

SELF-CALIBRATING PSEUDOLITE ARRAYS:
THEORY AND EXPERIMENT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Edward Alan LeMaster

May 2002

Copyright © 2002 by Edward Alan LeMaster
All Rights Reserved.

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Stephen M. Rock
(Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Per K. Enge

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Robert H. Cannon, Jr.

Approved for the University Committee on Graduate Studies:

Abstract

Tasks envisioned for future-generation Mars rovers — sample collection, area survey, resource mining, habitat construction, etc. — will require greatly enhanced navigational capabilities over those possessed by the 1997 Mars Sojourner rover. Many of these tasks will involve cooperative efforts by multiple rovers and other agents, necessitating both high accuracy and the ability to share navigation information among different users. On Earth, satellite-based carrier-phase differential GPS provides a means of delivering centimeter-level, drift-free positioning to multiple users in contact with a reference base station. It would be highly desirable to have a similar navigational capability for use in Mars exploration.

This research has originated a new local-area navigation system — a Self-Calibrating Pseudolite Array (SCPA) — that can provide centimeter-level localization to multiple rovers by utilizing GPS-based pseudolite transceivers deployed in a ground-based array. Such a system of localized beacons can replace or augment a system based on orbiting satellite transmitters. Previous pseudolite arrays have relied upon *á priori* information to survey the locations of the pseudolites, which must be accurately known to enable navigation within the array. In contrast, an SCPA does not rely upon other measurement sources to determine these pseudolite locations. This independence is a key requirement for autonomous deployment on Mars, and is accomplished through the use of GPS transceivers containing both transmit and receive components and through algorithms that utilize limited motion of a transceiver-bearing rover to determine the locations of the stationary transceivers.

This dissertation describes the theory and operation of GPS transceivers, and how they can be used for navigation within a Self-Calibrating Pseudolite Array. It presents new algorithms that can be used to self-survey such arrays robustly using no *á priori* information, even under adverse conditions such as high-multipath environments. It then describes the experimental SCPA prototype developed at Stanford University and used in conjunction

with the K9 Mars rover operated by NASA Ames Research Center. Using this experimental system, it provides experimental validation of both successful positioning using GPS transceivers and full calibration of an SCPA following deployment in an unknown configuration.

Acknowledgments

This dissertation would not have been possible without the assistance and encouragement of a great number of people. Foremost, I would like to thank my advisor Professor Stephen Rock for giving me the freedom to pursue what may have seemed like a crazy idea at the start, and for giving me feedback on my work which always seemed to cut to the heart of the issue. I also owe a special debt to Professor Robert Cannon for creating the Aerospace Robotics Lab and for fostering an environment where the students are encouraged to explore and create, and are mandated to have fun while doing so. Besides Professors Rock and Cannon, I would like to extend my appreciation to the other members of my reading and oral defense committees, specifically Professors Per Enge, Claire Tomlin, and Bill Newsome.

This project has been very generously funded through several different NASA grants. Personal funding for myself was also provided by the U.S. Department of Defense through an NDSEG Fellowship and by the Mustard Seed Foundation through a Harvey Fellowship. I would like to thank these organizations for their support and for making this research possible.

Although much of the specific work on this project was an individual endeavor, at two key times I was directly aided to a significant extent by other students. At the beginning, Melvin Ni was instrumental in building a significant set of the hardware necessary to get the project rolling; specifically our first rover prototype. The quality of his work is evident by the number of other research groups that have since copied his cafeteria-tray electronics mounting system. More recently, Masayoshi Matsuoka was of invaluable help during the final phase of field testing of the experimental system. I benefited greatly not only from his hard work and diligence under less-than-favorable operating conditions, but also from our many conversations and the innovative ideas he has expressed during our time working together.

While they did not directly work on this project, several people had extremely critical

roles in helping me to formulate my ideas and to develop the SCPA. First and foremost is Eric Olsen, who pointed me in the direction of using pseudolite transceivers for improved state observability. Eric also performed the key modifications of the Orion receiver to allow for pseudolite tracking. His insights and freely-given advice saved me great amounts of time, and without them this thesis would not have been possible. Jonathan Stone was also especially helpful, and greatly aided my understanding of the basics of pseudolites, GPS receivers, and transceivers. Franz Busse took the lead in the building of the Orion receivers themselves for both of our respective projects; it was a great pleasure working with him. Finally Chan-Woo Park provided very useful insights during our discussions of transceiver operations and the near/far problem.

I would like to thank my collaborators at NASA Ames Research Center for their generous donation of testing time aboard the K9 rover, and also for their extreme patience during the many experiments — many of them in scorching heat — in which my system was not operating as expected. I was always surprised at their willingness to head out for yet another test, even when they had nothing apparent to gain for themselves. Instrumental on this team were Maria Bualat, Mike Fair, and Anne Wright.

I was aided at Stanford by countless staff members: Several of these outstanding individuals deserve special note. Jane Lintott was always happy to go out of her way and deal with all of the ugly little details of ordering parts, booking travel arrangements, and wading through the resident bureaucracy so that I was able to focus on the research itself. I greatly appreciate both her help and her cheerful disposition. Godwin Zhang also contributed significantly, both in terms of his electronics knowledge and through the actual components that he constructed for the experimental system. Finally, Aldo Rossi went through countless iterations while machining parts for my transceivers, always with a smile and together with a great deal of practical advice for life itself.

My colleagues in the Aerospace Robotics Lab deserve many thanks, both for their ongoing support in technical matters and also for their continuing friendship. The HUMMING-BIRD autonomous helicopter group — Bruce Woodley, Eric Frew, and Hank Jones — was especially important to me. My time with them was the highlight of my Stanford career, and showed me just how much fun and how rewarding research can be. In addition, I would be amiss if I did not mention the members of the Stanford GPS lab, who have answered many basic questions, provided me with much advice, and fed me a great deal of pizza at their meetings.

I would also like to thank the many people who had no direct influence upon my research, but who greatly influenced my time at Stanford in that most important of ways: friendship. To all of my friends at IVGrad and YAF, I would like to thank you for pushing me and never letting me forget both God's perspective and His love, and for being of unfailing support in both the good times and the bad. To the disc-at-loki ultimate frisbee crew, I thank you all for running me into the ground and otherwise keeping me in shape. And to everyone from Danse Libre, SBDC, Jammix, and the other dance events I have been frequenting, thank you for keeping me sane, and for helping me to discover one of my true joys in life.

Finally, I would like to thank my parents Hanna and Dick, and my sister Michelle, for their unfailing support and encouragement. They never ceased believing in me, even when I did not believe in myself. This dissertation is in no small way a direct product of their love and continual investment in my life

Contents

Abstract	v
Acknowledgments	vii
List of Tables	xv
List of Figures	xvii
1 Introduction	1
1.1 Project Motivation	2
1.2 Mars Navigation Options	3
1.3 Self-Calibrating Pseudolite Arrays	5
1.4 Related Research	9
1.4.1 GPS Pseudolites	10
1.4.2 GPS Transceivers	10
1.4.3 Transceiver Arrays	11
1.5 Contributions	11
1.6 Dissertation Outline	14
2 GPS Transceivers	15
2.1 Transceiver Performance Tradeoffs	15
2.2 Transceiver Architectures	17
2.2.1 Synchrolite	17
2.2.2 Self-Differencing Transceiver	18
2.3 The Near/Far Problem	19
2.3.1 GPS Signal Characteristics	20

2.3.2	Near/Far Solution Methods	23
2.4	Pseudolite Pulsing	25
2.4.1	Synchronization	27
2.4.2	Pulse Duration	27
2.4.3	Pulsing Schemes	28
2.4.4	Power Level	29
2.4.5	Automatic Gain Control Effects	29
2.4.6	Interference Effects	30
2.4.7	Sidelobe Acquisition Effects	34
2.5	Summary of Near/Far Mitigation Techniques	39
3	Transceiver-Based Positioning	43
3.1	Conventional Differential GPS	44
3.1.1	Satellite-Based	44
3.1.2	Pseudolite Based	49
3.2	Bidirectional Ranging	51
3.2.1	Background	51
3.2.2	Fundamental Equation	54
3.2.3	Advantages	55
3.2.4	Disadvantages	56
3.3	Relative Positioning	56
3.3.1	2-Transceiver Positioning	57
3.3.2	3-Transceiver Positioning	58
3.3.3	N-Transceiver Positioning	60
3.3.4	4-Transceiver Positioning	62
3.4	Practical Considerations	65
3.4.1	Vertical Error	65
3.4.2	Multipath	67
3.4.3	Cycle Slip Detection and Correction	68
3.4.4	Synchronization	68
3.5	Summary	70

4	Self-Calibration Methods	71
4.1	Calibration Requirements	72
4.1.1	The Integer Ambiguity Problem	72
4.1.2	Pseudolite Locations	74
4.1.3	Line Biases	75
4.2	Self-Calibration Problem Formulation	76
4.2.1	Complete State Estimation	76
4.2.2	Batch Estimation	78
4.3	Linear Iterative Least Squares	80
4.4	Monte-Carlo Simulation	85
4.4.1	Simulation Description	86
4.4.2	Simulation Results	87
4.5	Quadratic Iterative Least Squares	87
4.6	Multiple-Estimate Solution	95
4.6.1	ILS/QILS Failure Modes	96
4.6.2	Multiple Seeding	97
4.6.3	Algorithm Comparison	97
4.7	Additional Considerations	102
4.7.1	Geometric Factors	104
4.7.2	Non-Geometric Factors	107
4.8	Summary	108
5	Experimental System	109
5.1	GPS Transceivers	110
5.1.1	Transceiver Architecture	110
5.1.2	Pseudolites and Receivers	112
5.1.3	Antennas	113
5.1.4	Secondary Transceiver Components	114
5.2	Base Station	119
5.3	GPSMixer Software	120
5.3.1	Receiver Interface	123
5.3.2	Pair Manager	123
5.3.3	Array Manager	123

5.3.4	Bias Estimator	126
5.3.5	Real-Time Considerations	127
5.3.6	Future Improvements	129
5.4	K9 Rover	129
5.5	Test Location	131
5.6	Summary	133
6	Experimental Results	135
6.1	Bidirectional Ranging	135
6.1.1	Stationary Ranging	136
6.1.2	Dynamic Ranging	138
6.1.3	Maximum Range	138
6.2	Array Localization	141
6.3	Self-Calibration	143
6.3.1	Coarse Calibration	143
6.3.2	Fine Calibration	146
6.4	Error Sources	148
6.4.1	Cycle Slips/Signal Dropouts	148
6.4.2	Clock Drift	149
6.4.3	Truth-System Error	152
6.5	Summary	153
7	Conclusions	155
7.1	Contributions and Results	155
7.2	Mars Considerations	157
7.3	Future Directions	158
7.3.1	Signal Structure	158
7.3.2	3-Dimensional Arrays	159
7.3.3	Large Arrays	160
7.3.4	Over-The-Horizon Arrays	161
7.3.5	Mobile Arrays	162
7.4	Final Thoughts	162
	Bibliography	163

List of Tables

4.1	Self-Calibration Sample Point Requirements	80
4.2	Algorithm Success Comparison	102
6.1	Range Measurement Stability	137
6.2	Initial Position Errors (Code-Phase)	145
6.3	Final Position Errors	146
6.4	Final (Removed) Ranging Biases	148

List of Figures

1.1	Mars Pseudolite Array	6
1.2	SCPA Architecture	8
1.3	Related Research	9
2.1	Self-Differencing Transceiver Architecture	19
2.2	Receiver Quantization Interference	22
2.3	Pulsing Effect on the Automatic Gain Control	26
2.4	Receiver Saturation — Fast AGC	31
2.5	Receiver Saturation — Slow AGC	32
2.6	Multiple-Pseudolite Pulse Interference Limit	35
2.7	Receiver Sidelobe Tracking	37
2.8	Pseudolite Pulsing Scheme Guidelines	41
3.1	Satellite-Based CDGPS	46
3.2	Pseudolite-Based Differential GPS	50
3.3	Self-Differencing	52
3.4	Bidirectional Ranging	53
3.5	3-Transceiver Lateral Positioning Errors	59
3.6	3-Transceiver No-Solution Region	61
3.7	Triangulation-Based Position Solutions	63
3.8	4-Transceiver Lateral Positioning Errors	66
4.1	Self-Calibration Procedure	77
4.2	Linear Iterative Least Squares Process	81
4.3	Monte-Carlo Simulation Configuration	88
4.4	Linear ILS Algorithm Success Rate	89
4.5	Quadratic Iterative Least Squares Process	93
4.6	Linear Versus Quadratic Contributions	94

4.7	Quadratic ILS (QILS) Algorithm Success Rate	95
4.8	Multiply-Seeded Quadratic ILS Algorithm Success Rate	98
4.9	Algorithm Success Rates WRT Transceiver 3 Location — Small Biases, Singly-Seeded	100
4.10	Algorithm Success Rates WRT Transceiver 3 Location — Small Biases, Multiply-Seeded	101
4.11	Algorithm Success Rates WRT Transceiver 3 Location - Large Biases	103
4.12	DOP Variation With Respect To Transceiver 3 Location	105
5.1	Transceiver Architecture	111
5.2	Transceiver Components	112
5.3	Custom Dipole Antenna	115
5.4	Antenna Tripod	116
5.5	Transceiver Tote	119
5.6	GPSMixer Software Architecture	121
5.7	GPSMixer Software Interface	122
5.8	Receiver Interface	124
5.9	Pair Manager	125
5.10	Array Manager Output	126
5.11	Bias Estimator Controls	127
5.12	Bias Estimator Output	128
5.13	K9 Rover	130
5.14	K9 Transceiver	131
5.15	NASA Ames Test Site	132
5.16	NASA Ames K9 Field Test	133
6.1	Measured Range Data	136
6.2	Dynamic Range Data	139
6.3	Maximum Range	140
6.4	Positioning of Mobile Transceiver	142
6.5	Pre-Calibration Position Estimates	144
6.6	Post-Calibration Positioning	147
6.7	Cycle-Slip Removal Through Multiple Self-Calibrations	150
6.8	Receiver Synchronization Errors	152

Chapter 1

Introduction

The planet Mars has long held the human imagination, and has intrigued people who have attempted to understand its unusual properties. Its mysterious retrograde motion has attracted the attention of observers dating back five millenia, and it was Kepler's study of Mars' motion that prompted his theory of elliptical orbits for the planets in 1609. With the introduction of the telescope in the early 17th century the 'Red Planet,' aided by its proximity to Earth and its exposed terrain features and recognizable 'seasons,' became a favorite object of scrutiny by astronomers. Speculation over extraterrestrial life on Mars has brewed for the past 300 years, beginning largely with Huygens' *Cosmotheoros* in 1698. Observations of 'canals' on the surface by Secchi (1858), Schiaparelli (1877), and Lowell (1906), provided further fuel for the public imagination, as did literature such as H.G. Wells' *War of the Worlds* in 1897 ([21][27]).

Over the last century much of the speculation was seemingly put to rest by better ground observations from Earth and by unmanned spaceprobes sent to the planet. The Mariner orbiters and the Viking landers revealed a dry, barren surface seemingly devoid of life. The past decade has seen a great increase in interest in Mars, however, spurred by a wave of orbiter, lander, and rover missions and the new insights they have provided. Arguments abound over the presence of microbial life, such as that purported to have been discovered in Mars meteorite ALH84001, while other debates focus on the presence of potential liquid water in subsurface reservoirs.

The scientific questions posed, and their implications, go far beyond the issue of life on Mars alone: they have great relevance towards our understanding of life on Earth as well. What are the conditions necessary for life to develop? To survive? Why did planets such

as Earth and Mars, seemingly so similar in their early development, take such drastically different courses? As humanity begins to understand the natural bounty and environmental interdependence in our own biosphere, we must also ask how vulnerable it is. How sustainable are our resources, and how dependant are we upon the survival of any given species? Could Earth, through human mis-management, one day come to look more like Mars? And if so, could we step to another planet and get a second chance? It is the quest to find answers to questions such as these that drives the exploration of Mars.

1.1 Project Motivation

Both NASA and the European Space Agency (ESA) have a strong current interest in Mars exploration. As of the time of this writing, NASA plans to launch a pair of Mars Exploration Rovers (MERs) in 2003. These will traverse the surface, examining the geology and looking for signs of water; and hence potential reservoirs for life. That same year, the ESA plans to launch the Mars Express spacecraft with the Beagle 2 lander for additional surface studies. Other landers and rovers will likely follow in the near future, with the greatest near-term goal being the collection of samples and their eventual return to Earth. Although mission goals over the long-term have not been defined in such detail, it is likely that they will be intended to lead up to an eventual human presence on the planet. In order to minimize the likelihood that accidents will result in the loss of human life, this will require extensive resource collection and infrastructure construction by robots before astronauts are ever launched from the Earth.

One of the many difficult problems associated with Mars rovers is the issue of navigation. Mars Pathfinder and its Sojourner rover are a case in point. Sojourner was able to make only limited traverses in the close vicinity of the lander. Its odometry built up tens of centimeters of positioning error over several-meter traverses, and as a result its position estimate had to be updated repeatedly from stereo camera fixes from the lander ([39][1]). Although Sojourner's navigation system was admittedly primitive, more-advanced navigation systems under development for future missions still rely upon some combination of inertial and vision sensors. The Global Positioning System (GPS) is used on Earth as a highly-accurate truth system for many rover prototypes such as JPL's Field Integrated Design and Operations (FIDO) testbed [57], but because of the absence of a GPS satellite constellation around Mars, GPS is not included on any of the designs intended for actual Mars surface use.

There are many potential Mars surface missions that could greatly benefit from improved rover navigation technology, especially when such technology facilitates cooperative operations between multiple units. Robotic surface exploration can immediately benefit from improved navigation, both as an aid to localizing discoveries and to help the robots return to previously explored sites. Construction of in situ propellant plants or coordinated mining of natural resources, either to prolong longevity on the surface or to facilitate a return trip to Earth, are other important applications. Habitat construction is perhaps the ultimate example due to the large scale of the endeavor and the stringent positioning requirements for successful mating of separate components and modules. Accurate knowledge of position on the surface may also aid future astronauts, preventing them from getting lost and greatly aiding human-robot interactions.

The purpose of the research reported here is to create an advanced navigation system that meets the needs of future-generation Mars surface missions. The critical design criteria met by this system are that it...

- Provides centimeter-level, drift-free positioning in a sizable local region (constrained by line-of-sight coverage only)
- Allows shared access and a common reference frame for multiple users: either robotic vehicles or astronauts
- Is not reliant upon other navigation systems
- Is suitable for autonomous deployment

In addition, it is also desirable that the system be easily and cheaply constructible using current technology and commercial off-the-shelf (COTS) components.

1.2 Mars Navigation Options

There are many different navigation technologies currently available for future Mars surface missions, each with its own potential advantages and disadvantages. This section summarizes some of the most attractive options, with a special emphasis on their suitability for Mars use.

Inertial sensors will doubtless play an important role in the future, much as they did for the Sojourner rover. Odometry together with accelerometers and rate gyros can generate

approximate knowledge of rover position, and has the advantage of being self-contained on the vehicle. Coupling the inertial system with sun sensors or star trackers for heading — Mars has no global magnetic field, rendering a compass useless — eliminates some of the associated drift, but a large component still remains. Although they are suitable for rough positioning over short trajectories, this drift makes inertial sensors alone unsuited for precise, repetitive tasks over long durations or for cooperative actions between rovers.

Computer vision is a very attractive sensor for robotic vehicles. Besides providing information of great scientific interest about the surroundings, it can also be used for navigational purposes. Vision allows a rover to find targets and avoid obstacles, and in certain conditions can be used for precise relative navigation through techniques such as visual servoing. A vision system tracking distinctive terrain features can even potentially be used to localize a robot's position with respect to those features [15]. While these benefits make computer vision systems almost indispensable, they still suffer many drawbacks. Because cameras are angle-based systems, accuracy degrades with target distance and localization of distant objects may be poor. Vision is often adversely affected by uneven or changing lighting conditions, darkness, or obscuration by blowing dust. Identifying landmarks can be difficult as well: In addition to the challenge of choosing features that are significantly distinguishable from one another, the same feature can also look quite different from differing viewpoints. All of these factors limit the ability of computer vision to provide precise navigational information.

Optical metrology is another potential sensor. Scanning lasers can make maps of surrounding terrain, and aimed lasers can determine the range and bearing to surrounding objects. As with vision, however, identifying reference points from different viewpoints is difficult. Moreover, optics and reflective markers would be vulnerable to coating with the ever-present Martian dust, making their longevity a questionable factor.

On Earth, a wide variety of radio navigation aids and services are available to aid in robot navigation. The most effective of these is the Global Positioning System, a constellation of Earth-orbiting satellites that provides precise position and timing information to terrestrial users. Techniques such as carrier-phase differential GPS (CDGPS) can yield subcentimeter-level relative positioning for users near a reference ground station. A Mars analog to the terrestrial GPS system could provide much of the desired navigation capability to Mars rovers. Such a 'Mars Positioning System' was proposed by Zimmerman [68]. Unfortunately, the high launch costs associated with Mars missions make it unlikely that a full satellite

constellation will be deployed in the near future. Current NASA discussions focus on a much smaller satellite constellation called the Mars Network, with only six navigation satellites. Such a sparse constellation would allow global position fixes, but only on a periodic basis (roughly every 2 hours) and with limited accuracy (10 meters) ([18][3]). Although this system would be extremely valuable for providing general infrastructure, it would be unable to provide the high-accuracy, real-time positioning capability desired for groups of mobile robots cooperating on the surface and performing complex, precise tasks.

1.3 Self-Calibrating Pseudolite Arrays

Although a Mars-based GPS satellite system is not likely in the near future, GPS-style navigation has been successfully demonstrated in the absence of GPS satellites through the use of small ground transmitters called pseudo-satellites, or ‘pseudolites.’ An array of these pseudolites could be deployed on the Martian surface as depicted in Figure 1.1, and would give centimeter-level, repeatable navigation to any number of rovers operating within the line-of-sight coverage area. Such a system could provide many of the navigation capabilities lacking in other technologies, and would be of great potential value in Mars exploration.

The primary difficulty with autonomous deployment of the pseudolites is determining where the pseudolites have actually been placed. In conventional satellite-based GPS, the satellites are periodically updated with their positions (orbital parameters) from the ground control segment, and then broadcast this information in turn to the users. Such knowledge is essential for accurate positioning. Similarly, in pseudolite arrays it is necessary to know the precise locations of the pseudolites themselves: If centimeter-level positioning of the rovers is desired, then the pseudolites locations must be known to centimeter-level accuracy as well. On Mars, with the pseudolite system most likely autonomously distributed by robots, this results in a very significant ‘chicken and egg’-type scenario: One must have a precise measurement system in order to survey in the pseudolite locations, but it is for the purpose of establishing such a system that the pseudolites are deployed.

In order to resolve this dilemma, this research has created a fundamentally new type of GPS pseudolite-based positioning system: a Self-Calibrating Pseudolite Array (SCPA). An SCPA is able to determine simultaneously both the locations of the pseudolites composing the array and the positions of one or more rovers moving in the vicinity of the array. Unlike previously-existing pseudolite systems, no *a priori* knowledge of the system configuration is




Figure 1.1: Mars Pseudolite Array

A conceptual illustration of a deployed pseudolite array being used for navigation on the surface of Mars. Multiple rovers share the precise navigation capability, which is provided by the pseudolite beacons within a local line-of-sight coverage region. (Photo courtesy JPL)

required.¹ It is therefore highly suitable for robotic deployment and calibration in uncertain environments. This advance in system capability is accomplished through the utilization of GPS transceivers — consisting of colocated pseudolites and receivers — instead of separate broadcast and receive elements. This collocation provides greatly increased observability of the locations of the stationary transceivers making up the array and makes accurate system self-calibration possible.

The SCPA is a distributed system composed of several stationary GPS transceivers within the same geographic region and within common line-of-sight. Multiple mobile receivers or transceivers may operate within the broadcast area of the array, obtaining centimeter-level navigation information as if they were in the presence of a satellite-based CDGPS system. The actual navigation is accomplished through a special solution technique called bidirectional ranging, which was created for GPS transceiver operations as part of this research and is described fully in Chapter 3. Data from the transceivers are collected over a wireless network, and are processed at a central computer. This system architecture is illustrated in Figure 1.2.

