions on Instrumentation and Measurement*, 43(1), 69–79.
- Wang, L., Groves, P. D., & Ziebart, M. K. (2012). Multi-constellation GNSS performance evaluation for urban canyons using large virtual reality city models. *The Journal of Navigation*, 65(3), 459–476.

- Wang, Y., Liu, P., Liu, Q., Adeel, M., Qian, J., Jin, X., & Ying, R. (2019). Urban environment recognition based on the GNSS signal characteristics. *Navigation*, 66(1), 211–225.
- Witchayangkoon, B. (2000). *Elements of GPS precise point positioning* [Doctoral dissertation, The University of Maine].
- Wojciechowski, K. E., & Olsson, R. H. (2015). A fully integrated oven controlled microelectromechanical oscillator—part ii: Characterization and measurement. *Journal of Microelectromechanical Systems*, 24(6), 1795–1802.
- Wu, J.-T., Wu, S. C., Hajj, G. A., Bertiger, W. I., & Lichten, S. M. (1993). Effects of antenna orientation on GPS carrier phase. *Manuscripta geodetica*, 18(2), 91–98.
- Xiang, Y., Gao, Y., & Li, Y. (2020). Reducing convergence time of precise point positioning with ionospheric constraints and receiver differential code bias modeling. *Journal of Geodesy*, 94, 1–13.
- Yang, S., & Bisnath, S. (2020). Gns pp navigation in obstructed environments using clock coasting. *Proceedings of the 2020 International Technical Meeting of The Institute of Navigation*, 521–537.
- Yang, S., Yi, D., Vana, S., & Bisnath, S. (2023). Resilient smartphone positioning using native sensors and PPP augmentation. *Navigation*, 70(2).
- Yi, D., Yang, S., & Bisnath, S. (2022). Native smartphone single-and dual-frequency GNSS-PPP/IMU solution in real-world driving scenarios. *Remote Sensing*, 14(14), 3286.
- Yuksel, Y., El-Sheimy, N., & Noureldin, A. (2010). Error modeling and characterization of environmental effects for low cost inertial mems units. *IEEE/ION Position, Location and Navigation Symposium*, 598–612.
- Zhu, K., Guo, X., Jiang, C., Xue, Y., Li, Y., Han, L., & Chen, Y. (2020). MIMU/odometer fusion with state constraints for vehicle positioning during BeiDou signal outage: Testing and results. *Sensors*, 20(8), 2302.
- Zhu, N., Marais, J., Betaille, D., & Berbineau, M. (2018). Gns position integrity in urban environments: A review of literature. *IEEE Transactions on Intelligent Transportation Systems*, 19(9), 2762–2778.
- Zumberge, J. F., Heflin, M. B., Jefferson, D. C., Watkins, M. M., & Webb, F. H. (1997). Precise point positioning for the efficient and robust analysis of GPS data from large networks. *Journal of Geophysical Research: Solid Earth*, 102(B3), 5005–5017.

Zwahlen, P., Dong, Y., Nguyen, A., Rudolf, F., Stauffer, J., Ullah, P., & Ragot, V. (2012). Breakthrough in high performance inertial navigation grade sigma-delta mems accelerometer. *Proceedings of the 2012 IEEE/ION Position, Location and Navigation Symposium*, 15–19.