The array self-calibration process involves driving a mobile transceiver around the array of stationary transceivers. Carrier-phase ranging data between the transceivers are collected and batch processed, yielding both the locations of the stationary transceivers and the actual trajectory traveled by the rover to centimeter-level accuracy. Determining the locations of all of the array elements is equivalent to resolving the carrier-phase integer ambiguities between each of the possible transceiver pairs.

Because the near-field geometry of the array cannot be completely described using linear relations, development of the SCPA has required the creation of new algorithms and methods more advanced than those typically used for GPS pseudolite applications. These nonlinear and stochastic algorithms enable successful array self-calibration under most possible array geometries and under extremely poor initial estimates of the system geometry. This robustness of the self-calibration process itself is another key property of the SCPA.

¹Orbital relaxation, a somewhat similar — but more limited — calibration process for the GPS satellite system, is discussed in Chapter 4.



Figure 1.2: SCPA Architecture

This illustration depicts a Self-Calibrating Pseudolite Array (SCPA) with four transceivers, one of them mounted on a roving platform. The solid arrows between the transceivers represent a subset of the bidirectional GPS ranging signals, which connect each transceiver pair. Data from the transceivers are collected over a wireless ethernet system (dotted lines) for processing on a central laptop computer. (Note that not all transmission paths are shown for the sake of clarity.)

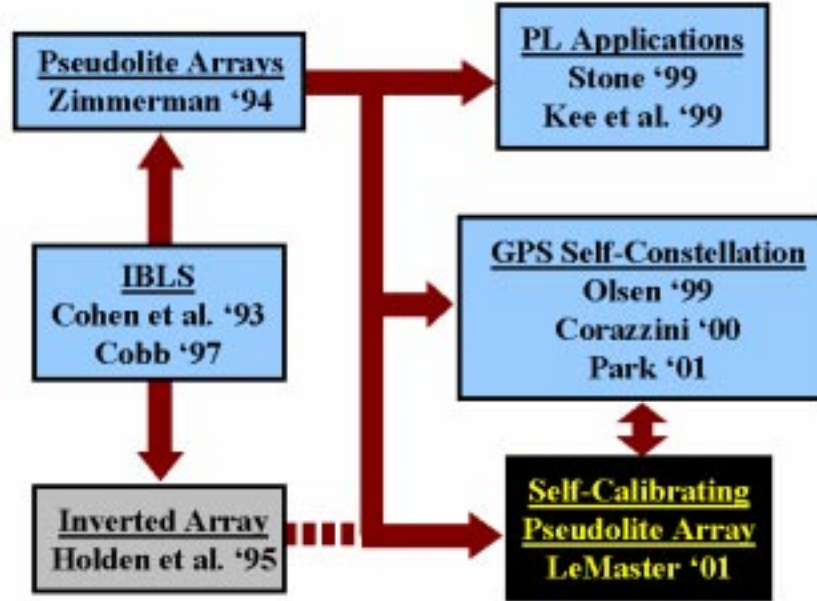


Figure 1.3: Related Research

Some of the principal stages in the development of GPS pseudolite systems together with the major paths of influence. Minor influence is indicated by the dashed line.

1.4 Related Research

Very few — if any — of the advances of science are accomplished in a vacuum, and the present research is no exception. The ideas presented in this thesis have built upon a foundation of previous accomplishments and ideas laid by other researchers, and have both been influenced by and have influenced several contemporary research projects. Although this section cannot describe all of these interactions, it does summarize the most relevant work in the field and describes the interaction between these various research efforts. This interaction is displayed in graphical form in Figure 1.3. Additional prior work not directly related to GPS pseudolite applications is described in later chapters as it becomes relevant.

1.4.1 GPS Pseudolites

The concept of GPS pseudolites is as old as GPS itself: Before the GPS satellite constellation was on orbit, the system was tested using ground-based transmitters. The use of pseudolites to augment the GPS satellite constellation was proposed in 1984 by Klein and Parkinson, who pointed out their benefits towards improved navigation geometry, availability, and reliability and robustness [31]. This paper also first proposed using pulsing as a method for alleviating the near/far problem. More recently, pseudolites were proposed to augment the GPS satellite constellation for the specific application of aircraft autoland by Cohen et al. [7]. The resulting system — known as the Integrity Beacon Landing System (IBLS) — later became part of the Federal Aviation Administration’s Local Area Augmentation System (LAAS) and marks the beginning of pseudolite research at Stanford University.

In 1994 Zimmerman proposed using the pseudolites under development at Stanford to completely replace the existing GPS satellites and enable indoor operations by placing several pseudolites in known locations in a room. This resulted in a series of experiments demonstrating GPS-based relative navigation between free-floating satellite simulators using an indoor pseudolite array [68]. This indoor pseudolite navigation work has since been duplicated by Kee et al. at Seoul National University in South Korea [30]. A different but related approach to satellite-less GPS was demonstrated by Holden and Morley [22]. This work demonstrated positioning of a mobile pseudolite with respect to an array of fixed and pre-surveyed GPS receivers located around the test area, effectively inverting the concept of a GPS pseudolite array. In each of these cases, the ‘near/far’ effects were minimized by operating the mobile units within a local area much smaller than the dimensions of the pseudolite array itself.

1.4.2 GPS Transceivers

GPS transceivers have also appeared several times in the navigation literature, although much less frequently than GPS pseudolites. One of the first instances of GPS transceivers appears in Cobb, wherein a GPS transceiver that is listening to a satellite rebroadcasts that received signal using a different PRN (pseudo-random noise) code, acting as a code-synchronous signal reflector [6]. This device, called a ‘synchrolite,’ forms a key part of the IBLS/LAAS system. The receiver component keeps the device synchronized to GPS time, while the transmit component provides an additional ranging source for the passing aircraft.

A 1999 paper by Stone and LeMaster et al. presented a summary of different kinds of pseudolites and transceivers, and also discussed a variety of potential applications [59]. This work introduced a new kind of transceiver architecture called a self-differencing transceiver, or ‘differlite.’ Rather than actively synchronizing the transmitted signal with the received one, in a self-differencing transceiver the pseudolite and receiver components operate independently using separate oscillators: The receiver monitors the broadcast signal from its associated pseudolite, and synchronization is performed after the measurement process by using modified double-differencing techniques. Although not used expressly as a stand-alone transceiver, this type of architecture has also been utilized by Olsen to augment formation-flying blimps with onboard signal generators [47].

1.4.3 Transceiver Arrays

The next stage in the evolution of satellite-less GPS came very recently, when researchers began using multiple GPS transceivers together to form transceiver arrays. The work mentioned in this section is contemporary with this dissertation, and has both influenced it and been driven by it as well.

Corazzini examined the use of onboard GPS transceivers to aid in relative positioning between formations of satellites. Although this system was intended primarily as an augmentation system for existing satellite coverage, it successfully demonstrated relative positioning indoors over small dynamic ranges using Zimmerman’s free-floating satellite simulator testbed in the absence of external signals [12]. Olsen and Park et al. expanded upon this work by using relative motion of onboard transceivers, in conjunction with external satellite/pseudolite signals, to determine the carrier-phase integer ambiguities between groups of blimps [48] and roving ground vehicles [51].

1.5 Contributions

Although the current work on Self-Calibrating Pseudolite Arrays is strongly related to the previous and contemporary research, it addresses many challenges that have not been previously encountered together in a single system. Specifically, an SCPA must confront, simultaneously, each of the following conditions:

- No *a priori* information about the system states before calibration

- No augmentation of the transceivers with external satellites/pseudolites
- Both mobile and stationary transceivers
- Extreme near/far conditions
- Potentially large code-phase multipath

In order to solve the issues listed above, this research contributed several important advances to the field of GPS pseudolite navigation. These contributions, which break into three separate categories, are listed below.

SCPA Development

- Invented the concept of an SCPA and proved that it is possible to self-survey the locations of all array elements with no *a priori* position information and only limited motion. An SCPA requires no satellite augmentation, and can provide centimeter-level CDGPS accuracy completely independently of any other navigation system. Once self-surveyed, an SCPA behaves like any standard pseudolite array.
- Developed methods for carrier-phase SCPA self-survey and integer determination using a single roving transceiver. Only one transceiver needs to be mobile, and any number of stationary transceivers may be surveyed using these methods. A new quadratic solution algorithm (Quadratic Iterative Least Squares) was developed to account for the inherent system nonlinearities, and when multiply-seeded provides 100% successful array self-calibration to cm-level accuracy under nominal array configurations and operating conditions. Under worst-case operating conditions — with multipath-induced errors in the initial state estimate potentially as large as the array itself — self-calibration success is still greater than 99%.
- Pioneered the use of bidirectional ranging as the sole observable for positioning with GPS transceivers. This is a special variant of conventional double-differencing that allows solutions to be built up on a pair-wise basis using time division multiple access (TDMA), with each partial solution presenting an actual physical quantity (range). Bidirectional ranging provides benefits over conventional double-differencing because it is significantly less vulnerable to signal dropouts and can support a greater number of transmitters in a given area. It also allows code-level positioning and localization

of the array without any special calibration. Chapter 3 describes the heritage of bidirectional ranging in other navigation systems.

- Identified and examined quantitatively the primary factors affecting the operation of SCPAs, including the SCPA configuration itself, the trajectory of the mobile transceiver, the required operational procedures, relevant error sources, and algorithm limitations, together with other practical considerations. Developed operational strategies that address, control, or otherwise overcome these factors.

GPS Transceiver Operations

- Together with Jonathan Stone, examined the tradeoffs between different transceiver architectures and formalized the terminology relating to GPS transceivers. This work is presented in [59].
- Expanded the theory relating to pseudolite pulsing as a solution to the ‘near/far’ problem. Determined a new limit on the maximum allowable pulse width for signal acquisition, and showed how receiver design influences the choice of effective pulsing patterns for pseudolite arrays.

Experimental Validation

- Designed, constructed, and validated experimentally an operational prototype GPS transceiver array based upon a self-differencing transceiver architecture and relying upon extensive use of commercial off-the-shelf components. This is the first complete, stand-alone, transceiver-only GPS positioning system in existence.
- Experimentally demonstrated bidirectional ranging at both code and carrier accuracies. This is the first time bidirectional ranging has been used as the sole measurement observable in an operational GPS positioning system.
- Experimentally demonstrated positioning of a mobile GPS transceiver in a stationary array using bidirectional ranging. This technique, performed for the first time, requires only two stationary transceivers to provide a position fix in 2 dimensions.
- Experimentally demonstrated complete self-calibration of a stationary transceiver array using a single mobile transceiver. This is the first time such a self-calibration has

ever been accomplished. The self-calibration used no external navigation or position sensors and no *á priori* information.

1.6 Dissertation Outline

This chapter has provided relevant background information and the motivation for this research on Self-Calibrating Pseudolite Arrays (SCPAs). Subsequent chapters provide details about the research itself. Because GPS transceivers play such an important role in the operation of an SCPA, Chapter 2 begins this work by describing transceiver architectures and the methods used to overcome practical difficulties such as the near/far problem. Chapter 3 then explains how GPS transceivers may be used — in the absence of carrier-phase integer ambiguities — for relative positioning. Chapter 4 is the core section of this dissertation, and describes the methods developed to determine the carrier-phase integer ambiguities and survey the locations of the array devices. Chapter 5 presents the experimental hardware/software system developed to test these concepts, and Chapter 6 shows experimental results from field tests conducted using the K9 rover at NASA Ames Research Center. These results culminate with a complete centimeter-level array self-calibration of the prototype system. Chapter 7 concludes by summarizing the achievements of this research and the new capabilities provided by the SCPA, and finally discusses the various additional challenges associated with developing a system for actual Mars deployment and offers recommendations for future research.

Chapter 2

GPS Transceivers

Although pseudolite-based GPS systems are becoming more commonplace, most pseudolite applications utilize non-collocated broadcast and receive elements: very few use any of the variety of potential transceiver architectures. Because this research relies heavily upon GPS transceivers, it is necessary to discuss the operation of pseudolite transceivers in some detail and describe the advances in transceiver operations developed during this research.

This chapter briefly examines some of the costs and benefits associated with GPS transceivers, and describes several of the different transceiver architectures utilized for different applications. It proceeds to review the near/far problem and how it relates to receiver and pseudolite architectures. This is followed by a new, intuitive description of how pseudolite pulsing interacts with the receiver automatic gain control and analog-to-digital converter to allow operations under extreme near/far conditions. This includes an analysis of the maximum allowable pulse width for successful acquisition of pulsed transceiver signals, a property that has not been examined in previous research. The chapter concludes with a brief summary chart to aid the designer in the proper choice of pseudolite pulsing schemes.

2.1 Transceiver Performance Tradeoffs

GPS transceivers provide many advantages over separate pseudolite/receiver systems. Their additional complexity, however, makes them inappropriate for some applications. This section discusses the advantages and disadvantages of pseudolite transceivers, and provides insights into appropriate situations for their use.

Some of the most important benefits of transceivers are...

- Improved geometric observability: By adding pseudorange and carrier-phase measurements from other transmitters to the transceiver, system states such as device locations and clock offsets become more observable. Certain geometric parameters — most notably the range between transceivers — may be solved for directly rather than by iterating on a set of nonlinear equations, provided that certain system biases such as carrier-phase integers and non-symmetric line biases are known .
- Partial solutions: The entire set of system states does not need to be solved for simultaneously, but can be broken down into smaller sub-problems. This yields important computational benefits, significantly aids system troubleshooting, and also opens up the possibility of successful navigation with only a portion of the array broadcasting at any given time.
- Robustness: The greater number of measurements renders the system more robust to signal dropouts and occlusions.
- Synchronization: The pseudolite is able to synchronize its transmission to an existing GPS signal, without the necessity of adding an external synchronization signal or line. The ability to synchronize or measure the pseudolite clocks directly eliminates the need for a separate reference ground station for differential GPS, and also reduces the necessary number of stationary broadcast elements by one.

These valuable — and sometimes essential — benefits are discussed in more detail in Chapter 3. The costs associated with using transceivers include...

- More-sophisticated hardware: Although GPS receivers are relatively cheap, their addition does make the overall system more expensive.
- Added data links: Because the receivers are distributed, an additional communications system is needed to allow them to share data. This system can either be wired, which may be difficult to deploy, or wireless, which may be expensive, have limited bandwidth, and potentially cause interference with other systems. Data could alternatively be transmitted on the GPS signal itself, although it would be restricted to low data rates.
- The near/far problem: Strong signals from the collocated pseudolite may prevent the receiver from successfully tracking GPS signals from other devices.

Although GPS transceivers provide many system improvements and capabilities beyond those of conventional pseudolite systems, it is recommended that they not be used unless they are absolutely necessary. This is primarily because of the difficulties associated with the near/far problem and the additional communications infrastructure required. Many of the same capabilities can be provided through the use of other navigation or measurement systems, especially outdoors on Earth with the presence of the satellite constellation and easy human intervention. GPS transceiver applications tend to be best suited to applications with little to no satellite coverage, extensive autonomy, limited human intervention potential, and a requirement for a minimal number of devices. Applications in which one or more of the pseudolites must be mobile are also prime candidates.

2.2 Transceiver Architectures

All GPS transceivers consist of two primary elements: the receiver and the pseudolite. Among the many ways in which these devices can be connected, most are merely implementation-specific hardware issues that do not effect the fundamental operation of the transceiver significantly. One issue that is very important, however, is how timing synchronization is handled between the pseudolite and the transceiver. The method of time synchronization effectively divides pseudolite transceivers into two classes, each of which is described below.

2.2.1 Synchrolite

The first clock-synchronized pseudolites, or ‘synchrolites,’ were developed as part of the Stanford Integrity Beacon Landing System/Local Area Augmentation System (IBLS/LAAS) program by Cobb [6]. The synchrolite listens to a single GPS satellite signal, remodulates it with a different pseudo-random noise (PRN) identification code, and then rebroadcasts it with a fixed and known phase delay from the original signal. This effectively creates another GPS satellite located on the ground, with the full level of clock synchronization possessed by the GPS satellites themselves. It is therefore essentially a GPS repeater. Another type of synchrolite does not rebroadcast a received signal, but rather ties together the pseudolite and receiver with a common oscillator. This provides similar benefits to the previous case, although the device is not synchronized to GPS time. It also requires an additional data link for the receiver. Both types of synchrolites demand high attention to

detail in the design process in order to achieve the sub-nanosecond-level timing required for centimeter-level positioning accuracy.

Other devices, which simply rebroadcast all of the satellite signals received, are used to ‘pipe’ GPS signals into buildings and to perform post-processed missile and projectile tracking [66]. These applications do not require precise knowledge of the processing delay in the devices, however, and so they cannot be truly classified as synchrolites.

2.2.2 Self-Differencing Transceiver

Most of the benefits associated with synchrolites can be achieved using conventional receivers and pseudolites with only minimal modifications. If the pseudolite transmission is monitored by a collocated receiver, that receiver is able to determine the instantaneous clock offset between the two halves of the device. This measured clock offset is then subtracted from the other receiver measurements, synchronizing the two oscillators in software. The resulting self-differencing transceiver, or ‘differlite,’ is much easier to construct than the previous synchrolite transceivers. Commercial off-the-shelf (COTS) components can generally be used, and the receiver and pseudolite oscillators are simply allowed to drift naturally with respect to each other.

The self-differencing transceiver architecture used in this research is illustrated in Figure 2.1. Two methods of monitoring have been employed, with the pseudolite output being fed directly into a dedicated radio-frequency (RF) front end on the receiver through a splitter and alternatively by airwave transmission to a single RF front end. The former method works well for non-pulsed pseudolites at low power, because the dedicated front end allows fine tuning of the monitored signal power. At higher signal power with pulsing, however, these benefits are lost due to RF leakage. In this case it is easier to use airwave transmission, eliminating the need for a second front end and somewhat simplifying the transceiver design.¹

Note that a self-differencing transceiver requires an additional data link in order to transfer the data from the receiver for processing. In addition, just as with a standard pseudolite array, an array of transceivers still needs to synchronize the receiver *sampling*

¹This architecture does introduce range biases from the unequal line biases in the transceivers and from the near-field transmission properties of the pseudolite signals. These biases are observed to be a couple of meters or less in the experimental system used for this research. Although methods exist to measure and correct for these biases in isolation, they are automatically eliminated as part of the carrier-phase array self-calibration process presented in Chapter 4.

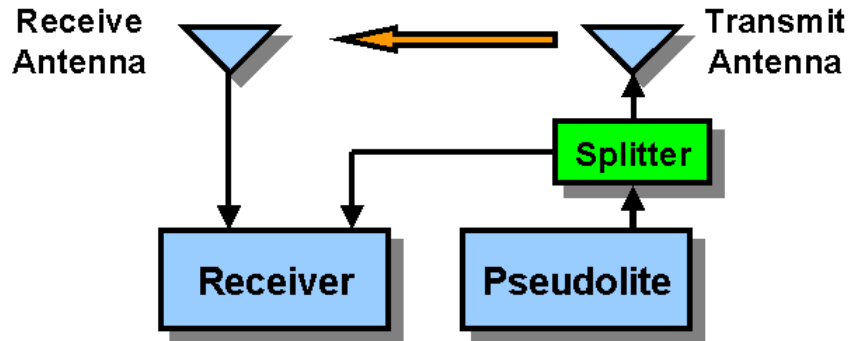


Figure 2.1: Self-Differencing Transceiver Architecture

The transceiver architecture utilized for this research. The receiver and pseudolite are independent commercial off-the-shelf devices with free-running local oscillators. The receiver can monitor its associated pseudolite either by direct measurement through a separate RF front end, or by picking up the signal through the receive antenna.

times. This is commonly done through the use of a ‘master’ pseudolite, a concept originated by Zimmerman [68] and discussed in more detail in Section 3.4.4.

2.3 The Near/Far Problem

The so-called ‘near/far problem’ is one that is inherent to transceiver-based GPS systems, and must be addressed before any successful transceiver system can be created. Simply stated, the near/far problem is the inability of standard GPS receivers to track simultaneously signals of widely-varying strengths, such as those from two different pseudolites at drastically different ranges. It occurs because GPS transceiver and pseudolite applications attempt to use the receivers to do something they were not designed to do, namely tracking a non-satellite signal with power levels considerably different from the nominal satellite signal strength. Although the near/far problem has been addressed in several other works, it is such a fundamental constraint on transceiver systems that it bears summarizing. Furthermore, this thesis provides additional insights into the methods used to combat the problem and explains some previously unrecognized restrictions on these methods.

In order to understand the near/far problem it is necessary first to understand the GPS

signal structure itself and the ways it is utilized by the GPS receivers. The former topic is described fully in ([24][23][63]) and the latter in ([16][62][64]), so only a brief summary is presented here.

2.3.1 GPS Signal Characteristics

GPS is a code-division multiple-access (CDMA) system that relies upon a bi-phase shift keyed (BPSK), 1023-bit Gold code modulated upon the carrier wave in order to distinguish between the signals broadcast by the different satellites. This is called the course-acquisition, or C/A, code.² These pseudo-random noise (PRN) codes are nearly orthogonal (-21.6 dB maximum cross-correlation), so a correlation process performed in the receiver using its own duplicate code sequences allows it to isolate each satellite signal from the others and from the background noise. In fact, the GPS system takes advantage of the correlation gain from this process (approximately 43 dB) to use very low broadcast power in the satellites. By the time it reaches the terrestrial user, the signal from each GPS satellite (~ -153 dBW max) is actually below the receiver thermal noise floor and is only detectable by using this correlation process.

Besides being very weak, GPS satellite signals are also relatively uniform in strength over the surface of the globe. The broadcast power of each satellite changes relatively little during the course of its operational life, and special beamshaping of its signal makes it fairly uniform for users directly under the satellite as well as for those who are viewing it near the horizon.

Receiver Design

GPS receivers incorporate these features of the GPS signal structure into their designs. Of primary importance for this discussion is the uniform signal strength, which allows the receiver to use a low number of bits in the analog-to-digital converter (ADC). Typical receivers use 2- or even 1-bit ADCs. In order to keep the signals within the working range of the ADCs, receivers with greater than 1-bit ADCs employ variable amplifiers upon the incoming signal in what is called the automatic gain control (AGC). For example, in the GP2000 chipset employed in this research, the AGC attempts to adjust the incoming signal until 70% of the received signal measurements fall within the first bit of the 2-bit ADC [58].

²Complete details of the GPS satellite signals and the codes used therein are found in [53].

The narrow dynamic range of most GPS receiver designs is chosen to simplify the design and reduce costs, and does not adversely affect the performance under standard operating conditions.

Most pseudolite systems — and especially transceiver systems — severely violate the assumption that signal strength is approximately constant. In free space, when removed from the local effects of the antenna transmission pattern, the received signal power drops off as $\frac{1}{R^2}$ with the distance from the transmitter. Ground-wave transmission paths are even more severely attenuated, sometimes dropping off as quickly as $\frac{1}{R^4}$ [28]. Thus, the signal strength from a transmitter located 10 cm away from a receiver will be at least 60 dB greater — and possibly up to 120 dB greater — than from one located 100 m away. The extra power from the nearby transmitter masks the signal from the far transmitter within the receiver through two mechanisms: the receiver quantizer and the Gold code structure itself.

Quantization Interference

The manner in which the receiver quantization masks low-strength signals is illustrated in Figure 2.2. The top two plots show two received pseudolite signals of moderately different strengths (10 dB), while the third plot presents the combined signal. The fourth plot represents what happens to the signal after it passes through the AGC — which endeavors to keep most of the signal in the lower bit of the ADC — and into the quantizer. In this case the rounding in the quantizer completely eliminates the first signal, leaving only the second high-strength signal remaining.³

The figure does not show exactly what happens in the GPS receiver because the two incoming signals are (nominally) buried under the receiver thermal noise floor, only to be recovered through the correlation process. The noise tends to ‘lubricate’ the quantizer, and actually lets through some of the low-strength signal. However, the illustration remains instructive, especially when dealing with high-strength pseudolite signals that may actually be above the noise floor. In that case the simplification presented here becomes closer to reality.

³It is sometimes said that the strong signal has ‘captured’ the the quantization nonlinearity.

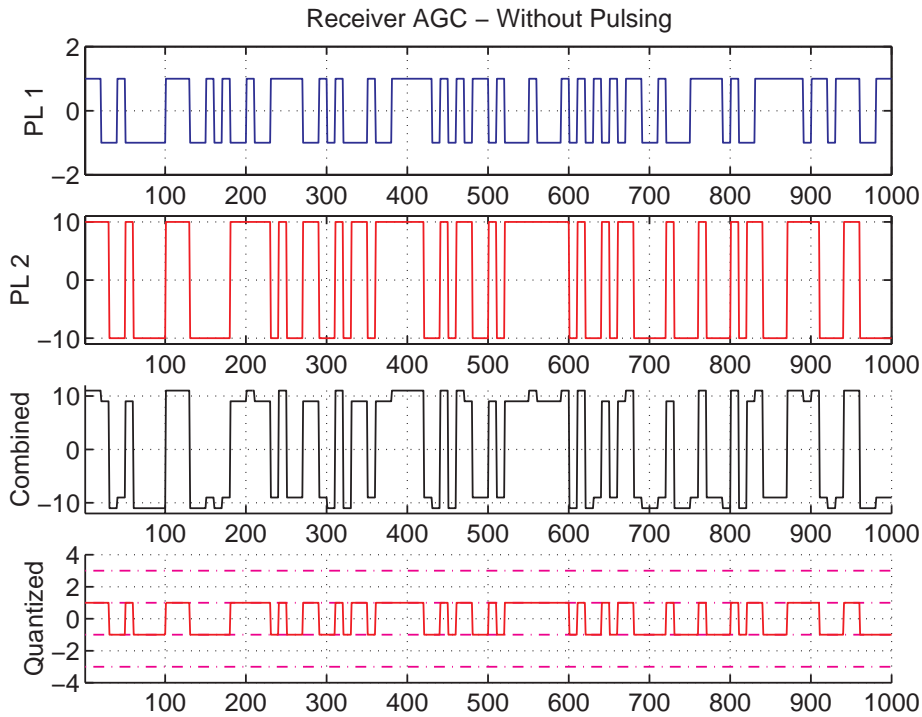


Figure 2.2: Receiver Quantization Interference

Combination of two signals in a 2-bit quantizer, without noise. The top two plots show two different raw coarse-acquisition (C/A) code signals at different strengths, corresponding to different ranges to the transmitters. The third plot shows the combination of the two. The final plot shows the final signal after passing through the automatic gain control (AGC) and being sampled in the analog-to-digital converter (ADC): only the first signal is recovered. The AGC is programmed to modify the gain such that 70% of the quantized samples lie at ± 1 .

Cross-correlation Interference

Even without quantization interference, higher-strength GPS signals would still interfere with low-strength signals. This is because of the GPS signal structure itself. It was mentioned previously that the GPS Gold codes are only ‘nearly’ orthogonal, rather than truly so. In fact, the auto- and cross-correlation processes yield many ‘sidelobes,’ which are only 21.6 dB lower than the main correlation peak (worst case, depending upon the relative Doppler). The presence of these sidelobes means that if one pseudolite were broadcasting a very strong signal, receivers attempting to track another pseudolite’s PRN code could in fact lock onto one of the sidelobes of this first (stronger) PRN code. Even when the receiver does not actually track the interfering signal by mistake, the high sidelobes can still reduce the ability of the receiver to obtain and maintain tracking lock on the desired signal.

2.3.2 Near/Far Solution Methods

Mitigation of the near/far problem has been a considerable challenge for designers of GPS pseudolite systems. Many of the preliminary techniques are set forth in [6]. This section gives a brief summary of some of the more common methods currently employed and some promising methods that have been proposed.

Tuning the Power Level

The most common approach to solving the near/far problem is simply to tune the pseudolite power levels to avoid the problem altogether. This is the strategy employed by much of the previous research in pseudolite arrays, such as by Zimmerman [68], Teague [60], and Olsen [47]. The pseudolite power levels are set so that within a limited working region, the received power from each different pseudolite is roughly the same. As the receiver moves towards the edges of this operating region the pseudolite signal strengths either increase or decrease dramatically, and the positioning system ceases to function properly. In a moderately-sized room, this operating region may be as small as a few square meters.

This method may be ineffective in a couple of different situations. If the pseudolites themselves are mobile, the individual power levels would have to be dynamically programmed to maintain the operating region. Although possible, this greatly adds to the system complexity. In addition, the presence of multiple users in different locations may preclude any instantaneous power setting that allows all of the users to track the pseudolites

simultaneously.

Modifying the GPS Hardware

Another solution to the near/far problem involves modifying the receiver to increase the dynamic range. This would involve increasing the number of bits in the ADC, increasing the range of the AGC, and increasing the maximum allowable input power to the RF front end. This can potentially be an expensive and time-consuming undertaking, and so tends to be the solution of last resort. Moreover, it does not affect the cross-correlation interference inherent with the GPS C/A code.

Modifying the Pseudolite Signal

Because one of the most problematic effects of near/far is associated with the presence of the sidelobes in the auto- and cross-correlation process, it is interesting to discuss the possibility of reducing these sidelobes through a change in the GPS signal structure itself. For example, using a signal based on the longer military P-code (10.23 MHz chipping rate as opposed to 1.023 MHz for C/A) gives a 23 dB reduction in cross-correlation with another C/A code [44]. For pseudolite-only systems, even greater improvements can be made by choosing proper codes. As receivers move from a more hardware-oriented processing approach towards a more software-oriented approach, this scheme could even conceivably be used without hardware modifications in any receiver using a software-only correlator. While longer codes have improved cross-correlation properties, however, they also increase acquisition time, and cannot overcome the receiver quantization limit if the received strength is above the receiver noise floor.

Van Dierendonck et al. have proposed offsetting the pseudolites from the nominal GPS L1 frequency of 1575.42 MHz, and placing them in the 10 MHz null of the GPS C/A and P-code spectrums [17]. While this would eliminate interference between the pseudolites and the satellite signals, any system with multiple pseudolites would require each pseudolite to be on a separate frequency, resulting in a complete frequency-division multiple-access (FDMA) system somewhat similar to the Russian GLONASS system. This would require extensive hardware modifications to allow the filters to pass all of the desired frequencies, and would make CDGPS solutions more difficult because of frequency-dependant group delays in the filters themselves.

Pseudolite pulsing is by far the most widely accepted method of mitigating the near/far problem. It will be discussed in detail in the next section. As opposed to the previously-discussed signal modifications that altered the CDMA or FDMA properties, pulsing adds time-division multiple-access (TDMA) properties to the pseudolite signals.

2.4 Pseudolite Pulsing

In conventional GPS, the satellites are broadcasting continuously. This continuous signal is then sampled by the receiver at discrete timesteps. Once a sufficient number of samples are taken — usually enough to fill one code epoch of 1 ms length — the receiver correlates the sampled data against an internal code replica to decode the timing and navigation data. This correlation process means that the receiver is not actually required to receive the signal at all times. Receiving a fractional part of a code epoch merely reduces the magnitude of the correlation peak, resulting in a lower signal-to-noise (S/N) and signal-to-interference (S/I) ratios. In many cases this has no noticeable effect on either the ability of the receiver to track the signal or the accuracy of the resulting measurements.

This property of the despreading process can be used to help combat the near/far problem by programming the pseudolites to broadcast in short pulses — less than one C/A epoch length — rather than continuously. To see how this works, consider again the example of a receiver tracking two pseudolite signals with magnitudes varying by 10 dB. The nearby (stronger) pseudolite is now pulsed at a low duty cycle, as is shown in Figure 2.3. The relative magnitudes are easily seen in the combined signal waveform. As in the previous example the combined signal now passes through the receiver AGC, which attempts to keep the majority of the samples within the lower bit of the ADC. Because the low-strength signal is the only one present most of the time, it is the signal that sets the resulting gain. The stronger pseudolite, during the times when it is broadcasting, completely overpowers the weaker signal and actually saturates the ADC.⁴

The result of this process is that the receiver is operating in two separate modes. When the strong pseudolite is actively broadcasting the receiver is saturated, and only the pulsed signal gets through. When it is not broadcasting, the receiver is able to track the other pseudolite signal at its preferred signal strength. The correlation gain for both signals is simply a function of the fractional duty cycle of the code epoch when each signal is received.

⁴The issue of AGC saturation and recovery is covered in Section 2.4.5.

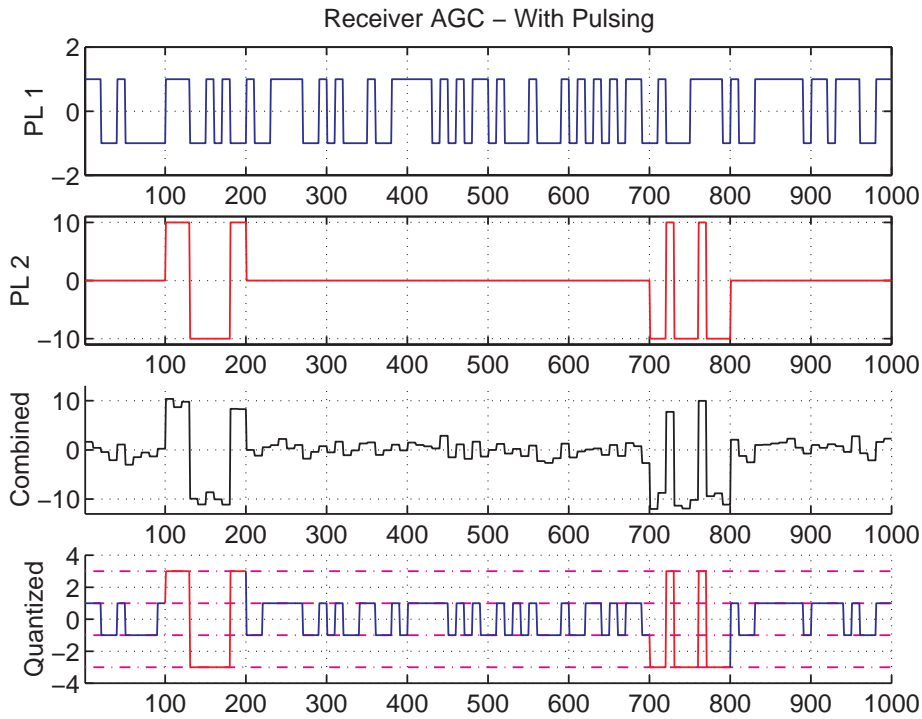


Figure 2.3: Pulsing Effect on the Automatic Gain Control

Combination of two signals in a 2-bit quantizer, without noise. The top two plots show two different raw C/A code signals at different strengths, corresponding to different ranges to the transmitters. In this case the stronger signal is pulsed at a low duty cycle. The third plot shows the combination of the two signals, while the final plot shows the final signal after passing through the AGC and being sampled in the ADC. The AGC is programmed to modify the gain such that 70% of the quantized samples lie at ± 1 . Because of the short duty cycle of the stronger pulsed signal, both signals are recoverable following quantization. Compare with Figure 2.2, in which the weaker signal is never recovered.

Although this example ignores the effect of receiver noise, it does show how pseudolite pulsing takes advantage of both the receiver AGC and ADC to allow it to track simultaneously signals of widely varying strength.

2.4.1 Synchronization

Pseudolite pulses may be either synchronized or unsynchronized. Unsynchronized pulses may overlap, and the resulting pulse interference reduces the ability of the receiver to track the signals and increases the measurement noise. In extreme cases with significant overlap, tracking lock may be lost altogether. With synchronized pulsing, on the other hand, an external clock is used to provide a unified time reference for the pseudolites. Each of the pseudolites is then given an allowed broadcast slot, and all other pseudolites are silent during this time.

Because synchronized pulsing is a true TDMA scheme, it offers many advantages. Since each pseudolite broadcasts only in its dedicated time slot, pulse interference between the different signals can frequently be eliminated altogether. The duration of each pseudolite pulse can be extended to the entire width of the slot, allowing for maximum correlation gain. If these slots are kept short, a large number of pseudolites can be used within the same geographic area. In addition, if the receiver is programmed to know when pseudolites are pulsing, it can turn off its correlator during times when no pulses are occurring. This is called pulse blanking, and increases the received S/I ratio by reducing the correlation with the background noise.

The primary disadvantage with synchronization is the more sophisticated infrastructure required to bring the external timing signal to the pseudolites, especially if these are distributed over a wide geographic area.

2.4.2 Pulse Duration

Pseudolite pulses may be of varying durations. In general, a longer pulse duty cycle with respect to the length of the PRN code will give greater correlation gain, and thus better tracking. There are several factors that may force one to use a short-duration pulse, however.

The first factor is simply the number of pseudolites present in the area. If many pseudolites are to be used simultaneously, then each must be given a shorter broadcast slot to minimize the interference time. Wide pulse widths under these conditions tend to force the

system designer to employ pulse synchronization to keep the pulses from overlapping for more than a minimal duration, thereby increasing the required infrastructure complexity. A second concern is the nature of the interaction between the receiver AGC and the pulsed signal. As will be discussed in more detail below, a receiver with a slow AGC requires a short pulse duration for successful tracking unless the received signal power from each of the pseudolites is relatively uniform, in general an unlikely occurrence. Finally, there is a maximum allowable pulse width for the successful acquisition of GPS signals under extreme near/far conditions.

2.4.3 Pulsing Schemes

There are several different pulsing schemes already used in pseudolite applications. The simplest involves merely broadcasting each pseudolite during a dedicated segment of its PRN code sequence. Another advances the segment broadcast during each epoch through the PRN code, so that over time the pseudolite broadcasts the entire PRN code [13]. Because the potential exists that two different pseudolites could always broadcast at the same instant, these schemes work best with synchronization.

One of the most well-known non-synchronized pulsing schemes is the RTCM (Radio Technical Commission for Maritime Use) pattern [61]. In the RTCM pulsing scheme each pseudolite broadcasts only a 93-chip (1/11 cycle) segment of its PRN code during each code epoch. Every 10th epoch the pseudolite broadcasts 2 segments instead of one, so that the mean duty cycle is 10%. The actual segment location in the PRN code itself varies in a pseudorandom pattern, so in general there will be some interference between multiple pulsed pseudolites. However, if two pseudolites are broadcasting at the same time during one epoch, it is very unlikely that they will interfere during the next. This limits the interference effects considerably. The mean correlation loss $E(\text{atten})$ from this interference (in addition to the -10 dB loss from the low duty cycle) is given by Equation 2.1, where N is the number of pseudolites employed.

$$E(\text{atten}) = 10 \log_{10} \left[\left(\frac{9}{10} \right)^{N-1} \right] \text{ dB} = -0.458 (N - 1) \text{ dB} \quad (2.1)$$

Obviously if two pseudolites start broadcasting at the same point in the pattern simultaneously they will interfere for all time (until their oscillator instabilities cause them to drift apart), although this is an unlikely occurrence.

Other non-RTCM pseudorandom pulsing schemes have also been developed, such as the RTCA (Radio Technical Commission for Aeronautics) pulsing scheme proposed for the LAAS system. This pattern utilizes very short ($\sim 13.7 \mu\text{s}$) pseudorandom pulses — sometimes multiple per epoch — and a low duty cycle ($\sim 2.7\%$) to minimize interference [55]. As these two examples show, pulsing schemes can be quite varied and can be adapted to suit a wide range of operating scenarios.

2.4.4 Power Level

The designer of a GPS pseudolite system has two primary choices for the broadcast power. The first option is to try to make the received power for all receivers uniform, and low enough that the receiver AGC/ADC does not saturate. This involves careful tuning of the power levels, requires a great deal of trial and error, and generally yields a relatively small workspace. It is therefore only usable in situations where the near/far problem is relatively benign. Sometimes pulsing can help extend this method to somewhat greater dynamic ranges, but numerous experiments show that system performance is inconsistent at best.

The best option for GPS transceivers, with their much greater dynamic range, is to set each pseudolite to broadcast at the maximum allowable or available power and then use pulsing to separate each pseudolite via TDMA to eliminate the interference. The goal of this approach, which is utilized by Cobb for the IBLS/LAAS system, is to have the receiver ADC operating in its saturation region *at all times* during a pseudolite pulse. This causes all of the received signals to appear to be of uniform strength with respect to the ADC and the correlator — regardless of received signal power at the RF front end — and ensures that no received pseudolite signal drops below the detection threshold of the least-significant bit (LSB) of the ADC.

2.4.5 Automatic Gain Control Effects

When using the maximum-power strategy just described, the received signal power is assumed to be high enough that the receiver ADC is completely saturated during each pseudolite pulse. Depending upon the actual hardware employed, this may or may not be the case. The most important factor is the recovery speed of the AGC, which sets a limit on the allowable duty cycle for the pulsing pattern. This property was first noted by Cobb, who uses it to describe the simultaneous tracking of satellites together with a single pulsed

pseudolite [6]. This section repeats his line of reasoning, but with modifications to account for the the desire to track multiple pulsed pseudolites simultaneously.

Consider first the case of a ‘fast’ AGC. Such a component can respond very quickly to changes in the received-signal amplitude, with a time constant much less than one code epoch. When the receiver receives pulses of widely-varying amplitudes, the AGC responds quickly and keeps each pulse either within the working range of the ADC or at saturation, depending upon the pulse strength and the minimum gain in the AGC. This is independent of the width of the pulses themselves, as is illustrated in Figure 2.4. The blocks show the actual power of both the noise floor and the pseudolite pulses (which are rising above the noise). The line shows the threshold set by the AGC into the ADC, which is always equal to or less than the pseudolite signal amplitude, as desired.

In receivers with a relatively slow AGC — i.e. with a time constant slower than one code epoch — this is not always the case, specifically when employing long pulse durations. Figure 2.5 illustrates this effect. In the top plot the receiver is tracking four different pseudolites, each with different signal strengths due to near/far. Each is broadcasting at a level above the noise floor, so by the earlier discussion it would appear at first glance that this pulsing scheme should enable full tracking. The threshold line for the AGC, however, reveals that this does not in fact occur. The fourth pulse, although above the noise floor, is much lower strength than the other pseudolite pulses and the AGC cannot bring its power level up fast enough. The fourth pulse is thus in danger of not registering in the least-significant bit of the ADC.

The solution to this problem is always to use very short pulses, as is shown in the second plot. Because each pulse is very short the AGC does not raise its threshold very high, and thus each pseudolite pulse is always in saturation, and therefore trackable. This also keeps satellite signals in range of the ADC and allows joint satellite/pseudolite operations. Note that receivers with 1-bit ADCs have no AGC and always operate in the saturation region, and so exhibit the same behavior as if they had a fast AGC. Receivers with slow AGCs, but in which the AGC has reached its gain limit and is itself always saturated (because of very high signal strength), also behave like fast AGCs.

2.4.6 Interference Effects

Even with a fast AGC, there can be limits on the allowable pseudolite pulse width for certain system configurations. Cobb presents a limit based upon the ability of the receiver

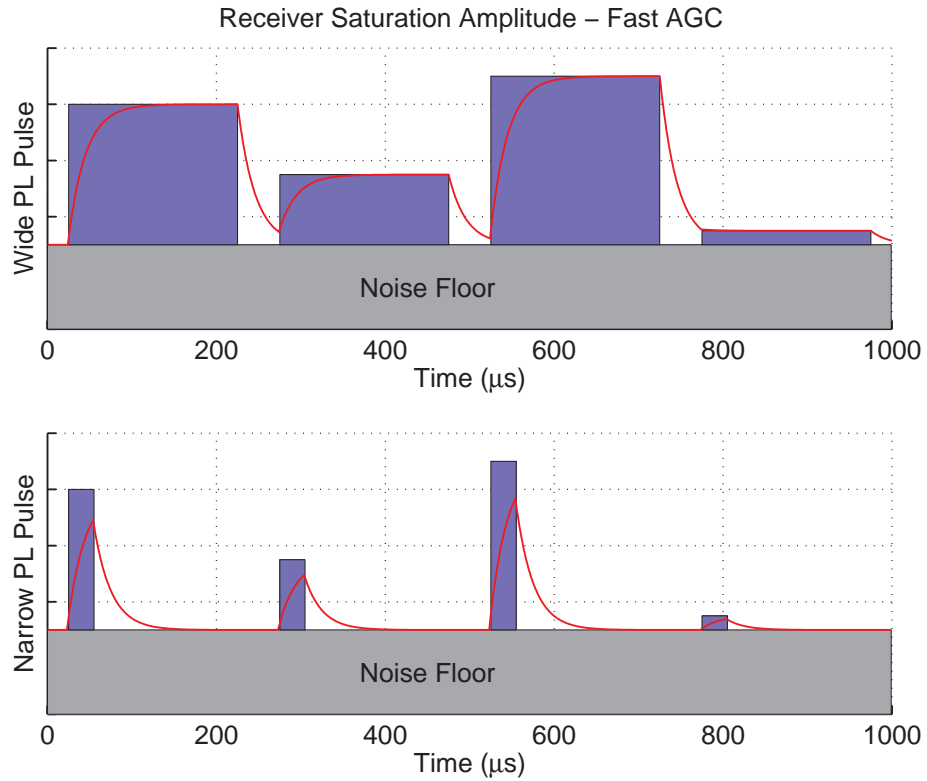


Figure 2.4: Receiver Saturation — Fast AGC

An illustration of how a receiver with a fast AGC is affected by pulsed signals. Four different pulsed signals are present, each one broadcasting in a separate time slot and each one stronger than the noise floor. The solid line shows the AGC saturation value as it attempts to track the received signal strength. For both the wide and narrow pulses, the AGC is able to successfully track the varying signal power. This demonstrates how fast-AGCs are unaffected by pulse width, assuming that the pulses do not overlap.

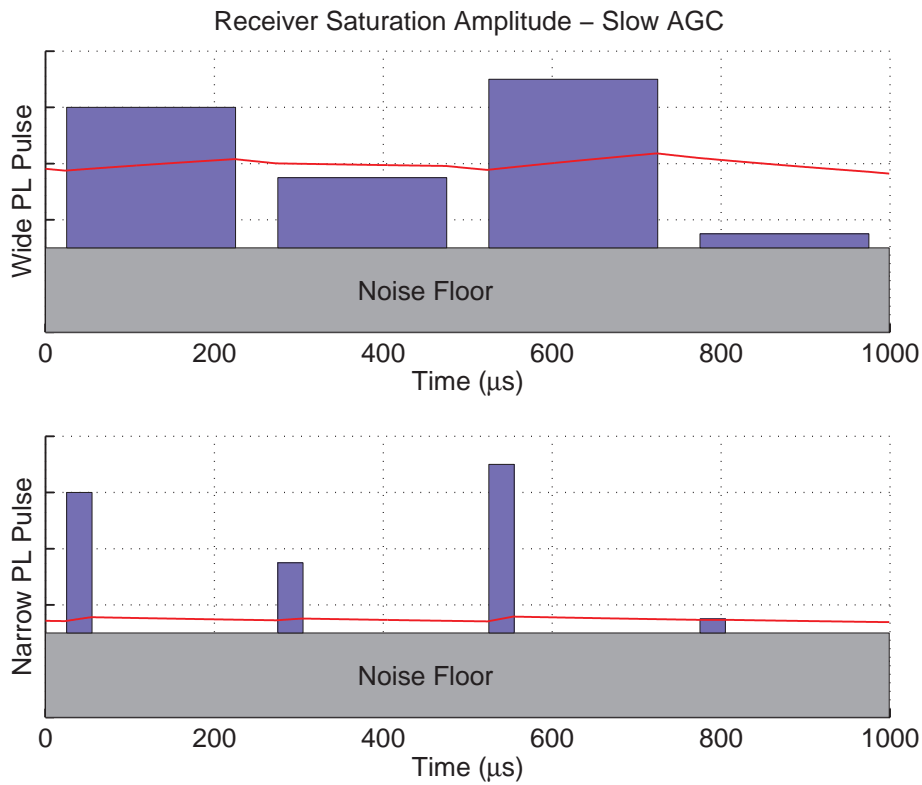


Figure 2.5: Receiver Saturation — Slow AGC

An illustration of how a receiver with a slow AGC is affected by pulsed signals. Four different pulsed signals are present, each one broadcasting in a separate time slot and each one stronger than the noise floor. The solid line shows the AGC saturation value as it attempts to track the received signal strength. With the wide pulses the AGC saturation level rises significantly higher than one of the weaker pulses, resulting in a loss of tracking ability. With the narrow pulses the AGC saturation level stays low, and each of the pulses causes saturation and successful tracking. Compare with Figure 2.4.

to track a satellite signal in the presence of interfering, high-strength pseudolite signals. This decrease in tracking ability arises from two factors: First, there is a loss of correlation with the desired signal as the pulse — which is assumed to saturate the receiver — blocks it out. Second, the pulse interference also replaces (and effectively raises) the receiver thermal noise floor during the pulse.

The pulse interference relation derived by Cobb for a single pseudolite interfering with satellite tracking is repeated here in Equation 2.2 [6].

$$(S/I)_{avg} = 10 \log_{10} \left[\frac{s_{typ} \cdot (1 - d)}{p \cdot d + (1 - d)} \right] \quad (2.2)$$

$$\begin{aligned} (S/I)_{avg} &= \text{Average signal-to-interference ratio (dB)} \\ s_{typ} &= 10^{(S/I)_{typ}/10} \\ p &= 10^{(P/I)/10} \\ d &= \text{Fractional pseudolite pulse duty cycle} \end{aligned}$$

The relevant measures of signal strength are defined below...

$$\begin{aligned} (S/I)_{typ} &= \text{Typical tracked (satellite) signal-to-interference ratio (dB)} \\ (S/I)_{max} &= \text{Maximum tracked (saturated) signal-to-interference ratio (dB)} \\ P/I &= \text{Maximum cross-correlation pulse interference } ((S/I)_{max} - 21.6 \text{ dB}) \end{aligned}$$

Note that these noise and interference values are *not* referenced with respect to a 1 Hz bandwidth, so that these ratios express the actual power contained within the receiver bandwidth.

This equation assumes that only one pulsed pseudolite is operating and that the receiver is non-blanking, so that it continues its correlation process even during the interval of the pulse interference. The numerator reflects the loss of correlation gain on the satellite signal — nominally tracked at $(S/I)_{typ}$ — during the pulse. The denominator shows how the noise level increases during the pulse up to the maximum S/I of the receiver (assuming the pseudolite pulse is in saturation), minus the cross-correlation loss from the differing PRN. Once the average S/I drops below the minimum tracking threshold, satellite tracking is lost. For example, in the Mitel GP2000 chipset used for this research $(S/I)_{max} \approx 33 \text{ dB}$ and the effective tracking threshold is roughly 6 dB (nominally at 3 dB plus a 3 dB factor of safety). If we assume that the satellite signals are tracked at $(S/I)_{typ} \approx 17 \text{ dB}$, this yields a maximum pseudolite duty cycle of approximately 45%.

Cobb's analysis must be modified before it may be used to predict pulse limits for a pseudolite or transceiver array. This is because the number of pseudolites is now greater than one, thereby increasing the baseline interference, and also because the tracked pseudolite signal is itself at a low duty cycle. In addition, *all* of the tracked signals are assumed to saturate the receiver. Equation 2.3 shows the new relation that is applicable to the SCPA. Each of the N pseudolites is broadcasting at the same duty cycle, and the receiver AGC is again in saturation mode during each of the pulses. It is also assumed that the pulses are synchronized, and so do not overlap. A similar relation may be derived for non-synchronized pulsing, but would depend upon the specific pulsing scheme used. For small N , the effect of non-synchronization is minimal.

$$(S/I)_{avg} = 10 \log_{10} \left[\frac{s_{max} \cdot d}{p \cdot (N - 1) \cdot d + (1 - (N - 1) \cdot d)} \right] \quad (2.3)$$

$$s_{max} = 10^{(S/I)_{max}/10}$$

Figure 2.6 plots the results of Equation 2.3 with respect to the pseudolite duty cycle for different numbers of pseudolites, together with results of Cobb's satellite interference analysis for a single pseudolite for comparison. Although $(S/I)_{avg}$ decreases as the number of pseudolites increases, it actually increases as the duty cycle increases, at least up to the limit where the pulses start to overlap. This is because of the increased correlation time for the tracked signal. Of special interest in this plot is the result that even with large numbers of pseudolites operating in a local area, tracking can be maintained with pulse duty cycles of less than 1%. Therefore, while it is advisable from a pulse-interference standpoint to operate a pseudolite array at the maximum duty cycle obtainable without significant pulse overlap, this is not a necessity. This fact becomes important when one considers the acquisition of variable-strength CDMA signals with significant cross-correlation sidelobes.

2.4.7 Sidelobe Acquisition Effects

While pulse interference places only a moderate limit on the maximum pulse duty cycle for pseudolite arrays (due to pulse overlap), there is a more stringent limit due to the signal acquisition process. This limit — which has not been previously addressed by other researchers — is caused by the GPS signal structure itself, specifically the cross-correlation of the C/A code sidelobes. The problem stems from the fact that a receiver with standard

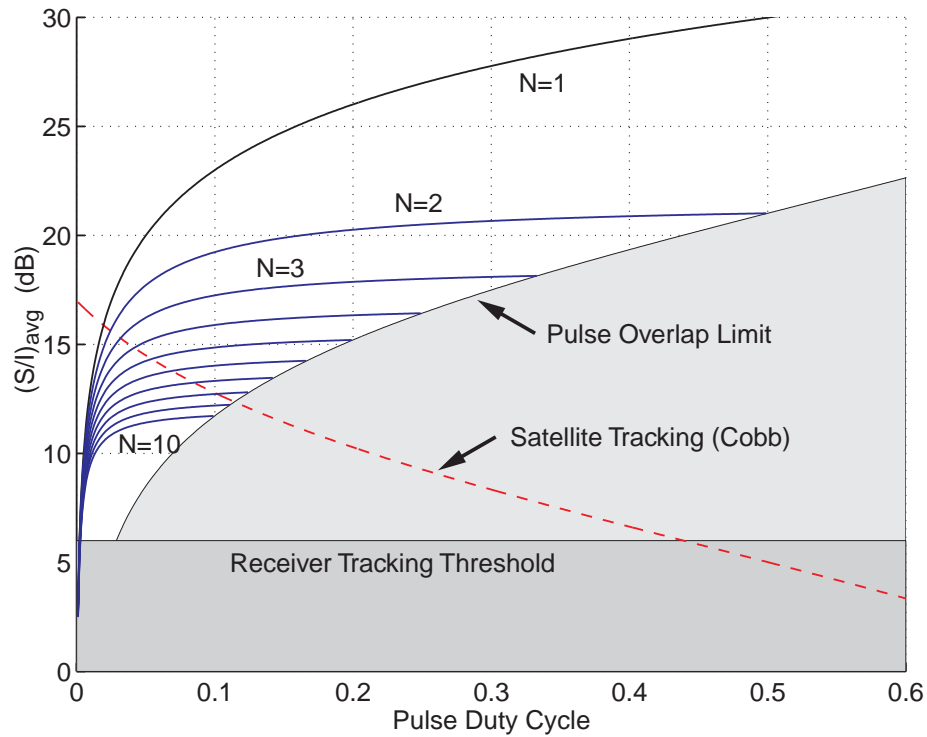


Figure 2.6: Multiple-Pseudolite Pulse Interference Limit

Interference limits for tracking of pulsed pseudolite signals within an all-pseudolite array. It is assumed that all of the pseudolite signals are in saturation (33 dB), and that the pulses are synchronized so that they do not overlap. The tracked S/I decreases slightly as the number of pseudolites increases. Unlike the satellite tracking degradation described by Cobb, which is also plotted here for comparison, pseudolite tracking improves as the pulse width of each pseudolite increases, up to the limit of pulse overlap. In addition, it does not degrade significantly for shorter duty cycles until those duty cycles become extremely small, allowing the pulses to be kept short if this is otherwise necessary.

programming, upon first acquisition of a GPS signal, does not know what strength signal it is looking for. It therefore assumes that any large increase in correlation above a certain detection threshold corresponds to the main lobe on the desired PRN. This works well in satellite-based GPS, when all of the signals are roughly uniform strength and the sidelobes are small.

With pseudolites broadcasting at high power, however, the cross-correlation sidelobes from other pseudolite PRNs can be large enough that they appear to be the main lobe of the desired PRN. When an SCPA is initially deployed and turned on, the receivers do not know their relative locations and therefore do not know at what signal power to look for the pseudolite signals. Because nearby pseudolites signals are much higher strength, it is very likely that any receiver looking for a far-away pseudolite signal will successfully correlate to one of the many high-strength sidelobes from a nearby pseudolite. This results in the receiver erroneously tracking the wrong PRN, with a corresponding failure of subsequent navigation.

Figure 2.7 shows this phenomenon in data from several tests performed using the Mitel receivers employed in the SCPA. The non-pulsed pseudolite signal (\circ) is fed through a variable attenuator into the receiver, starting at a low nominal signal strength. As the power level increases so does the measured S/N, until the AGC saturates and the curve levels off. The receiver is then reset so that it momentarily loses lock. When it regains lock, it starts tracking a sidelobe instead of the main lobe because of the high signal strength, resulting in a low S/N and poor tracking. Because the sidelobes are much more numerous than the main lobe, this is the typical response. Once the signal strength is lowered via the variable attenuator, the sidelobe drops below the receiver acquisition threshold and the receiver resumes tracking the main lobe.

There are several potential methods to correct the problem of sidelobe tracking. The first, which is suggested by the previous experiment, is to start the pseudolites at low signal power and then slowly increase to maximum power. This means that the receiver will automatically track the main lobe because it is the first to rise to the detection level in the correlator, while the sidelobes are still too small to detect. This is what happens automatically in systems such as IBLS/LAAS, in which the aircraft starts outside of the broadcast radius of the pseudolite and then flies in to a region of high power. In practice it is sometimes more convenient to change the amplification in the receiver input path rather than in the pseudolite output path. Employing variable amplification works well

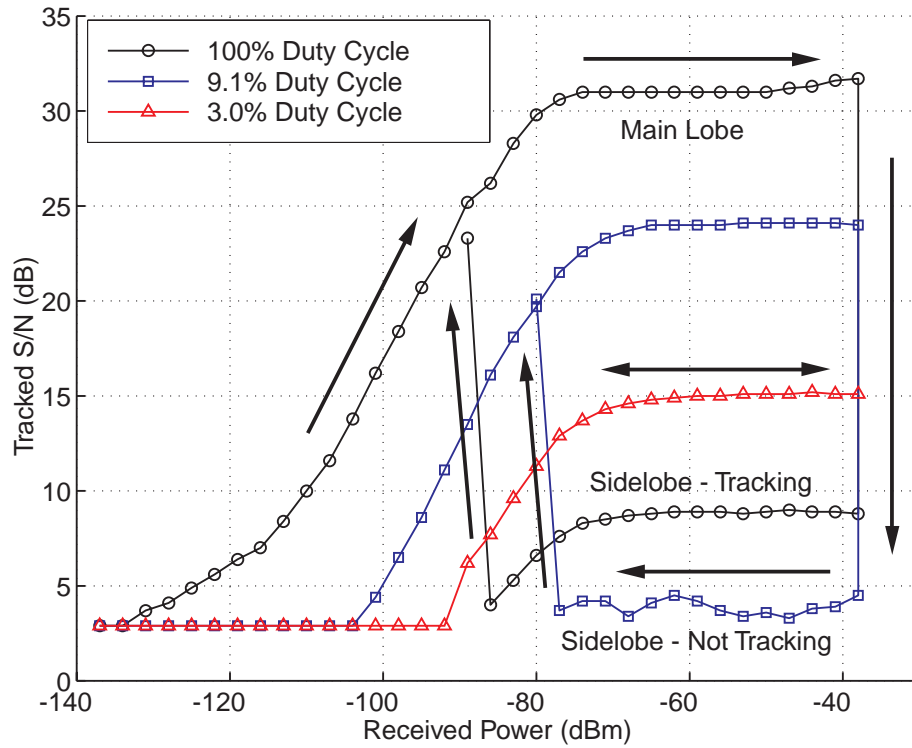


Figure 2.7: Receiver Sidelobe Tracking

Three benchtop experiments showing receiver tracking and acquisition of high-strength pseudolite signals at different pulse duty cycles. The receiver is a Mitel Architect receiver utilizing the GP2000 chipset, and the pseudolite is in IntegriNavatics IN200C. The pseudolite signal is fed directly into the receiver RF front end via cables. In each case the pseudolite is started at low signal strength through the use of a variable attenuator. Attenuation is gradually decreased (following the path of the arrows, left to right), allowing the receiver to begin tracking the pseudolite signal. Eventually the receiver AGC saturates, after which the pseudolite is briefly turned off. After it resumes broadcasting, the receiver re-acquires signal lock. The attenuation is then increased again to lower the received signal strength (right to left).

In the first experiment at 100% pulse duty cycle (\circ), the receiver re-acquires a high-strength sidelobe. Once the sidelobe drops enough in strength, the receiver loses lock on it and re-acquires the main lobe instead. In the second experiment (9.1% pulse duty cycle, \square) the sidelobe is too weak to track, but is still strong enough to confuse the receiver and prevent tracking of the main lobe. In the final experiment (3% duty cycle, \triangle) the sidelobe is suppressed enough by the narrow pulse width that it does not interfere, and only the main lobe is tracked.

for solitary pseudolites, but is less effective with multiple pseudolites or receivers because of the difficulty in isolating the attenuation for a single receiver channel in conventional receiver architectures. In addition, this technique may be difficult to manage if the receivers experience frequent signal loss.

A second strategy is to perform a more exhaustive search in the GPS receiver to find the correlation peak with the maximum amplitude. Because of the large number of cross-correlation peaks (approximately 250) in the GPS Gold code and the necessity to examine each potential peak long enough to determine its S/I, this yields a noticeable increase in acquisition time, possibly up to several minutes. For many applications, however, a moderate increase in deployment and acquisition time is not a major concern. A fully operational SCPA would likely utilize a combination of these two active sidelobe suppression methods, because they directly address the root of the problem and do not otherwise affect tracking accuracy.

A third method, which is the simplest to implement and is the one utilized in this research, is to decrease the duty cycle of the pseudolite pulses. Decreasing the duty cycle decreases the maximum correlation peak amplitude of both the main lobe and the sidelobes. While the main lobe is still trackable, the sidelobe drops below the detection threshold and no longer confuses the receiver tracking loops, as is illustrated by curve (Δ) in Figure 2.7. This same effect can also be accomplished by simply raising the detection threshold in the receiver, at the expense of making the receiver unable to track lower-strength satellite signals. A disadvantage of using shorter pulse widths to suppress the sidelobes passively is that the tracked S/I decreases, resulting in a slight degradation in positioning accuracy and tracking-loop stability.

Equation 2.4 presents a new relation for the maximum allowable pseudolite pulse duty cycle η_{max} in order to prevent tracking of the PRN sidelobes. This may be derived from examination of Figure 2.7, with the stipulation that the low duty cycle must reduce the strength of the largest possible cross-correlation sidelobe to below the tracking threshold.

$$\eta_{max} \leq 10^{\left\{-\frac{1}{10}((S/I)_{max} + T_{det} - 21.6 \text{ dB})\right\}} \quad (2.4)$$

$$\begin{aligned} (S/I)_{max} &= \text{Maximum tracked S/I (dB)} \\ T_{det} &= \text{Detection threshold (dB)} \end{aligned}$$

For the Mitel Orion receivers used for this research $(S/I)_{max} \approx 33 \text{ dB}$ and $T_{det} = 3 \text{ dB}$,

yielding a maximum pulse duty cycle of 14.5%. Experiments with this particular receiver have shown that the gain at low signal strengths is slightly nonlinear, however, probably due to compensation in the AGC. This allows tracking at slightly lower signal strengths than would be otherwise predicted. When a 3 dB factor of safety is subtracted from the detection threshold to account for these nonlinearities, the maximum pulse duty cycle is reduced to only 7.2%. While this is a significantly more stringent duty cycle restriction than for satellite tracking in the presence of pulse interference, it is not a cumulative effect, and so does not change as the number of pseudolites increases.

It is important to note that although previous research has utilized wider pulse widths, this was accomplished under different conditions. The IBLIS/LAAS system involves a receiver gradually moving into the pseudolite bubble, making acquisition unambiguous. In addition, only two pseudolites are being tracked. Other work with pulsed pseudolites in indoor conditions has utilized pulses as wide as 14% [10]. This was accomplished on a multiplexing TANS Quadrex receiver, however, resulting in only a 3.5% effective duty cycle for each of the four antennas.

2.5 Summary of Near/Far Mitigation Techniques

The near/far problem is one of the greatest technical challenges associated with operating GPS pseudolite — and especially transceiver — arrays. This is because the potential power ratio between distant and nearby pseudolites is very large, forcing the off-the-shelf receiver hardware to handle dynamic ranges that it was never designed for. Many techniques are available to combat the near/far problem, one of the most effective of which is pseudolite pulsing. Synchronization, pulse width, pulse patterns, and power level are all characteristics available to the designer, and can be modified to improve system performance. Of special interest is the realization that there is a limit on the maximum allowable pulse width during acquisition in order for the receivers to avoid locking onto the pseudolite sidelobes. This is a different effect from the previously recognized pulse-width limit that applies when attempting to track satellite signals in the presence of pulsed pseudolites, or when the pulses begin to overlap.

Based upon the research segment described in this chapter, a set of simple guidelines toward designing an effective pulsing scheme for SCPA operation is presented in Figure 2.8. In general, knowledge of the receiver AGC properties, the pseudolite synchronization, and

the cross-correlation sidelobe suppression method will dictate the design choices. Care must be taken, however, in verifying that the chosen pulsing scheme is compatible with both the desired system architecture and all possible modes of operation. A poor choice of pulsing schemes may work adequately under certain benign conditions, and then fail under more adverse dynamic ranges or with additional pseudolites in the local area. When in doubt, it is generally advisable to use narrow pulses. Although this does degrade the positioning accuracy slightly due to the reduced S/I of the tracked signals, it does add the most robustness with respect to unknown interference or other factors.

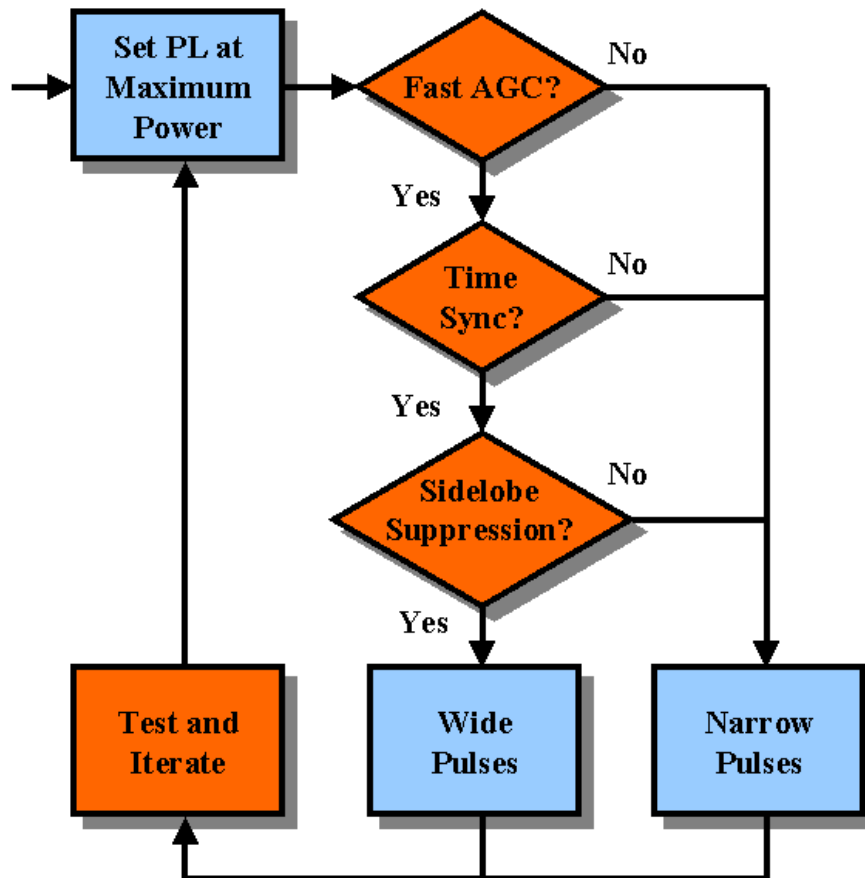


Figure 2.8: Pseudolite Pulsing Scheme Guidelines

Key considerations for the design of a pulsing scheme for pseudolite arrays. The primary decision points involve the speed of the receiver AGC, whether or not the pseudolite pulses are time synchronized, and whether or not an active method is used to detect or eliminate the tracking of cross-correlation sidelobes. Unless the system components are all highly tuned for use with pseudolite applications, the general approach is to use very narrow pulses.

Chapter 3

Transceiver-Based Positioning

This chapter examines how the GPS transceivers described previously may be used to determine the relative locations of the devices comprising a Self-Calibrating Pseudolite Array (SCPA). Because a GPS transceiver is capable of both receiving and transmitting GPS signals, an array composed of transceivers — either in whole or in part — provides many more measurements for use in relative positioning of the devices in the array than when compared to a conventional pseudolite or satellite system. In general, the number of measurements increases from $\mathcal{O}(N)$ to $\mathcal{O}(N^2)$. This gives the designer additional latitude in the choice of positioning algorithms and also creates new solution methods that were previously unavailable. In the special case when the integer ambiguities and line biases are assumed to be known or are relatively small, these transceiver-based positioning techniques yield the expected centimeter-level positioning accuracies. Before array self-calibration has been performed (see Chapter 4), they yield only approximate results.

Of special interest in this chapter is a new basic solution method called bidirectional (or inter-transceiver) ranging. Unlike conventional techniques, which require multiple satellites for any usable solution, bidirectional ranging generates a physically meaningful result — the range between a transceiver pair — using only the local measurements involving that pair. The improved state observability associated with bidirectional ranging is one of the key enabling technologies for the SCPA. Bidirectional ranging also opens up interesting possibilities for operations with sparse arrays and for time-division multiple-access (TDMA) operations with very large arrays.

3.1 Conventional Differential GPS

In order to understand how the addition of GPS transceivers alters the measurement equations, it is instructive to review briefly how GPS positioning is accomplished using conventional satellite and pseudolite arrays. This brief overview focuses on the aspects that relate to navigation using SCPAs, specifically the techniques used for differential GPS (DGPS). More details on DGPS can be found in ([52][14]).

3.1.1 Satellite-Based

Satellite-based GPS was originally designed so that stand-alone users can determine their positions with approximately 10-20 meter accuracy (without selective-availability [SA]) anywhere on the globe. This is done by tracking signals that are continuously broadcast by the satellites, and then determining the time-of-flight of the signal by examining the embedded timing information. The resulting measure of the range to the satellites is called pseudorange, because it is in error due to uncertainty in the local time within the receiver reference oscillator. Accumulating simultaneous pseudorange measurements to four or more satellites (with reasonable sky separation) allows the user receiver to determine its four unknown states: x-y-z position and time. In practice, the accuracy of the satellite-based system may be 5-10 meters or better, depending upon the local signal environment.

The accuracy of the basic GPS system can be improved beyond its basic level by performing relative differential positioning with respect to another fixed reference receiver in a nearby, accurately-known location. This process is called differential GPS, and works because many of the error sources in GPS such as ionospheric, tropospheric, and satellite clock and ephemeris errors are highly correlated with user location. The removal of these correlated error sources on a large scale is the concept behind the FAA Wide Area Augmentation System (WAAS) ([19][29]). Using WAAS or other DGPS techniques, the user location relative to these reference stations can be determined to an accuracy of a few meters. This accuracy degrades outside the coverage area of the reference network because of the spatial decorrelation of the error sources.

Even greater improvements in positioning accuracy can be made by tracking the satellite carrier wave itself, rather than just the modulated code signal. Carrier-phase differential GPS (i.e. CDGPS) uses this approach to achieve centimeter-level positioning accuracy. Figure 3.1 shows the measurement concept of CDGPS for a single satellite. Subtracting the

reference station measurement of the number of carrier cycles to the satellite from the user measurement yields the component of the range between the two receivers that is parallel with the line-of-sight vector to the satellite. This process is repeated with other satellites to obtain the remaining displacement vector components.

There are three characteristics of CDGPS that will play important roles in the discussion to follow. First, the great distance to the satellites together with the relative proximity of the two receivers results in nearly identical line-of-sight vectors for the two receivers, and therefore an essentially linear set of measurement equations. In pseudolite-based GPS, this advantage is generally lost. Second, CDGPS introduces a new system state — the carrier-phase cycle or integer ambiguity — that must be resolved in order for positioning to occur. This ambiguity stems from the fact that the receivers can measure their relative phases to the incoming satellite signal only modulo 2π radians. Much of the research in CDGPS systems involves methods for the determination of these integers. Finally, in order for the differencing process to be successful, the measurements must be taken nearly simultaneously. This is because the errors are temporally as well as spatially correlated, especially the receiver clock biases. The use of bidirectional ranging, which is described later in this chapter, alleviates some of the difficulties associated with this constraint.

Within the general satellite-based DGPS framework there are two primary solution methods, single-differencing and double-differencing. These two techniques are described below.

Single-Differencing

The most basic solution method for CDGPS is called ‘single-differencing,’ and involves the subtraction of ranging measurements from the user and the reference station on a satellite by satellite basis. In the development below this is performed using carrier-phase measurements, because this is the more general case. For code-based differential GPS (DGPS) the equations are the same, with the exception that the carrier-phase integer $\mathcal{N}_u^{(k)}$ is zero. Note that in the discussion to follow, superscripts in parenthesis refer to the satellite or transmitter index, while subscripts refer to the receiver index. This notation is used extensively in this and in future chapters.

For a single receiver making a carrier-phase ranging measurement to a single satellite, the describing equation is




Figure 3.1: Satellite-Based CDGPS

Both the user and the reference receiver take simultaneous phase measurements of the carrier-wave from the given satellite. Assuming the integers are known, this yields the projection ΔR of the range between the two receivers along the line-of-sight (LOS) vector to the satellite.

$$\phi_u^{(k)} = \left(r_u^{(k)} + \tau_u - \tau^{(k)} \right) + \mathcal{N}_u^{(k)} \lambda + b_u^{(k)} + A_u^{(k)} + \nu_u^{(k)} \quad (3.1)$$

where

$$\begin{aligned} \phi_u^{(k)} &\triangleq \text{carrier-phase measurement} \\ r_u^{(k)} &\triangleq \text{range from satellite } k \text{ to user } u \\ \tau_u &\triangleq \text{clock bias for user } u \\ \tau^{(k)} &\triangleq \text{clock bias for satellite } k \\ \mathcal{N}_u^{(k)} &\triangleq \text{integer ambiguity} \\ \lambda &\triangleq \text{carrier-phase wavelength} \\ b_u^{(k)} &\triangleq \text{line bias} \\ A_u^{(k)} &\triangleq \text{atmospheric error} \\ \nu_u^{(k)} &\triangleq \text{noise and other errors} \end{aligned}$$

A similar equation describes the measurement of the same satellite by the reference receiver r . The single-difference measurement is simply the subtraction of the equation for the user from that for the reference station.

$$\Delta\phi_{r,u}^{(k)} \triangleq \left(\phi_r^{(k)} - \phi_u^{(k)} \right) = \left(r_r^{(k)} - r_u^{(k)} \right) + (\tau_r - \tau_u) + \left(\mathcal{N}_r^{(k)} - \mathcal{N}_u^{(k)} \right) + \left(b_r^{(k)} - b_u^{(k)} \right) + \left(\nu_r^{(k)} - \nu_u^{(k)} \right) \quad (3.2)$$

In this process the satellite clock bias, which is identical for both users, is eliminated. The atmospheric error is also assumed to be nearly identical and is therefore eliminated. Noting that the reference receiver and the user have a nearly identical line-of-sight (LOS) vector to the satellite (which is known for the reference station), this can be simplified to

$$\Delta\phi_{r,u}^{(k)} = \underline{\mathbf{1}}_r^{(k)} \cdot \underline{\mathbf{r}}_{r,u} + \Delta\tau_{r,u} + \Delta\mathcal{N}_{r,u}^{(k)} \lambda + \Delta b_{r,u}^{(k)} + \Delta\nu_{r,u}^{(k)} \quad (3.3)$$

with

$$\begin{aligned} \underline{\mathbf{1}}_r^{(k)} &\triangleq \text{unit LOS (row) vector towards satellite } k \\ \underline{\mathbf{r}}_{r,u} &\triangleq \text{vector from reference } r \text{ to user } u \\ \Delta\tau_{r,u} &\triangleq \tau_r - \tau_u \\ \Delta\mathcal{N}_{r,u}^{(k)} &\triangleq \mathcal{N}_r^{(k)} - \mathcal{N}_u^{(k)} \\ \Delta b_{r,u}^{(k)} &\triangleq b_r^{(k)} - b_u^{(k)} \\ \Delta\nu_{r,u}^{(k)} &\triangleq \nu_r^{(k)} - \nu_u^{(k)} \end{aligned}$$

When single-difference measurements are available from many satellites, these equations can be combined to form a set of linear equations (Equation 3.4) that are solvable by taking the pseudoinverse, assuming the integer biases $\Delta N_{r,u}^{(k)}$ and the line biases $\Delta b_{r,u}^{(k)}$ are known.¹ For 3-dimensional positioning, this requires range measurements to four common satellites with significantly different LOS vectors. The problem of determining the integer biases is discussed fully in Chapter 4.

$$\begin{bmatrix} \Delta\phi_{r,u}^{(1)} \\ \Delta\phi_{r,u}^{(2)} \\ \vdots \\ \Delta\phi_{r,u}^{(n)} \end{bmatrix} = \begin{bmatrix} \underline{1}^{(1)} & 1 \\ \underline{1}^{(2)} & 1 \\ \vdots & \vdots \\ \underline{1}^{(n)} & 1 \end{bmatrix} \begin{bmatrix} r_{r,u} \\ \Delta\tau_{r,u} \end{bmatrix} + \begin{bmatrix} \Delta\mathcal{N}_{r,u}^{(1)} \\ \Delta\mathcal{N}_{r,u}^{(2)} \\ \vdots \\ \Delta\mathcal{N}_{r,u}^{(n)} \end{bmatrix} \lambda + \begin{bmatrix} \Delta b_{r,u}^{(1)} \\ \Delta b_{r,u}^{(2)} \\ \vdots \\ \Delta b_{r,u}^{(n)} \end{bmatrix} + \begin{bmatrix} \Delta\nu_{r,u}^{(1)} \\ \Delta\nu_{r,u}^{(2)} \\ \vdots \\ \Delta\nu_{r,u}^{(n)} \end{bmatrix} \quad (3.4)$$

Double-Differencing

In practice, the equations above are further reduced through a process called double-differencing by subtracting the single-difference equations from different satellites. The basic equation is

$$\nabla\Delta\phi_{r,u}^{(k,l)} \triangleq \Delta\phi_{r,u}^{(k)} - \Delta\phi_{r,u}^{(l)} = \left(\underline{1}^{(k)} - \underline{1}^{(l)}\right) r_{r,u} + \nabla\Delta\mathcal{N}_{r,u}^{(k,l)} \lambda + \nabla\Delta\nu_{r,u}^{(k,l)} \quad (3.5)$$

with

$$\begin{aligned} \nabla\Delta\mathcal{N}_{r,u}^{(k,l)} &\triangleq \Delta\mathcal{N}_{r,u}^{(k)} - \Delta\mathcal{N}_{r,u}^{(l)} \\ \nabla\Delta\nu_{r,u}^{(k,l)} &\triangleq \Delta\nu_{r,u}^{(k)} - \Delta\nu_{r,u}^{(l)} \end{aligned}$$

In this formulation the clock biases between the receivers are eliminated because they are common between each single-difference. More importantly, assuming that each receiver has only one antenna, the differential line bias between the two receivers is the same for each satellite, and therefore also cancels. This greatly simplifies the practical aspects of the design, because no line bias calibration is needed.

In a similar manner to the single-differences, the double-differences are then combined to form a linear system of equations, as is shown in Equation 3.6. Note that the double-difference equation is still linear, and that measurements from the same number of satellites

¹Note that if each receiver uses only one antenna to receive all of the satellite signals, then the differential line bias $\Delta b_{r,u}^{(k)} = \Delta b_{r,u}$ is identical for all satellites.

are still required.

$$\begin{bmatrix} \nabla \Delta \phi_{r,u}^{(1,n)} \\ \nabla \Delta \phi_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \phi_{r,u}^{(n-1,n)} \end{bmatrix} = \begin{bmatrix} \underline{\mathbf{1}}_r^{(1)} - \underline{\mathbf{1}}_r^{(n)} \\ \underline{\mathbf{1}}_r^{(2)} - \underline{\mathbf{1}}_r^{(n)} \\ \vdots \\ \underline{\mathbf{1}}_r^{(n-1)} - \underline{\mathbf{1}}_r^{(n)} \end{bmatrix} \underline{\mathbf{r}}_{r,u} + \begin{bmatrix} \nabla \Delta \mathcal{N}_{r,u}^{(1,n)} \\ \nabla \Delta \mathcal{N}_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \mathcal{N}_{r,u}^{(n-1,n)} \end{bmatrix} \lambda + \begin{bmatrix} \nabla \Delta \nu_{r,u}^{(1,n)} \\ \nabla \Delta \nu_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \nu_{r,u}^{(n-1,n)} \end{bmatrix} \quad (3.6)$$

3.1.2 Pseudolite Based

Using pseudolite-based positioning the same basic principles hold. The major difference, however, is that the pseudolite transmitters are located nearby the receivers. This means that the LOS vector to a given pseudolite is *not* the same for the different receivers within the array. This is illustrated in Figure 3.2, in which it is apparent that the difference in range between the two receivers cannot be accurately represented by a projection of one range vector onto the other. Although double-differencing is still used to eliminate the clock and line biases, the resulting equations are nonlinear with respect to the receiver locations. The double-difference equations for near-field pseudolite positioning are presented in Equation 3.7.

$$\begin{bmatrix} \nabla \Delta \phi_{r,u}^{(1,n)} \\ \nabla \Delta \phi_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \phi_{r,u}^{(n-1,n)} \end{bmatrix} = \begin{bmatrix} r_r^{(1)} - r_u^{(1)} - r_r^{(n)} + r_u^{(n)} \\ r_r^{(2)} - r_u^{(2)} - r_r^{(n)} + r_u^{(n)} \\ \vdots \\ r_r^{(n-1)} - r_u^{(n-1)} - r_r^{(n)} + r_u^{(n)} \end{bmatrix} + \begin{bmatrix} \nabla \Delta \mathcal{N}_{r,u}^{(1,n)} \\ \nabla \Delta \mathcal{N}_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \mathcal{N}_{r,u}^{(n-1,n)} \end{bmatrix} \lambda + \begin{bmatrix} \nabla \Delta \nu_{r,u}^{(1,n)} \\ \nabla \Delta \nu_{r,u}^{(2,n)} \\ \vdots \\ \nabla \Delta \nu_{r,u}^{(n-1,n)} \end{bmatrix} \quad (3.7)$$

$$\begin{aligned} r_i^{(k)} &\equiv \left\| \underline{\mathbf{p}}^{(k)} - \underline{\mathbf{p}}_i \right\| \\ \underline{\mathbf{p}}^{(k)} &\triangleq \text{Position of satellite } k \\ \underline{\mathbf{p}}_i &\triangleq \text{Position of receiver } i \end{aligned}$$

One of the most common methods used to solve this set of nonlinear equations is to linearize the equations about an initial estimate of the receiver locations, and then perform an Iterative Least Squares (ILS) calculation to determine the actual locations. This tends to work well when all of the receivers are inside the physical bounds of the pseudolite array, and if the initial guess is not too far from the actual positions.

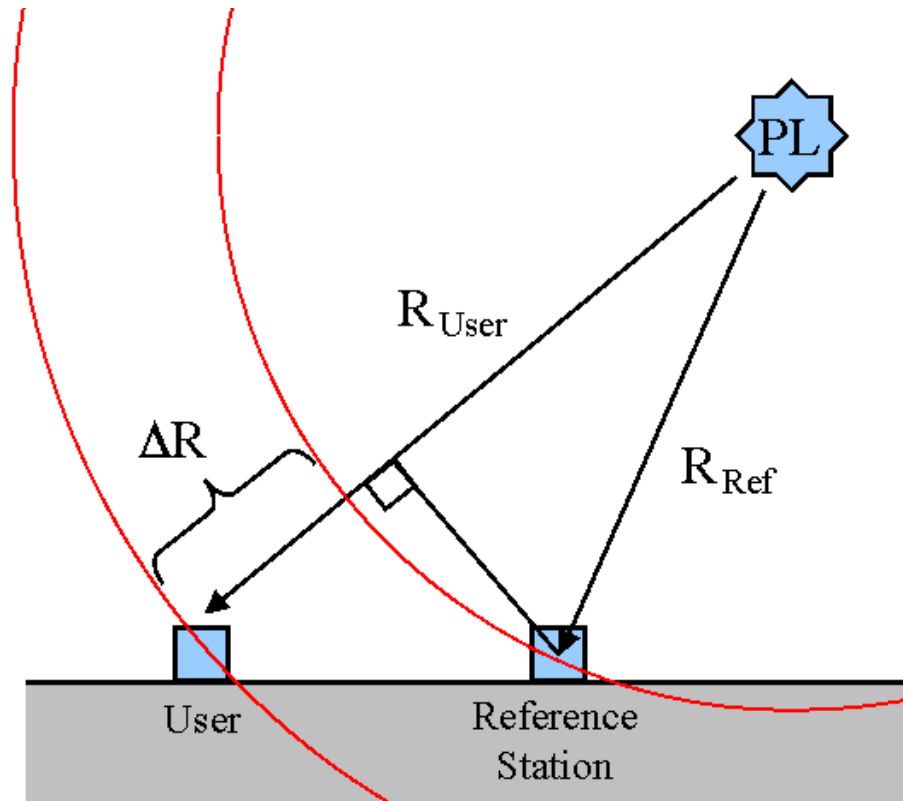


Figure 3.2: Pseudolite-Based Differential GPS

Both the user and the reference receiver take simultaneous phase measurements of the carrier-wave from the given pseudolite (PL). Unlike in the satellite-based case (Figure 3.1), the line-of-sight vector from each user is not identical. This means that ΔR is no longer a simple projection, and the more complicated non-linear geometry must be retained.

Note that achieving accurate positioning using pseudolites requires that the locations of the pseudolites be known to great accuracy (centimeter-level in order to achieve centimeter-level positioning of the receivers). Methods of determining pseudolite locations using non-GPS sensors or by using CDGPS measurements using receivers in *known* locations are presented in [68] and [30]. This thesis addresses the more general problem of determining the pseudolite locations without *a priori* knowledge from another source.

3.2 Bidirectional Ranging

When transceivers are incorporated into a GPS system, an additional measurement type becomes available. In addition to the standard pseudorange measurements and both single- and double-differencing, the system designer may also use a powerful new capability via ‘self-differencing’: including the collocated pseudolite and receiver pair within a transceiver as part of both the transmit and receive components of a double-difference solution. This has the benefit of removing some of the geometric complexity from the system, since the range between the two transmit and receive antennas for the transceiver is known (and can frequently be zero, depending upon the transceiver design). Figure 3.3 illustrates this concept.

This self-differencing capability is especially useful if applied to the transceivers not just once, but rather twice. This generates a solution called bidirectional ranging, wherein only two transceivers are involved in the double-difference. Bidirectional ranging between GPS transceivers yields the range between the two devices as a direct observable (assuming the integers and line biases are known), as is illustrated in Figure 3.4.

3.2.1 Background

As was discussed in Chapter 1, concepts relating to transceiver self-differencing have been extant for several years. Self-differencing has been discussed in relation to formation flying spacecraft using GPS by Corazzini et al. ([10][11][12]), Olsen et al. [48], and Park et al. [51], all of whom use onboard transmitters to augment external satellite or pseudolite signals. Both Olsen and Park use single-differences exclusively for carrier-phase ranging, and compute both the vehicle ranges and clock biases using a batch process. Corazzini uses single-differences for code-based ranging and double-differencing for carrier-based ranging.

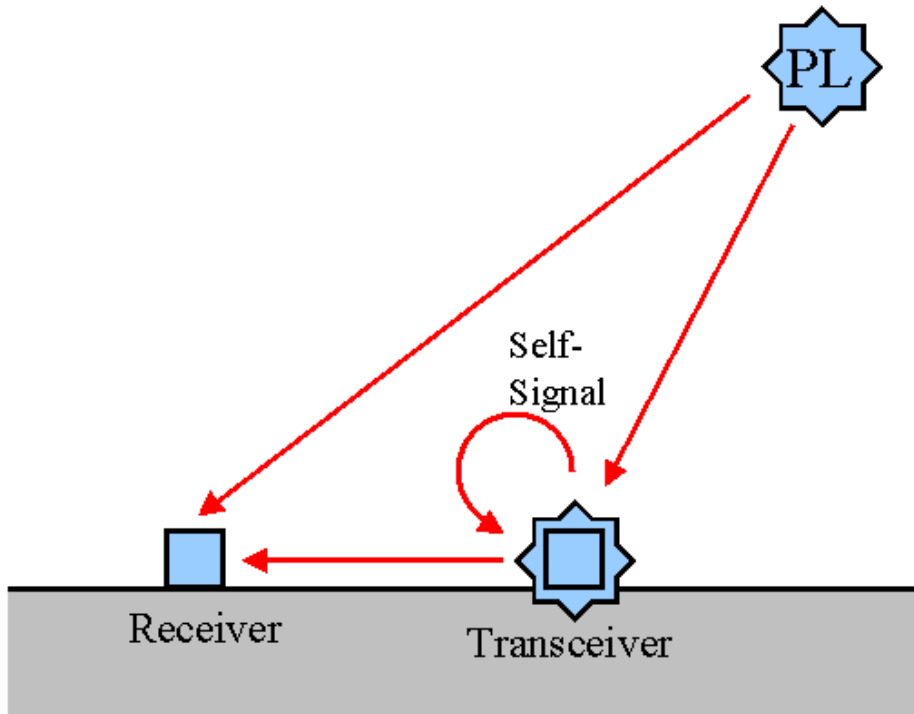


Figure 3.3: Self-Differencing

This modification of conventional double-differencing involves a transceiver together with an additional pseudolite and receiver. Both the stand-alone receiver and the transceiver-based receiver monitor both the pseudolite and the transceiver signal. These two measurement sets are then combined in a double-difference. Using the transceiver in this fashion effectively reduces the number of unknown system states, because there are only three device locations to be determined instead of four.




Figure 3.4: Bidirectional Ranging

Bidirectional ranging results when the self-differencing technique of Figure 3.3 is applied twice to a pair of transceivers. Each transceiver monitors both its own signal and the signal of the other transceiver, ultimately yielding the range between the two devices.

In the latter case bidirectional ranging can appear as a possible permutation on the indexing in the measurement equations, but is not explicitly invoked and is not used between all vehicles.

The concept of bidirectional ranging has been explicitly addressed in many diverse contexts. The Distance-Measuring Equipment (DME) used in general and commercial aviation since the 1950's utilizes a bidirectional signal between the aircraft and a ground transponder to give the range between the two. DME operates on a pulse-response principle in the lower L-band, with accuracies down to 20 meters [28]. Similar acoustic transponders are used for underwater positioning [42]. In more related work, Lau et al. proposed using a GPS-like continuous signal in the Ka-band for bidirectional ranging between formation-flying satellites, although few details are provided [33]. A publication by Olsen also advocates the use of bidirectional ranging as an augmentation source for formation flying [49]. The JPL proposal for the Mars Network utilizes a UHF bidirectional ranging signal for positioning, but again with few implementation details [3]. Finally, Weiser discusses the possibility specifically for GPS systems, but dismisses the idea because of the perceived difficulties in using GPS transceivers [65]. While there are many other applications and technologies well-suited to utilizing bidirectional ranging — Loran chain calibration, Ultra-Wideband (UWB) beacons,

etc. — a complete list and comparison is beyond the scope of this dissertation.

Although a few of these works touch upon the concept of bidirectional ranging using GPS transceivers, none of them uses GPS-based bidirectional ranging as the sole measurement observable. As a result, several of the implications of this solution have not been fully addressed. This will be remedied in the sections that follow.

3.2.2 Fundamental Equation

Since bidirectional ranging is a modification of conventional double-differencing techniques, the development of the governing equations parallels the previous discussion. The starting point is the raw carrier observable previously presented in Equation 3.1, which is repeated below for convenience. In this case it is assumed that the transceivers are close to each other, so atmospheric signal delay is negligible.

$$\phi_i^{(j)} = \left(r_i^{(j)} + \tau_i - \tau^{(j)} \right) + \mathcal{N}_i^{(j)} \lambda + b_i^{(j)} + \nu_i^{(j)} \quad (3.8)$$

Receiver i collects raw carrier-phase measurements from two transceivers: itself and transceiver j . These two measurements are single-differenced internally to give Equation 3.9, where it is assumed that the receive and transmit antennas for device i are collocated and so $r_i^{(i)} = 0$.

$$\overset{\text{b}}{\Delta} \phi_i^{(j,i)} \triangleq \left(\phi_i^{(j)} - \phi_i^{(i)} \right) = r_i^{(j)} + \left(\tau^{(j)} - \tau^{(i)} \right) + \left(\mathcal{N}_i^{(j)} - \mathcal{N}_i^{(i)} \right) \lambda + \left(b_i^{(j)} - b_i^{(i)} \right) + \left(\nu_i^{(j)} - \nu_i^{(i)} \right) \quad (3.9)$$

Note that in general, the integers and line biases along the signal path from pseudolite i to the collocated receiver i will not be zero, and so they are retained in the formulation.

The single-difference above combines with a similar one from receiver j to form the fundamental bidirectional ranging equation below.

$$\overset{\text{b}}{\nabla} \Delta \phi_{i,j} \triangleq \overset{\text{b}}{\Delta} \phi_i^{(j,i)} - \overset{\text{b}}{\Delta} \phi_j^{(i,j)} = 2 \left\| \underline{p}_j - \underline{p}_i \right\| + \overset{\text{b}}{\nabla} \Delta \mathcal{N}_{i,j} \lambda + \overset{\text{b}}{\nabla} \Delta b_{i,j} + \overset{\text{b}}{\nabla} \Delta \nu_{i,j} \quad (3.10)$$

$$\begin{aligned}
\nabla^b \Delta \mathcal{N}_{i,j} &\triangleq (\mathcal{N}_i^{(j)} - \mathcal{N}_i^{(i)} + \mathcal{N}_j^{(i)} - \mathcal{N}_j^{(j)}) \\
\nabla^b \Delta b_{i,j} &\triangleq (b_i^{(j)} - b_i^{(i)} + b_j^{(i)} - b_j^{(j)}) \\
\nabla^b \Delta \nu_{i,j} &\triangleq (\nu_i^{(j)} - \nu_i^{(i)} + \nu_j^{(i)} - \nu_j^{(j)})
\end{aligned}$$

3.2.3 Advantages

Several interesting observations can be made about Equation 3.10. Of primary importance, the range (not pseudorange) between the two transceivers appears as a direct observable, meaning that it can be used as a stand-alone component of the positioning algorithm. This is in contrast to both single- and double-differencing — wherein a set of equations utilizing many different transmitters must be assembled before the system is solved in a batch manner — and has many useful implications that are discussed below.

First of all, the observability of the position states of the system is improved substantially. This enhances algorithm performance and widens the space upon which such algorithms converge, which is important for the array self-calibration algorithms described in Chapter 4. (Note, however, that using conventional single-differences with transceivers can also have this same property.)

Second, using bidirectional ranging allows sparse arrays to be used effectively. Although a single pair of transceivers will not give actual position, the measured range alone may be useful for some applications. More importantly, 2-dimensional relative positioning may be done with only 3 transceivers: for example, one mobile transceiver and two stationary transceivers. Although there are limitations to the area over which such positioning is possible, as will be discussed in the following sections, this is still a large improvement over conventional double-differencing, which requires at least 5 devices for relative 2-dimensional positioning (3 stationary pseudolites, a stationary reference station, and the mobile user).

Third, the fact that the ranging measurements can be solved for piecewise — rather than in a simultaneous batch process for the entire array — allows for the utilization of very large arrays. This is difficult to accomplish using conventional techniques, because the pulsing schemes used to overcome the near/far problem limit the number of pseudolites that may simultaneously broadcast in a given area. With bidirectional ranging, however, individual pairs can be operated on a TDMA (time-division multiple-access) basis and broadcast in turn, assuming that there is no instantaneous motion of the transceivers between the sample times.

Fourth, bidirectional ranging is well suited to over-the-horizon operations. Transceivers do not need to be in line-of-sight with all of the other devices in the array but merely with those closest to them, effectively forming sub-arrays. The subsections of the overall array are computed separately, and can then be ‘stitched together’ via the range measurements between the transceivers on the boundaries of the sub-arrays.

3.2.4 Disadvantages

Bidirectional ranging does have several disadvantages when compared to more conventional techniques. The first, which is revealed in Equation 3.10, is that the line biases do not cancel. In contrast with conventional double-differencing, the signals composing the measurement equations pass down different transmission lines, and so cannot be eliminated through this technique. For carrier-phase positioning this is not a critical problem, because the line biases can be added to the integer ambiguities and determined during that calibration process. For code-based positioning they play a more important role, and to the extent in which they cannot be accurately determined they limit the overall positioning accuracy of system. Note however that these line biases appear in a differential manner: Assuming that the radio-frequency (RF) signal path delays are roughly uniform between different transceivers, this does mitigate their overall effect.

A second major disadvantage with bidirectional ranging is the added system complexity, specifically in the datalink requirements. Rather than collecting data at two receivers, requiring a single data link, data are collected from N receivers, requiring at least $N - 1$ datalinks. Besides the physical hardware required, this also may require increased communications system bandwidth. The hardware aspect of this problem may be alleviated by integrating these data links directly upon the data message of the transmitted GPS signals. Because of the low data rate (nominally 50 bps, with a potential maximum of 1 kbps), however, such a technique is still bandwidth limited unless the underlying GPS signal structure is modified to accommodate a faster data rate. Note that this disadvantage applies to all GPS transceiver applications, and not just to those using bidirectional ranging.

3.3 Relative Positioning

This section examines how the bidirectional ranging signals between transceiver pairs can be used to determine the relative positions of the transceivers, and therefore to reconstruct

the entire array geometry. As with the range measurements themselves, this geometric reconstruction may be performed at either code- or carrier-level accuracies. In the case of carrier-phase positioning, it is assumed at this point that the integer ambiguities have been successfully resolved using the methods of Chapter 4.

It is a relatively straightforward task to determine the relative locations of a large network of receivers when using satellite-based GPS. Assuming that each of the receivers sees the same set of satellite signals, standard differential solution methods will generate a redundant set of relative vector displacements between the receivers. These vector displacements may then be combined in an optimal manner such as least squares to determine the overall array geometry and eliminate independent measurement errors. This process is called network adjustment, and is discussed in some detail in [34]. The simplicity and reliability of this method is largely due to the fact that the satellite locations are known, and so only a subset of the overall system states must be determined.

Reconstructing the array geometry of an SCPA using bidirectional ranging is not this straightforward because none of the geometric states of the system is known. This difficulty is compounded because the relative measurements are scalar ranges instead of vector displacements. This loss of directional information greatly reduces the observability of the states of interests, and also makes the measurement equations nonlinear. Triangulation-based solution methods work well for truly planar or unambiguously non-planar configurations as long as the measurement error is relatively small. In the presence of measurement noise or biases, however, or when the array geometry is poor — such as when the array is slightly non-planar — the problem is complicated significantly.

3.3.1 2-Transceiver Positioning

Positioning (i.e. ranging) between a single transceiver pair using bidirectional ranging is a relatively trivial operation. A direct application of Equation 3.10 on the measured code- or carrier-phase measurements yields the range between the two devices.

Bidirectional ranging lends itself well to code-range solutions because of the lack of integer ambiguities. If the two transceivers are experiencing little relative motion, time averaging is very effective at reducing the random measurement noise. Line biases and multipath can limit accuracy significantly, however, and cannot be eliminated by time averaging. Alternatively, bidirectional ranging via carrier-phase measurements provides much greater accuracy, assuming the integers can be determined via some other method.

If only changes in range are of interest, the rough nominal range can be determined using code-ranging — possibly averaged over time — and then later changes in range can be very accurately determined using carrier-phase ranging.

3.3.2 3-Transceiver Positioning

Determining the relative locations of three transceivers using bidirectional ranging presents a somewhat more interesting problem. Obviously, only relative positioning in a plane can be accomplished. At its simplest level, positioning is performed via triangulation. This is an unambiguous process — with the exception of a possible inverted (mirrored) solution — provided that the measurement errors are negligible. When the measurement errors are not negligible, such as when the carrier-phase integers are not known exactly, this triangulation solution method can run into severe difficulties. The errors in the measurements effectively warp the solution space, resulting in both positioning errors and potential solution singularities. Standard covariance analysis techniques can account for some of this effect by placing probabilistic bounds on the locations of the transceivers, but it cannot be used in the singular regions. This section examines how this solution space warping occurs and how it can cause positioning errors much larger than the magnitude of the range errors themselves.

Figure 3.5a shows one such example of solution-space warping. Two stationary transceivers are placed at *known* locations 1.0 units apart. A third mobile transceiver, marked by the dot, then triangulates its resulting position based upon its measured ranges to the two stationary transceivers. These measurements are both corrupted by biases, in this case +0.1 units. The contour plot shows the magnitude of the positioning error for the third transceiver at all possible locations around the array. This error varies from 0.1 units on the x-axis outside of the array to a maximum of 0.332 units at the midpoint between the two stationary transceivers. Thus, the magnitude of the positioning error can be over three times that of the measurement errors themselves.

Figure 3.5b shows the same situation but with both measurements in error by -0.1 units. In this case the same type of warping occurs, with the same minimum and maximum position errors, but with the added condition that if the mobile transceiver is inside the gray ellipse, no triangulation solution is possible at all because the two ranging vectors are too ‘short’ to meet at any point. Standard covariance-based location analyses are ineffective in degenerate cases such as this, and so special heuristics must be developed to handle these

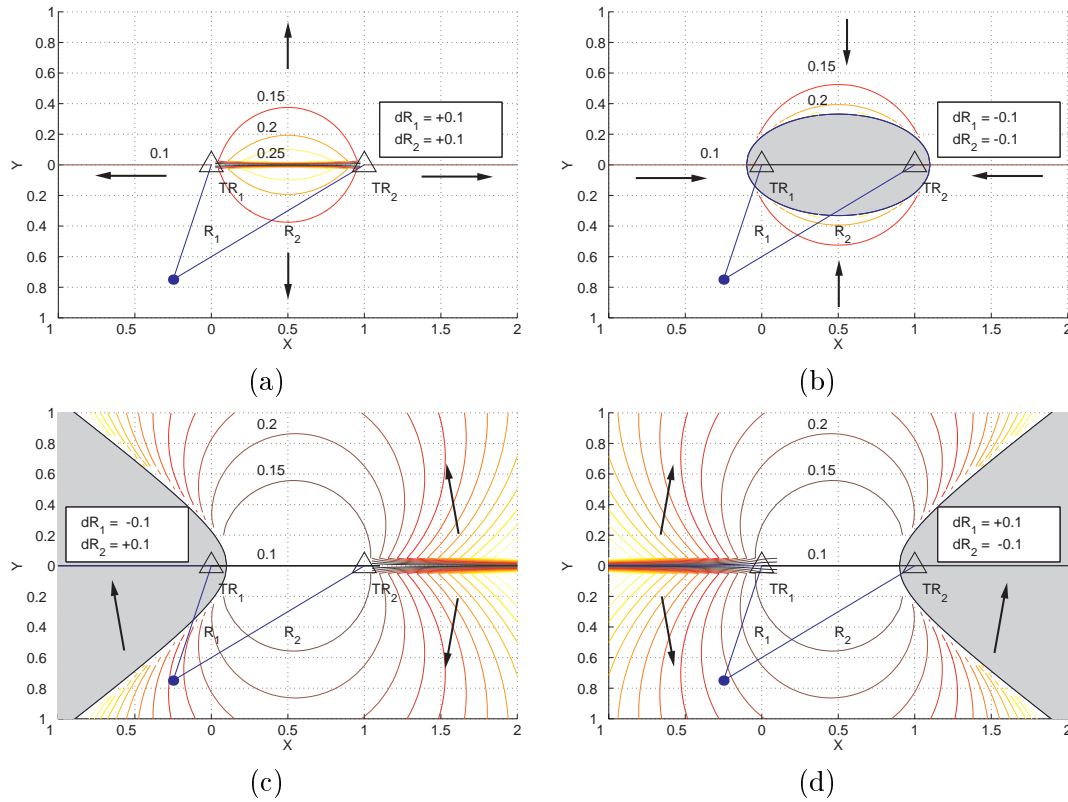


Figure 3.5: 3-Transceiver Lateral Positioning Errors

The magnitude of the resulting positioning errors when bidirectional ranging from two stationary transceivers is used to triangulate the position a third mobile transceiver, assuming that each of the bidirectional range measurements is in error. The stationary transceivers are located at known locations $(0,0)$ and $(1,0)$, and the array is assumed to be 2-dimensional.

Plot (a) shows error contours with respect to the true transceiver position when both range measurements are too long by 0.1 units. The position error ranges from 0.1 units up to a maximum of 0.332 units at the midpoint between the stationary transceivers (The contour spacing is too close to label all of the contours near the center). The arrows show the approximate direction of the position offset. Plot (b) shows similar contours when both range measurements are too small by 0.1 units. The gray ellipse marks the central region in which no triangulation solution exists. Plot (c) and (d) show the results when one range measurement is too long, and the other too short. The region with no solution is now an unbounded hyperbola, and the maximum positioning error is also unbounded. See Figure 3.6 for more detail.

A general set of range/position errors may be roughly approximated as a superposition of these conic sections.

situations.

Figures 3.5c and 3.5d show the worst-case scenerio, wherein one measurement is increased by 0.1 units while the other is decreased. This creates a large unbounded hyperbolic region to one side of the array within which no solution is possible, while in other directions the positioning error magnitude grows without bound as the third transceiver gets farther away from the array. Figure 3.6 shows additional detail of the no-solution region from case (d), this time with range biases of 0.2 units magnitude for additional clarity.

In general, a random set of measurement errors will give positioning errors that can be described by some (nonlinear) superposition of either case (a) or case (b) together with one of case (c) or (d). Because of the nonlinear nature of the geometry, the magnitude of these positioning errors tends to grow faster than the bias errors themselves. These warping characteristics, and especially the degenerate regions without possible triangulation solutions, make it highly desirable to either have more transceivers in the array to give redundant measurements or to blend the GPS transceiver system with other sensors.

Note that regardless of the error in the ranging measurements, it is impossible to tell whether the third transceiver is in the positive or negative y half-plane. If this distinction is important, it must be provided by other sensors or via *á priori* information. This also means that it is impossible to determine whether the mobile transceiver has actually crossed the x -axis using these ranging measurements alone, lending another motivation for the addition of extra transceivers to the array.

3.3.3 N-Transceiver Positioning

When confronted with an arbitrary array of N transceivers, the geometry reconstruction process involves patching together all possible combinations of the triangular solutions to obtain an optimal network of the devices. In general, there will be $\frac{N(N-1)}{2}$ possible ranges between the devices and $\frac{N(N-1)(N-2)}{6}$ possible triangles [32].

Under ideal conditions the range measurements are exact and the determination of the array geometry can be cast into the Euclidean Distance Matrix (EDM) problem. An EDM is a symmetric matrix, with the rows and columns directly relating to the indices of the transceivers, and the off-diagonal terms representing the ranges between the corresponding devices. In the more general case presented by an SCPA, the ranges between the devices are colored both by noise and by range biases, including the integers. This means that the EDM cannot be solved exactly, but must be solved for using an optimization strategy that

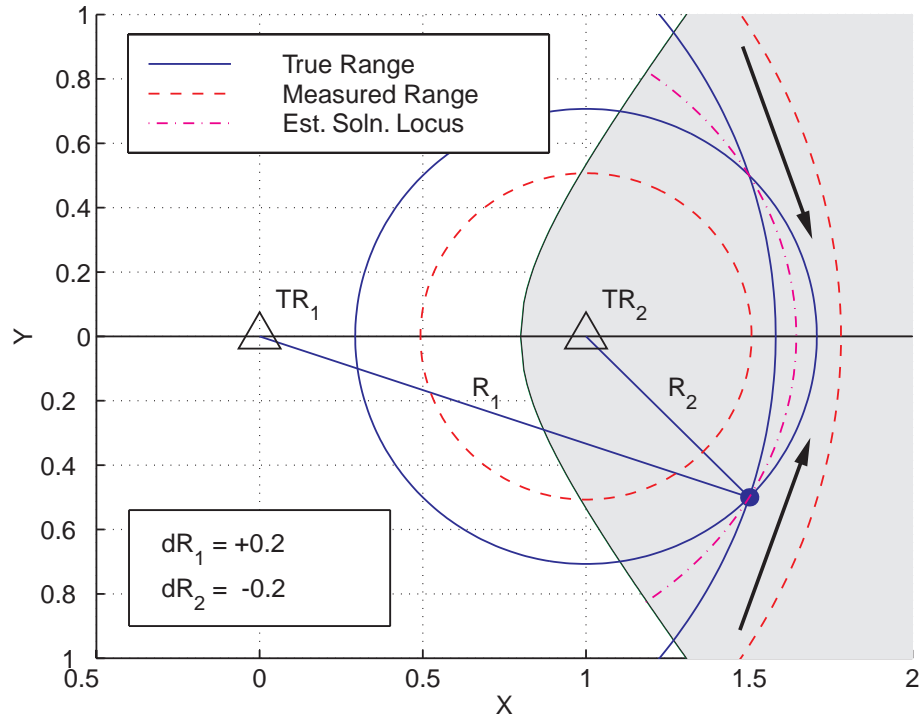


Figure 3.6: 3-Transceiver No-Solution Region

A more detailed examination of the gray no-solution region from Figure 3.5d. The solid circles present the true range to the mobile transceiver from the two stationary transceivers. The dashed circles present the measured ranges, biased by ± 0.2 units. It is clear that the range measurements do not intersect, and that no triangulation solution is possible.

The dash/dot ellipse presents a locus of possible locations that are equidistant from both sets of ranging circles. For all points along this locus, the maximum-likelihood solution for the location of the mobile transceiver — assuming a Gaussian distribution of range errors — is at the point where the locus intersects the x-axis (i.e. where the range measurements are closest). Similar loci exist for other points within the gray region, with the maximum-likelihood solutions always lying along the x-axis.

minimizes the overall system error.

3.3.4 4-Transceiver Positioning

This dissertation does not attempt to solve the generalized EDM problem. Rather, it follows an application-specific approach to solve for the relative positions of the transceivers in the nominal experimental setup. This configuration includes three stationary transceivers arrayed in a triangle, together with a single mobile transceiver. The exact shape of the stationary triangle is not critical, provided that no two of the transceivers are near to each other and they are not collinear, both of which would cause singularities. An equilateral configuration gives the best overall observability of the device locations, so it can be considered the nominal array geometry.

Determination of the transceiver positions follows a cascaded approach. First, the relative locations of the three stationary transceivers are determined using simple triangulation from the bidirectional range measurements between themselves. Second, a rough estimate of the location of the mobile transceiver is determined using triangulation from the stationary transceivers, which are assumed — for the moment — to be in known relative locations. Because triangulation with respect to any two points gives two (mirrored) solutions, this process gives a total of six possible locations from the three pairwise combinations of the stationary transceivers (Figure 3.7). These six solutions are examined to take the three that are most consistent with each other, one from each pair. The initial position estimate of the mobile transceiver is then the mean of these three points.

This triangulation technique provides a good estimate of the location of the mobile transceiver provided that there are no measurement errors. In order to compensate for inconsistent range measurements, the algorithm employed in this research utilizes an additional nonlinear iterated least squares (ILS) technique to refine the position of the mobile transceiver.

The basic nonlinear measurement equation for this case is

$$\begin{bmatrix} r_m^{(1)} \\ r_m^{(2)} \\ r_m^{(3)} \end{bmatrix} = \begin{bmatrix} \left\| \underline{p}_m - \underline{p}^{(1)} \right\| \\ \left\| \underline{p}_m - \underline{p}^{(2)} \right\| \\ \left\| \underline{p}_m - \underline{p}^{(3)} \right\| \end{bmatrix} + \begin{bmatrix} e_m^{(1)} \\ e_m^{(2)} \\ e_m^{(3)} \end{bmatrix} \quad (3.11)$$

or

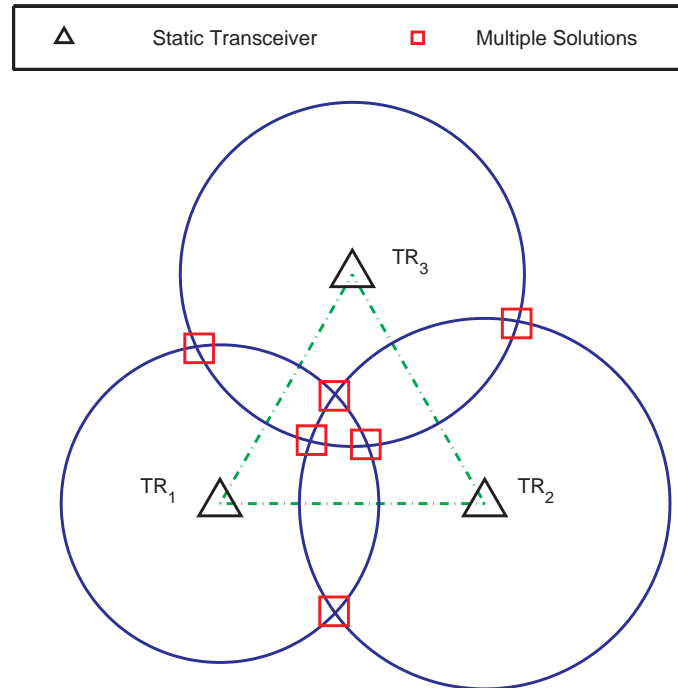


Figure 3.7: Triangulation-Based Position Solutions

Possible solution locations from range-based triangulation from three stationary transceivers in known locations, when the multiple bidirectional range measurements between the mobile and stationary transceivers are inconsistent with each other due to measurement errors. The true location of the mobile transceiver is at the center of the array, and the circles indicate the measured range from each stationary transceiver. Solutions consistent with a pair of measurements are indicated: No solution is consistent with all three measurements.

Note that for a three-dimensional array the surfaces of measured range are spheres, and the intersections are circles.

$$\underline{r} = F\left(\underline{p}_m, \underline{p}^{(1)}, \underline{p}^{(2)}, \underline{p}^{(3)}\right) + \underline{e} \quad (3.12)$$

with

$$\begin{aligned} r_m^{(i)} &\triangleq \text{measured range from stationary transceiver } i \text{ to the mobile transceiver} \\ \underline{p}_m &\triangleq \text{position of the mobile transceiver} \\ \underline{p}^{(i)} &\triangleq \text{position of stationary transceiver } i \\ e_m^{(i)} &\triangleq \text{error on measurement } i \end{aligned}$$

This is solved by linearizing around the estimated position of the mobile transceiver derived from triangulation to give

$$\delta \underline{r}_k = G_k \cdot \delta \underline{p}_{m,k} \quad (3.13)$$

at iteration step k , where

$$G_k \triangleq \nabla F\left(\hat{\underline{p}}_{m,k}, \underline{p}^{(1)}, \underline{p}^{(2)}, \underline{p}^{(3)}\right) = \left. \frac{\partial F\left(\underline{p}_m, \underline{p}^{(1)}, \underline{p}^{(2)}, \underline{p}^{(3)}\right)}{\partial \underline{p}_m} \right|_{\hat{\underline{p}}_{m,k}}$$

is the vector gradient (Jacobian matrix) of F at the estimated position $\hat{\underline{p}}_{m,k}$ and

$$\delta \underline{r}_k = \underline{r}_{meas} - F\left(\hat{\underline{p}}_{m,k}, \underline{p}^{(1)}, \underline{p}^{(2)}, \underline{p}^{(3)}\right)$$

Equation 3.13 can then be solved for $\delta \underline{p}_{m,k}$ by taking the left pseudoinverse of G_k .

$$\delta \underline{p}_{m,k} = G_k^L \cdot \delta \underline{r} = \left(G_k^T G_k\right)^{-1} G_k^T \cdot \delta \underline{r}_k \quad (3.14)$$

A new $\hat{\underline{p}}_{m,k+1}$ is then computed, and the process is repeated until convergence is reached.

The technique employed above does not give a globally optimal solution in the sense that the positions of the stationary transceivers relative to each other are determined via triangulation using only the range measurements between themselves, and these positions are not allowed to change due to the additional ranging information provided by the fourth mobile transceiver. This procedure is advantageous, however, because the measured arrangement of the stationary transceivers is therefore not dependant upon the motion of the mobile transceiver and thus stays constant with time, yielding an immobile base from which

to provide relative navigation to any number of mobile users. It would of course be possible to use this additional ranging information to occasionally refine the estimate of the relative locations of the stationary transceivers, perhaps by applying an average correction after the mobile transceiver has moved to many different locations. Such a refinement is very closely related to the self-calibration methods presented in Chapter 4.

As with the 3-device triangulation described above, positioning with four devices in the presence of measurement errors also produces a warped solution space regardless of the actual algorithm employed. Depending upon the magnitude and direction of these errors and the relative geometry, this warping may or may not be amenable to simple covariance analysis. Figure 3.8 shows the magnitude of the positioning error of a fourth mobile transceiver based upon biased ranging measurements to three stationary transceivers in known locations. The stationary transceivers are arranged in an equilateral triangle 1.0 units per side, and the range biases are permutations of ± 0.1 units. Positioning was performed using the cascaded application of triangulation followed by ILS. As with the 3-device case, positioning errors are much worse for dissimilar directions in the range biases. In plots (a) and (b), with uniform direction biases, the position error is bounded, while in plots (c) and (d), with biases in different directions, the errors outside the array grow without bound.

3.4 Practical Considerations

There are several additional practical considerations that must be taken into account when positioning using an actual transceiver array. The most critical are outlined below.

3.4.1 Vertical Error

The previous discussion has assumed a planar configuration for the transceiver array. With only three transceivers the triangulation process will naturally generate a plane, but it may not coincide with the local horizontal. With four or more transceivers full 3-dimensional positioning is possible, but only if one of the transceivers is located sufficiently out-of-plane for the z-dimension to be observable. In the target application — positioning of a rover on the surface of Mars relative to a ground-based array — this will not be the case.

In such situations it must be assumed that the transceivers are all in a plane, because the vertical dimension is essentially unobservable. Therefore any vertical displacement due

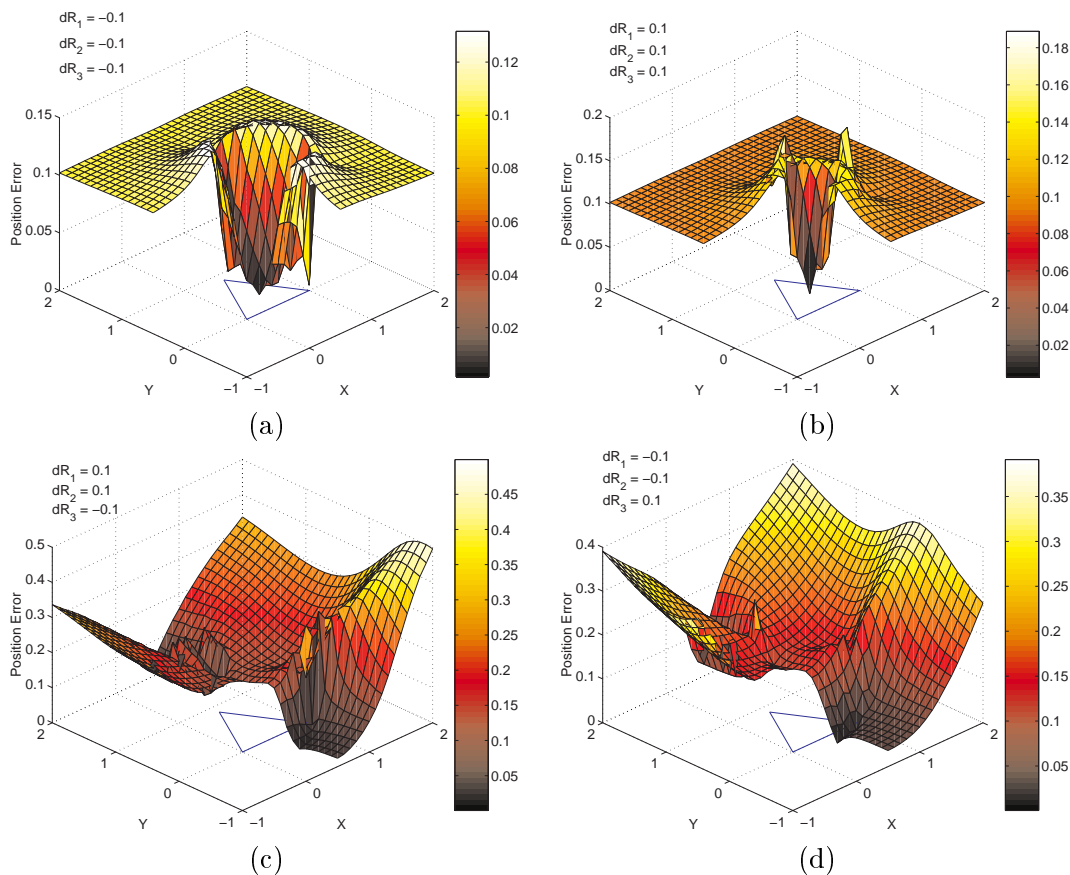


Figure 3.8: 4-Transceiver Lateral Positioning Errors

The magnitude of the resulting positioning errors when bidirectional ranging from three stationary transceivers is used to triangulate the position of a fourth mobile transceiver, assuming that each of the bidirectional range measurements is in error. The stationary transceivers are arranged in an equilateral triangle (see Figure 3.7) at known locations $(0,0)$, $(1,0)$, and $(0.5,0.866)$, and the array is assumed to be 2-dimensional. The estimated position is computed via iterative least squares to minimize the RMS range error to all three stationary transceivers.

The individual plots (a)-(d) show different permutations of range errors, in a similar fashion to Figure 3.5. The near quadrant of each surface is removed to facilitate viewing. As before, range errors of common sign yield bounded position errors, while range errors of differing signs yield unbounded errors. A general set of range/position errors may be roughly approximated as a superposition of these four combinations.

to terrain will result in positioning errors. Fortunately these are of the type presented in Figures 3.5a and 3.8a, in which all of the range errors are positive and therefore the position errors are bounded. In fact, the ranging error resulting from a vertical position error tends to be an order of magnitude smaller than that from an error in the ranging measurement itself. For example, a 10% vertical error results in a horizontal positioning error of only 1.0%. Therefore the errors associated with assuming a true planar configuration tend to be on the order of centimeters to tens of centimeters rather than meters, except in extremely uneven terrain. Although such errors could result in difficulties for very fine navigation, it is important to note that they are repeatable errors and so will be common-mode for multiple users utilizing the same GPS transceiver array for navigation.

3.4.2 Multipath

While multipath errors are limited to a few centimeters for carrier-phase positioning, code-phase multipath errors can be over 100 meters [4]. Although generally much smaller than this, they still pose a significant challenge.

The planar geometry associated with transceiver arrays has both positive and negative effects on the multipath error. On one hand the transceivers are tracking signals nearly parallel with the ground, rather than tracking satellite signals coming from a sizable oblique angle above the horizon. Because the multipath signal is nearly parallel with the direct signal, it cannot be easily rejected via choice of antenna pattern or by using groundplanes. On the other hand this ground-reflection multipath is only delayed by a very short amount by the extra transmission pathlength. Thus this multipath, although potentially strong, is at a very short delay with respect to the direct signal. It should therefore cause only a small error, except in the case when it destructively interferes with the direct signal and causes it to fade.

In practice a much more detrimental form of multipath comes from buildings, cliffs faces, or other vertical reflectors located some distance away from the array. These can generate very strong multipath signals parallel to the ground, but also at path delays sufficient to cause significant ranging error. In addition, this multipath error cannot be eliminated via time averaging when ranging between the stationary transceivers.

Because of the difficulty associated with measuring multipath except in controlled situations, the true impact of code-based multipath on planar transceiver arrays has not been accurately characterized. Although code/carrier divergence such as is typical of multipath

has been observed in the experimental data, more-targeted experiments must be conducted to truly assess the threat of this phenomenon.

3.4.3 Cycle Slip Detection and Correction

Carrier-phase cycle slips (instantaneous, integer-valued jumps in the integrated carrier phase in the receiver tracking loops) and other brief losses of tracking are another common difficulty in GPS positioning. With the presence of additional sensors or redundant measurements the resulting jumps in measured range can usually be eliminated. When using a stand-alone transceiver system without available redundant measurements, however, they become much more disruptive.

The experimental system used in this research (described fully in Chapter 5) relies upon the presumed slow motion of the mobile transceiver to attempt to detect and fix these slips. When using bidirectional ranging, a single cycle slip of 19 cm (one cycle) results in a measured range jump of 9.5 cm. With an update rate of 5 Hz, this translates to a slip-induced displacement rate of only 0.48 m/s. Since the transceivers are frequently moved at up to 1 m/s during testing, and sometimes with reasonably quick accelerations, such short cycle slips are essentially undetectable using the ranging measurements alone. Arrays with greater numbers of transceivers to provide redundant measurements, integration with other sensors, and faster data rates would all help to alleviate this problem, but are outside the scope of this work.

3.4.4 Synchronization

As with all DGPS applications, bidirectional ranging requires that the measurements at the two receivers be latched nearly simultaneously. If the receivers are not properly synchronized, the range measurement between the two transceivers will be in error. The issue is not that the two receivers must register the same absolute time, since the steady-state clock offset $\delta\tau_{r,u}$ is resolved in the double-difference measurement equations. Instead, this is a higher-order effect caused by the integrated frequency difference in the receiver oscillators over the interval between the instants when the receivers latch their respective measurements. Without proper clock synchronization, this latching interval will slowly increase and the measured range will change from its original value.

Ideally the receiver clocks would be perfectly synchronized and would latch their measurements truly simultaneously. In practice, a very small synchronization error yields ranging errors that are still acceptable. The maximum allowable latching interval depends upon the short-term clock stability of the receivers, and can be roughly estimated using the following formula

$$\Delta R \lesssim \left(\frac{\delta f}{f} \right) \cdot \Delta T \cdot c \quad (3.15)$$

where

$$\begin{aligned} \Delta R &\triangleq \text{resulting range measurement error} \\ \frac{\delta f}{f} &\triangleq \text{short-term receiver clock stability} \\ \Delta T &\triangleq \text{difference between sampling times} \\ c &\triangleq \text{speed of light} \end{aligned}$$

For example, for an inexpensive temperature-controlled crystal oscillator (TCXO) with a typical short-term temperature stability of ± 1.0 ppm and a desired range error of one mm, this leads to an acceptable sampling offset of $3.3 \mu\text{s}$, or roughly the duration of three code chips.

The method by which this level of synchronization is achieved in the prototype SCPA is through the use of a designated master pseudolite, a strategy developed by Zimmerman [68]. Each receiver in the array tracks the master pseudolite signal, and latches all of its measurements at the start of the code epoch of the master. Because receiver code tracking is good to about 10 ns, and the maximum transmission time delay for the master pseudolite signal to any receiver over a 100 m array is 330 ns, the maximum synchronization error is roughly 350 ns. The desired range accuracy is therefore easily achievable using the current TCXOs. Greater ranges would result in proportionally larger range errors, which could be countered through the use of more stable oscillators.

Note that using a master pseudolite for receiver synchronization requires that it be broadcasting at all times. If a TDMA pulsing scheme is used in which not all of the transceivers broadcast at all times, the master must be exempted from the scheme. Alternatively, the receivers could be programmed to use different pseudolites as the ‘master’ depending upon which are broadcasting at any given time. Using multiple master pseudolites for different transceiver pairs is also necessary if the array extends over the horizon, in which case a single master pseudolite would not be visible to all of the transceivers.

3.5 Summary

The bidirectional ranging formulation presented in this chapter gives designers of GPS pseudolite transceiver systems considerable new capabilities for the operation of pseudolite arrays. Although the concept of ranging between transponders is by no means new, the application of bidirectional ranging as the sole measurement observable in a GPS-based system is unique to this research. Besides giving improved observability of system states such as transmitter location, this approach also enables operation in sparse arrays that would normally be incapable of CDGPS positioning, and enables TDMA operation to preserve system bandwidth. The greatest disadvantage with this system formulation is the necessity of using GPS transceivers, with the corresponding near/far issues and the added datalink requirement. For extremely difficult applications such as Self-Calibrating Pseudolite Arrays, however, bidirectional ranging is the only solution method known to give adequate system performance.

Within the bidirectional ranging formulation, determination of the relative locations of the GPS transceivers is nominally a simple matter of triangulation. As this chapter has shown however, inconsistencies in the range measurements from noise, multipath, or system biases can result in severe warping of the solution space, to such an extent that the positioning errors are much greater than the measurement errors themselves. While covariance analysis can describe some of this effect, the potentially widespread regions in which no simple triangulation solution is possible necessitate the use of more complicated heuristics for relative positioning in the presence of large biases and other errors. Carrier-phase integer ambiguities present a similar error source that must be resolved before a transceiver array may be used for centimeter-level positioning. The elimination of these system biases is the object of the array self-calibration process, which is described in detail in the next chapter.

Chapter 4

Self-Calibration Methods

The positioning methods discussed in the previous chapter rely upon the assumption that all of the range measurements between the transceiver pairs are reasonably accurate and that the stationary transceivers are in known locations. Upon initial deployment of a Self-Calibrating Pseudolite Array (SCPA), this is not the case. Code-phase multipath causes errors in the location estimates of the stationary transceivers, which in turn result in incorrect determination of the carrier-phase integers between the transceivers. This chapter describes in detail the algorithms and methods used to eliminate these two primary error sources, thereby enabling centimeter-level navigation. It is these self-calibration methods that make the autonomous deployment of pseudolite arrays a practical option. Once the array self-calibration is successfully completed, accurate navigation of multiple mobile users may be accomplished by using the transceiver-based positioning techniques described in Chapter 3, or alternatively by using standard pseudolite-based CDGPS with respect to a stationary reference receiver.

This work on SCPAs has developed several innovative strategies for the self-calibration process. The utilization of GPS transceivers instead of separate pseudolites and receivers gives an extra degree of system state observability that is critical for successful self-calibration. In addition, the algorithms themselves present useful extensions to the common practice of GPS navigation. Of special interest is the new Quadratic Iterative Least Squares (QILS) algorithm, which offers a significant performance improvement over the linear methods commonly used in GPS applications. Used together with a stochastic multiple-seeding approach, QILS demonstrates a 100% self-calibration success rate under nominal array configurations, and a 99.80% success rate under worst-case error conditions. The details of the

these algorithm performance metrics appear towards the end of this chapter.

4.1 Calibration Requirements

There are many error sources that can affect the raw range measurements used for positioning in a pseudolite array. Some of those that are of concern in satellite-based GPS, such as atmospheric delays, are negligible over the short ranges associated with most pseudolite applications. Others, such as multipath, are severe problems that can be difficult to eliminate. This section examines three of the error sources that can be eliminated through careful system calibration:

- Unknown carrier-phase cycle (integer) ambiguities
- Unknown locations of the broadcast elements
- Unknown line biases and delays in the transceiver hardware

The first two of these error sources comprise the greatest effect on positioning accuracy within the pseudolite array, and successful determination of these parameters allows one to accurately determine the trajectories followed by mobile elements within or in the vicinity of the array.

A Self-Calibrating Pseudolite Array is able to calibrate itself autonomously in order to successfully eliminate these error sources, thereby determining the locations of the GPS transceivers to centimeter-level accuracy. Because understanding the sources and significance of these errors is critical towards understanding the SCPA self-calibration methodology itself, this section briefly describes these error sources and summarizes the methods used to compensate for their effects in standard GPS applications. The self-calibration algorithms employed in an SCPA are based in part upon these common methods, but extend them in significant ways: These modifications are described in detail in the subsequent sections of this chapter.

4.1.1 The Integer Ambiguity Problem

Determination of the carrier-phase cycle ambiguities $\nabla \Delta \mathcal{N}_{r,u}^{(k,l)}$ (Equations 3.3, 3.5, and 3.10) is necessary for any carrier-phase GPS application. There are three general techniques used

to determine these integers. This section contains a brief summary of each and discusses their applicability to an SCPA.

Integer Search Methods

The brute-force method is to take advantage of the integer nature of the bias and perform a comprehensive search through all possible integer combinations, and then choose the integers that give the smallest residual for the given cost function. Even though the double-differencing process greatly reduces the number of possible integer combinations, the computational burden is still prohibitive for all but the shortest baselines between receivers.

Frequency Methods

A second solution method is multiple-frequency Cascaded Integer Resolution (CIR), which is also called ‘widelaning.’ The GPS satellite signal does not actually consist only of a single carrier at L1 (1575.42 MHz), but also includes another signal at L2 (1227.6 MHz). The code modulated upon the L2 signal (P-code) is unusable except to users with access to a special encryption key. The L2 carrier, on the other hand, is readily trackable by dual-frequency receivers. Beating these two signals against each other yields a new widelane signal with frequency $L1-L2 = 347.84$ MHz and a resulting wavelength of 86.2 cm.

The basic technique in CIR is to start with a rough range measurement based upon C/A code, and use that to determine the correct integer in the widelane signal. This is easier to do with the widelane (WL) signal than with the raw L2 carrier because of the longer wavelength. Once the WL integer is known, range can be determined to within a few centimeters, giving enough information to determine unambiguously the L2 integer value. Because the jump from C/A code accuracy ($\sim 2-3$ m) to the widelane integer (86.2 cm) is difficult, other frequencies are needed as well to fill in the gap and generate small enough jumps to be made unambiguously. For example, Jung et al. present a CIR strategy using the newly created civil GPS signal at L5 (1176.45 MHz) to create an Extra Widelane (EWL) beat signal of 5.86 meter wavelength to bridge this gap [26].

Because a Mars-based pseudolite system can be created from scratch without regard to existing system infrastructure, the frequency plan can potentially be chosen such that unambiguous CIR is the only method needed for integer determination. This would be especially useful with bidirectional ranging, since each integer affects only a single pairwise measurement. As the new GPS civil frequencies become widely available, it may even be

possible to do this cheaply using commercial off-the-shelf (COTS) components. In the near-term, however, multi-frequency receivers and pseudolites are more expensive and difficult to use, and so it is still desirable to have an integer determination method for single-frequency use.

Motion-Based Methods

The final technique for integer determination is to use relative motion of either the transmitters or the receivers. As the devices translate with respect to each other, the measured ranges change. If there is sufficient motion and redundant measurements, then it eventually becomes apparent that only one set of integers is consistent with the measured range changes among the devices during this motion. This technique is used by surveyors by placing a receiver at a fixed location and waiting for the GPS satellites to move until the integers (and hence the receiver position) are very accurately known. In CDGPS systems — both satellite and pseudolite — relative motion between the receiver of interest and the differential reference station can be used to determine the relative integers very quickly. This has been done to resolve integers in a wide variety of systems including autonomous tractors [45], the Integrity-Beacon Landing System (IBLS) [7], and formation-flying spacecraft [68]. It has also been used to determine attitude in multiple-antenna vehicles such as an autonomous airplane [43] and a robot helicopter [9].

Because motion-based integer determination does not require expensive multiple-frequency receivers and pseudolites, it is the technique employed in the self-calibration formulation presented in this thesis.

4.1.2 Pseudolite Locations

The second major set of parameters that must be determined before operation of a pseudolite array is possible is the location of the transmitting elements. In satellite-based GPS, this information is known and is broadcast to the users on the satellite data message. For pseudolite systems, they must be determined through other means.

Existing localization methods vary widely in scope, complexity, and accuracy. Zimmerman's pseudolite array was originally surveyed by hand using a tape measure [68]. Later, both Zimmerman [68] and Kee [30] demonstrated more-precise surveying by utilizing receivers placed at precisely known locations, thereby inverting the problem to solve for transmitter instead of receiver locations. Other methods employing theodolites and laser

range finders are also viable, although angular measurement errors may become unacceptable for long baselines. There may also be difficulty in determining the actual phase center of the GPS antennas when using other calibration systems, so it is generally beneficial to use the radio-frequency (RF) components of the pseudolite system to survey in the array whenever possible.

When using GPS transceivers instead of isolated receivers and pseudolites, range measurements exist between each pair of devices. As was discussed in Chapter 3 this means that code-level calibration of the array may be done instantaneously, yielding a coarse meter-level survey. Most applications will benefit from a more accurate carrier-phase survey, however.

The primary purpose of carrier-phase surveying is to determine the carrier-phase cycle ambiguities present in the single- or double-difference solutions. Once all of the integer values are known the ranges between the transceivers are likewise known, and the locations can then be determined using the methods of Chapter 3. The problems of determining the integers and surveying in the transceiver locations are therefore identical, at least in concept. In practice, it is difficult to cast the self-calibration problem into an all-integer or all-position formulation because direct integer determination between stationary transceivers is difficult without relative motion. The self-calibration algorithms presented in this thesis therefore retain a mix of position and integer states. Analysis of algorithm effectiveness, however, may still be reduced to the correct determination of the integers themselves.

4.1.3 Line Biases

A third set of unknown system parameters, although of secondary importance, are the line biases present in the RF signal paths in the system hardware. These result from the signal propagation delay in the cables, amplifiers, and connectors within the transceivers themselves. Although nominally constant, these delays do vary slightly with time, primarily due to temperature fluctuations. In a portable system that involves frequent assembly and disassembly of RF signal cables, such as the one utilized for these experiments (described fully in Chapter 5), the line biases will change with every tightening or loosening of the connectors. Line biases tend to be on the order of centimeters, so they primarily affect CDGPS systems.

Line biases may be determined by placing both the broadcast and receive elements in known locations and observing the discrepancy between the predicted and measured carrier phases. They may also be determined using integer determination techniques by focusing

on the non-integer remainders as opposed to the integers themselves. In practice many floating-point integer determination methods simply lump together the integer ambiguities and the line biases, solving for them simultaneously and never separating them thereafter. This is the technique employed in this dissertation.

4.2 Self-Calibration Problem Formulation

The array self-calibration for an SCPA utilizes a cascaded, multi-layer approach, which is presented in Figure 4.1. First a rough estimate of the the starting configuration of the array, including the starting location of the mobile transceiver, is generated using averaged code-range measurements. This step includes the bidirectional ranging between the stationary transceivers, and is used to set the initial integer estimates for the carrier-based ranging measurements. Second, the rover carrying the mobile transceiver starts a trajectory around and/or between the stationary transceivers, collecting carrier-range measurements periodically. The rover position along the trajectory is estimated using these measurements. Finally, the carrier-range data collected during the trajectory are utilized to generate a new estimate of both the locations of the stationary transceivers and the actual track followed by the mobile transceiver. This also generates an estimate of the corresponding carrier-phase integers/line biases. The algorithms used to perform this self-calibration step are described in detail below.

4.2.1 Complete State Estimation

In most conventional GPS applications the locations of the broadcast elements — either satellite or pseudolite — are known. If they are not known, such as during a surveying process, then the location of the receivers are usually known instead. This divide-and-conquer approach generally yields solution formulations that are easily linearizable and present few solution difficulties, and also greatly reduces the potential size of the solution space. In contrast, the determination of an independent pseudolite array from scratch cannot leverage off this *á priori* information, yielding a more complicated and difficult solution process.

Self-consistent solutions of the complete state of a system have been performed many times in the past in a variety of fields. In GPS applications, a process called orbital relaxation is used to refine the estimate of the satellite locations at the same time that it

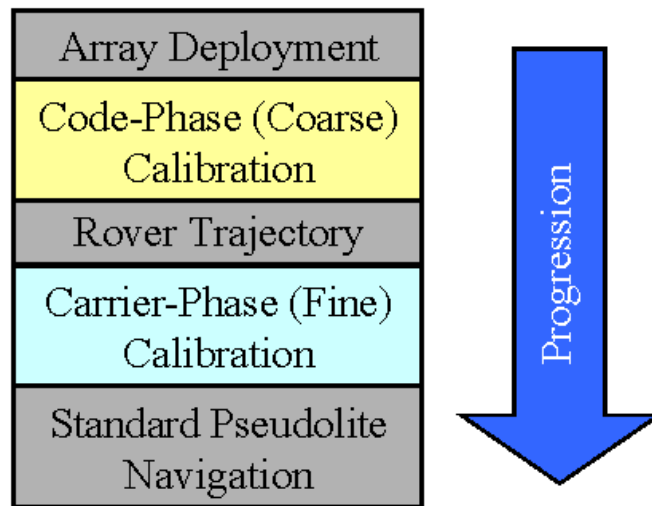


Figure 4.1: Self-Calibration Procedure

The self-calibration procedure used for determining the true locations of the devices comprising an SCPA. Calibration is conducted at two levels: code-phase gives meter-level accuracy, and carrier-phase gives centimeter-level accuracy. Following successful self-calibration, the SCPA may be used like any standard pseudolite array.

solves for the location of the reference receivers on the ground [34]. In this case the initial estimates tend to be quite good, and only small corrections are sought. The system is thus easily linearizable and the solution is well-behaved. Another related application is Simultaneous Localization and Mapping (SLAM), wherein a robot maneuvers in an area and generates a map of its surroundings and simultaneously determines its location with respect to that map using a host of onboard sensors, generally including vision, inertial measurements, and/or dead reckoning. Many SLAM applications are linear, although this is not a necessity. An example of a nonlinear SLAM algorithm applied to underwater robotics can be found in [20].

There are a couple of important differences that make the self-calibration problem presented here more difficult than the preceding applications. The geometry is much more nonlinear than is present in either orbital relaxation or most SLAM applications, making linearization difficult and inaccurate. This is because of the near-field geometry and the relatively large size of the measurement errors. In addition, the initial estimates of the position states of interest are relatively poor, compounding the problem of the nonlinearities.

Another challenge in the self-calibration problem is that none of the position states of interest is directly observable: Instead, each must be inferred from a time-history of highly-biased range measurements between the transceivers. Ideally an SCPA would be augmented with other sensors such as vision or inertial systems, with these providing complete (although poor) state estimates at all times and the GPS signals providing the accuracy and stability. The goal of this research, however, is to determine the best performance of a stand-alone SCPA without additional sensor augmentation. It will be shown that the techniques in this chapter are adequate for array self-calibration without additional sensor augmentation; system performance can only improve with the presence of additional sensors.

4.2.2 Batch Estimation

Most estimation processes involving mobile robots use Kalman filters. Because the system states are only partly observable and are related in a highly nonlinear manner, however, and because the initial state estimates are so poor, it is difficult to generate an extended Kalman filter (EKF) for this application that converges reliably. Array self-calibration in the absence of additional sensors can therefore best be accomplished through a batch process by collecting range data from many different points along the rover trajectory

and then solving for the positions of the stationary transceivers, the path of the mobile transceiver, and the corresponding integer/line biases simultaneously. This batch process is very effective because the system states, although not completely observable at any given instant, are completely observable in an aggregate sense following an adequate traversal of the array.

A simple analysis of the number of available measurements and the number of unknown system states sets the requirements for both the array size and the number of trajectory sample points. The variables of interest are

$$\begin{aligned} N &\triangleq \text{Number of stationary transceivers} \\ \mathcal{D} &\triangleq \text{Number of dimensions (2-D or 3-D)} \\ S &\triangleq \text{Number of trajectory sample points} \end{aligned}$$

Because the system employs GPS transceivers capable of bidirectional ranging, there is no need to determine the relative clock biases between the pseudolites and receivers as part of the self-calibration process. Moreover the measurements are geometrically more closely related to the states of interest than is possible using near-field (very nonlinear) single- or double-differencing techniques. This greatly simplifies the estimation process by reducing the number of both states and measurements, and eliminating one level of calibration nonlinearities.

Utilizing this transceiver-based architecture, there are $N \cdot S$ available independent range measurements between the mobile and the stationary transceivers along the trajectory, while the total number of states that must be resolved is given by

Positions of stationary transceivers:	$\mathcal{D} \cdot N$
Location of trajectory sample points:	$\mathcal{D} \cdot S$
Integer/Line Biases:	N
Total:	$\mathcal{D}(N + S) + N$

Note that the ranges between the stationary transceivers themselves are not included among the measurements, nor are the associated range biases. This is because these ranges are constant except for phase noise, and there is no associated geometry change during the trajectory to make the biases observable. It is therefore necessary to calibrate the locations of the stationary transceivers without benefit of the cross-link range measurements, except as an initial estimate.

It is also necessary to constrain the system through application of a coordinate system,

	2-D	3-D
N	S	S
2	N/A	N/A
3	6	N/A
4	5	6
5	4	5
6+	4	4

Table 4.1: Self-Calibration Sample Point Requirements

The minimum number of trajectory sample points S necessary for array self-calibration, as a function of the number of stationary transceivers N and the dimension of the array. Situations in which it is impossible to self-calibrate the array are noted by N/A.

because the range measurements between the transceivers are capable of providing only relative positions: The array as a whole is able to translate and rotate freely in space. The number of additional constraints is three for 2 dimensions (two translations and one rotation) or six for 3 dimensions (three translations and three rotations).

For successful array self-calibration the number of measurements must exceed the number of unknowns. The required number of transceivers for 2-dimensional and 3-dimensional SCPAs are therefore given by the following inequalities:

$$N \cdot S \geq 2(N + S) + N - 3 \quad (2\text{-D})$$

$$N \cdot S \geq 3(N + S) + N - 6 \quad (3\text{-D})$$

Table 4.1 presents the minimum sample requirements for a given number of stationary transceivers for both 2- and 3-dimensional arrays. The samples points must be reasonably spaced, and singularities and degenerate geometries must be avoided for successful self-calibration to occur. Even then successful self-calibration is not assured. As a practical matter, most trajectories that are effective utilize more than the minimum number of samples.

4.3 Linear Iterative Least Squares

The basic solution method for array self-calibration is linear Iterative Least Squares (ILS). Starting from the nominal estimate of array configuration and trajectory, the nonlinear

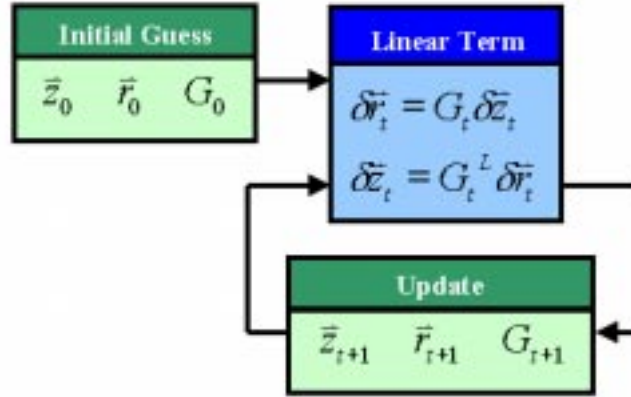


Figure 4.2: Linear Iterative Least Squares Process

The basic flow of the standard linear Iterative Least Squares (ILS) algorithm. The initial estimate of the system state is represented by \underline{z}_0 , and the corresponding ranges between the transceivers by \underline{r}_0 . The system is linearized about this estimate, yielding the linearized geometry matrix G . This matrix is inverted using the left pseudo-inverse and multiplied by the difference between the estimated and measured range, yielding a state correction term $\delta \underline{z}$ that is then added to the initial state estimate. This process is repeated until it converges to a steady final estimate. The effect of linear ILS is to minimize the RMS measurement error residual $\delta \underline{r}^T \delta \underline{r}$, assuming that it converges to the correct system state and not a local minimum.

equations describing the dependence of the measurements upon the system states are linearized about that estimate to determine the local gradient. The difference between the measured ranges and the expected ranges generates a displacement vector with respect to the gradient, resulting in a new estimate of the array configuration. The measurement equations are again linearized about this new point, and the process is repeated until convergence is achieved. Figure 4.2 shows this process in block diagram form. Although ILS is a standard technique, the method with which it is applied to the self-calibration process bears examination. Additionally, the ILS process provides necessary background information and notation for the more sophisticated quadratic solution algorithm.

The ILS self-calibration algorithm operates as a batch process on range data collected over the course of the rover trajectory. At each sample point along the trajectory, the states of interest are the positions of the N stationary transceivers, the integer/line biases

in the ranges from the stationary to the mobile transceivers (all constant), and the current position of the mobile transceiver. The measurements of interest are the instantaneous ranges between the stationary and the mobile transceivers. To contain these variables, the instantaneous state vector \underline{z}_s and the instantaneous range measurement vector \underline{r}_s are defined below.

$$\begin{aligned} \underline{z}_s &\triangleq \left[b_m^{(1)} \quad \dots \quad b_m^{(N)} \quad \underline{p}^{(1)T} \quad \dots \quad \underline{p}^{(N)T} \quad \underline{p}_{m,s}^T \right]^T \\ \underline{r}_s &\triangleq \left[r_s^{(1)} \quad \dots \quad r_s^{(N)} \right]^T \end{aligned}$$

- $b_m^{(i)} \triangleq$ Range bias between mobile transceiver and stationary transceiver i
- $\underline{p}^{(i)} \triangleq$ Position of stationary transceiver i
- $\underline{p}_{m,s} \triangleq$ Position of mobile transceiver at point s
- $r_s^{(i)} \triangleq$ True range between mobile transceiver and stationary transceiver i at point s

Note that the ranges between pairs of *stationary* transceivers are not included in the measurement vector because they are both constant and biased, and therefore do not contribute to the observability of the system states: The positions of the stationary transceivers must be deduced entirely from the range measurements to the mobile transceiver.

The relationship between the system states and the measurements is given by Equation 3.11, which is repeated below with the modification that the collective measurement error \underline{e}_s has now been replaced by the constant biases $b_m^{(i)}$ — which are to be estimated and are thus part of the system state — and the zero-mean random vector \underline{v}_s . In addition, the formulation is expanded to include N transceivers.

$$\underline{r}_s \equiv F(\underline{z}_s) = \begin{bmatrix} \left\| \underline{p}_{m,s} - \underline{p}^{(1)} \right\| + b_m^{(1)} \\ \vdots \\ \left\| \underline{p}_{m,s} - \underline{p}^{(N)} \right\| + b_m^{(N)} \end{bmatrix} + \underline{v}_s \quad (4.1)$$

At any given step k in the iteration process it is necessary to have estimates of the transceiver positions and resulting range measurements. The initial position estimate of the locations of the stationary transceivers can be estimated using averaged code-range measurements and triangulation. The initial estimate of the rover trajectory is computed using the triangulation methods of Chapter 3, beginning with averaged code-range at the starting point and then reverting to changes in measured carrier-range for subsequent sample points along the trajectory. At each sample point s the estimated ranges are computed as

$$\hat{\mathbf{r}}_{s,k} \equiv F(\hat{\mathbf{z}}_{s,k}) = \left[\begin{array}{c} \|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(1)}\| + \hat{b}_m^{(1)} \\ \vdots \\ \|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(N)}\| + \hat{b}_m^{(N)} \end{array} \right]_k \quad (4.2)$$

The initial position and range estimates are

$$\hat{\mathbf{z}}_{s,0} = \left[0 \quad \dots \quad 0 \quad \hat{\underline{\mathbf{p}}}^{(1)T} \quad \dots \quad \hat{\underline{\mathbf{p}}}^{(N)T} \quad \hat{\underline{\mathbf{p}}}_{m,s}^T \right]^T \Big|_0 \quad (4.3)$$

$$\hat{\mathbf{r}}_{s,0} = \left[\begin{array}{c} \|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(1)}\| \\ \vdots \\ \|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(N)}\| \end{array} \right]_0 \quad (4.4)$$

because there is no á priori information available about the biases $\hat{b}_m^{(i)}$.

In order to refine these initial estimates, the measurement equation $F(\mathbf{z}_s)$ is linearized about the position estimate to give the local Jacobian of the system at the given sample point

$$G_{s,k} \triangleq \nabla F(\hat{\mathbf{z}}_{s,k}) = \left[\begin{array}{ccccc} & \frac{\partial \hat{r}_s^{(1)}}{\partial \hat{\underline{\mathbf{p}}}^{(1)}} & \dots & 0 & \frac{\partial \hat{r}_s^{(1)}}{\partial \hat{\underline{\mathbf{p}}}_{m,s}} \\ I_N & \vdots & \ddots & \vdots & \vdots \\ & 0 & \dots & \frac{\partial \hat{r}_s^{(N)}}{\partial \hat{\underline{\mathbf{p}}}^{(N)}} & \frac{\partial \hat{r}_s^{(N)}}{\partial \hat{\underline{\mathbf{p}}}_{m,s}} \end{array} \right]_k \quad (4.5)$$

$$\left. \frac{\partial \hat{r}_s^{(i)}}{\partial \hat{\underline{\mathbf{p}}}^{(i)}} \right|_k = - \left. \frac{\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(i)}}{\|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(i)}\|} \right|_k$$

$$\left. \frac{\partial \hat{r}_s^{(i)}}{\partial \hat{\underline{\mathbf{p}}}_{m,s}} \right|_k = \left. \frac{\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(i)}}{\|\hat{\underline{\mathbf{p}}}_{m,s} - \hat{\underline{\mathbf{p}}}^{(i)}\|} \right|_k$$

such that

$$\delta \mathbf{r}_{s,k} = G_{s,k} \cdot \delta \mathbf{z}_{s,k} \quad (4.6)$$

with the corresponding range measurement and state-update equations

frame such that transceiver #1 is located at the origin (0,0) of the coordinate frame and transceiver #2 is along the x-axis ($y_2 = 0$). The corresponding entries in the global state vector are eliminated, along with the corresponding columns of the global state matrix.

At this point the update equations are now solvable using the standard ILS iteration equations.

$$\delta \underline{r}_k = \underline{r}_{meas} - \hat{\underline{r}}_k \quad (4.10)$$

$$\delta \underline{z}_k = \mathcal{G}_k^L \cdot \delta \underline{r}_k \equiv \left(\mathcal{G}_k^T \mathcal{G}_k \right)^{-1} \mathcal{G}_k^T \cdot \delta \underline{r}_k \quad (4.11)$$

$$\hat{\underline{z}}_{k+1} = \hat{\underline{z}}_k + \delta \underline{z}_k \quad (4.12)$$

Iteration is continued either until the 2-norm of the measurement residuals $\delta \underline{r}_k^T \delta \underline{r}_k$ drops below some specified bounds, until it grows above another specified limit indicating solution divergence, or until some maximum number of allowed iterations is reached.

4.4 Monte-Carlo Simulation

It is necessary to have a reliable method for determining the success of the ILS self-calibration algorithm. With conventional satellite-based GPS the observability of all states of interest is frequently measured using the Geometric Dilution of Precision (GDOP), which is a measure of how ranging errors translate into final receiver positioning errors. If the measurement equation is

$$\delta \underline{r} = G \cdot \delta \underline{z} \quad (4.13)$$

then the DOP for an individual parameter (or GDOP for the entire array) is given by

$$\begin{aligned} iDOP &= \sqrt{A_{ii}} \\ GDOP &= \sqrt{\text{trace}(A)} \end{aligned} \quad (4.14)$$

where

$$A \equiv \left(G^T G \right)^{-1}$$

Multiplying the DOP for a given parameter by the raw range-measurement standard deviation gives the expected standard deviation for the desired state [53].

Examining the DOP for satellite-based systems is an extremely useful evaluation tool because it is easy to compute and shows the accuracy with which any given parameter can be computed. DOP is somewhat less useful for evaluating SCPA configurations, however, for two reasons. First, the number of unknown parameters required to represent a SCPA configuration and trajectory runs from tens to hundreds, as opposed to the four unknowns for a satellite GPS solution. This large number makes the interaction between parameters difficult to understand and visualize, especially when they are strongly cross-coupled. Second and more important, DOP is only truly descriptive of linear — or nearly-linear — systems. In an SCPA with large initial measurement biases, the system can be nonlinear enough that DOP ceases to become a useful parameter for determining algorithm effectiveness. DOP can still be useful to determine the expected final accuracy, however, if the algorithm does in fact converge to the correct configuration.

Because of the large number of parameters and the highly nonlinear nature of the SCPA, Monte-Carlo simulations become the most useful method for evaluating algorithm performance and convergence. This chapter will present results from several such simulations to show the performance of various self-calibration algorithms under a variety of conditions.

4.4.1 Simulation Description

Figure 4.3 shows the standard configuration used for the Monte-Carlo simulations. The stationary transceivers are arrayed in a triangle and the mobile transceiver then moves in the looping trajectory shown. Range measurements are taken at each of the sample points indicated, and are corrupted by measurement biases of uniform magnitude but random sign. These biases affect the ranging between the stationary transceivers as well the rover, so the initial estimate of their locations will be in error as well. Therefore the unknown measurement bias maps directly to the error in the initial position estimate: hereafter the terms are used interchangeably to refer to the same property. Although multipath and other biases would not actually be of the same magnitude between all transceiver pairs in an operational SCPA, using the same bias magnitude for all pairs in the simulation makes it easier to isolate the effect of the bias magnitude. Typically between 20 and 100 tests with different bias directions are run for each configuration to give a statistical metric with which to assess the algorithm performance. Because of the one-to-one mapping between the

true ranges and the actual array geometry, successful determination of these range biases results in an accurate determination of the array geometry and the rover trajectory.

Many different parameters can be changed in these simulations. The most useful include the location of transceiver #3, the radius of the trajectory curves, the spacing of the sample points, and the range-measurement bias magnitude.

4.4.2 Simulation Results

Figure 4.4 presents simulation results from a 6250-run Monte-Carlo simulation of the ILS self-calibration algorithm with respect to both the trajectory curve radius and the range-measurement bias magnitude. All units are normalized with respect to the spacing between transceivers #1 and #2, so a bias of magnitude 1.0 gives range measurements in error by the full size of the array. Successful convergence is determined if the final estimated range biases and the positions of the stationary transceivers are within a small distance (a few centimeters) of the actual values. Failure can occur either because the iteration process diverges or because it converges to a false local minimum, the latter being the more likely failure mode.

From the plot it is apparent that with the trajectory passing only a small distance ($> 20\%$) outside the array, the ILS algorithm converges to the correct array geometry and rover position nearly 100% of the time for range biases less than 10% of the array size. For example, for a 100 meter array the algorithm would perform successfully with range biases of 10 meters, well within the accuracy of nominal code-phase ranging. For a somewhat smaller array the margins start becoming low enough to cause concern, however, especially in the presence of potential multipath errors greater than several meters; a 50 meter array with 10 meter range biases due to multipath would yield successful convergence only 80% of the time. With biases greater than 20%, success drops dramatically. This margin is further reduced if the rover passes too close to one of the transceivers, a potential possibility before the array is surveyed and the locations of the transceivers are truly determined.

4.5 Quadratic Iterative Least Squares

The linear ILS algorithm begins to fail with large initial biases because the initial state estimate becomes far removed from the relatively narrow range in which the system linearization is valid. In these situations it is therefore much more vulnerable to falling into

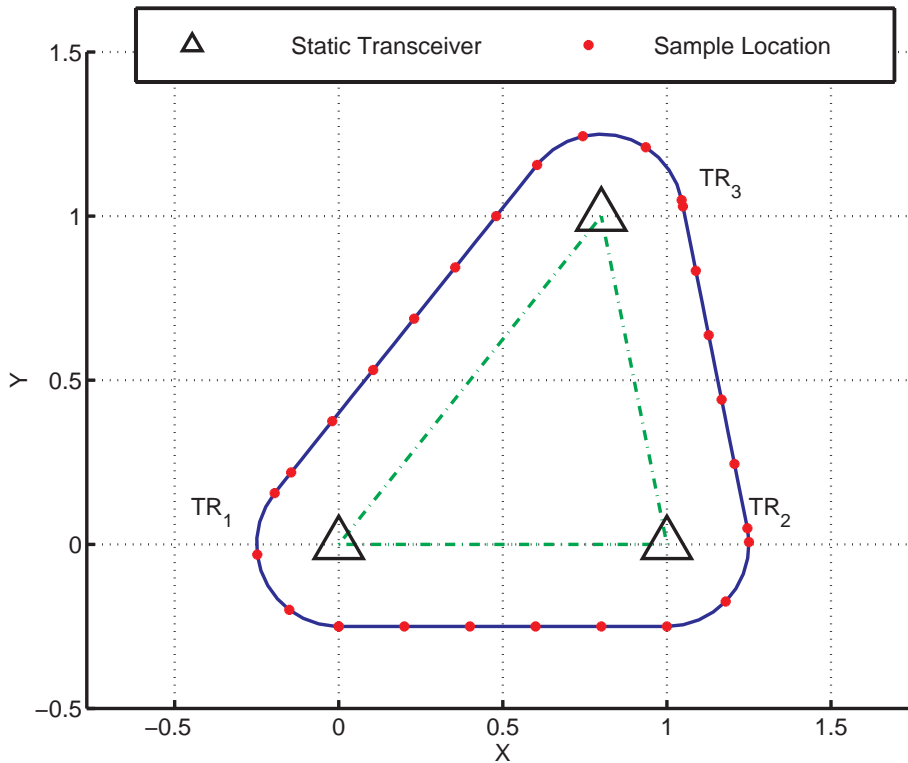


Figure 4.3: Monte-Carlo Simulation Configuration

The simulation configuration consists of three stationary transceivers together with a fourth mobile transceiver that makes a looping trajectory around the outside of the array with constant radius turns with respect to the stationary transceivers. Transceiver #1 is located at point $(0,0)$ and transceiver #2 is located at $(1,0)$. Transceiver #3 is nominally located at $(0.5,1.0)$ forming a nearly-equilateral triangle, although its location can be changed to alter the shape of the array as is shown.

The actual trajectory followed varies with the locations of the stationary transceivers. Range measurements between the stationary transceivers and the mobile transceiver are taken periodically: nominal sampling is every 0.2 units along the trajectory, although this varies at the curve interfaces.

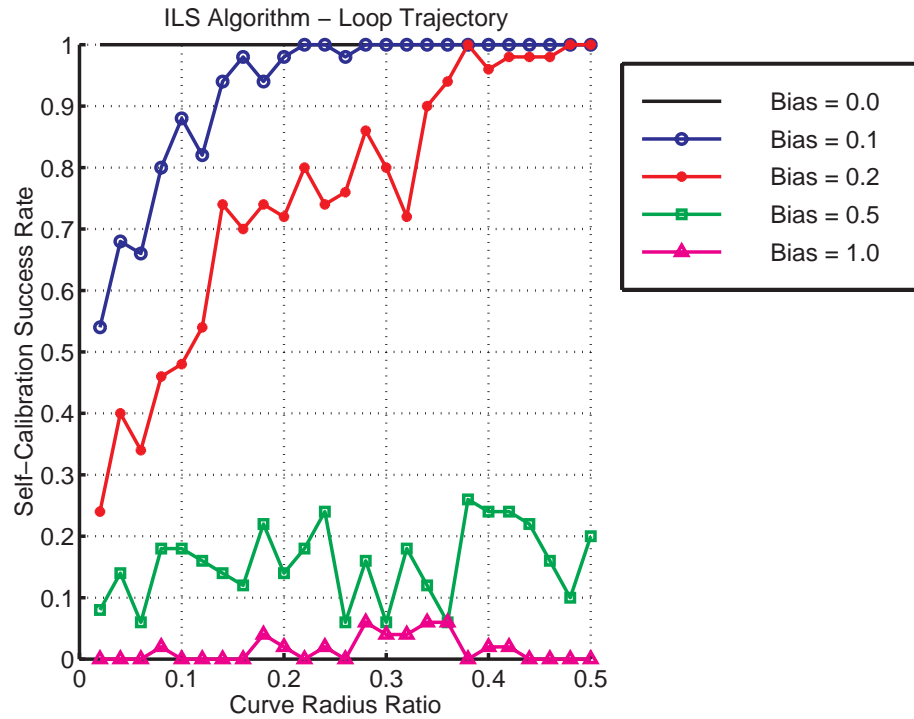


Figure 4.4: Linear ILS Algorithm Success Rate

The success rate of the linear ILS algorithm as a function of the trajectory curve radius and the magnitude of the range-measurement bias. The array configuration and trajectory are as shown in Figure 4.3, with transceiver #3 located at point (0.5,1.0). A successful self-calibration is one that correctly determines all of the system states of interest: the locations of the stationary transceivers, the trajectory of the mobile transceiver, and the carrier-phase integer biases.

local minima or to divergence. Because larger biases are possible in high-multipath environments, it is highly desirable to have an improved algorithm that is better able to handle the system nonlinearities. The method developed in this research still involves an iterative process of making local approximations around the state estimate and then descending the gradient towards the minimum residual. Rather than merely linearizing, however, second-order terms are retained as well.

The second-order expansion of range between the mobile transceiver and one of the stationary transceivers is

$$\delta r_s^{(i)} = \nabla F \left(b_m^{(i)}, \hat{\underline{p}}^{(i)}, \hat{\underline{p}}_{m,s} \right) + \nabla^2 F \left(b_m^{(i)}, \hat{\underline{p}}^{(i)}, \hat{\underline{p}}_{m,s} \right) \quad (4.15)$$

The linear component — which was used earlier in the linear ILS formulation — is

$$\nabla F \left(b_m^{(i)}, \hat{\underline{p}}^{(i)}, \hat{\underline{p}}_{m,s} \right) = \begin{bmatrix} 1 & \frac{\partial \hat{r}_s^{(i)}}{\partial \hat{\underline{p}}^{(i)}} & \frac{\partial \hat{r}_s^{(i)}}{\partial \hat{\underline{p}}_{m,s}} \end{bmatrix} \begin{bmatrix} \delta b_m^{(i)} \\ \delta \hat{\underline{p}}^{(i)} \\ \delta \hat{\underline{p}}_{m,s} \end{bmatrix} \quad (4.16)$$

and the quadratic component is

$$\begin{aligned} \nabla^2 F \left(b_m^{(i)}, \hat{\underline{p}}^{(i)}, \hat{\underline{p}}_{m,s} \right) &= \frac{1}{2} \begin{bmatrix} [\delta \hat{\underline{p}}^{(i)}]^T & [\delta \hat{\underline{p}}_{m,s}]^T \end{bmatrix} \begin{bmatrix} \frac{\partial^2 \hat{r}_s^{(i)}}{[\partial \hat{\underline{p}}^{(i)}]^2} & \frac{\partial^2 \hat{r}_s^{(i)}}{\partial \hat{\underline{p}}^{(i)} \partial \hat{\underline{p}}_{m,s}} \\ Sym & \frac{\partial^2 \hat{r}_s^{(i)}}{[\partial \hat{\underline{p}}_{m,s}]^2} \end{bmatrix} \begin{bmatrix} \delta \hat{\underline{p}}^{(i)} \\ \delta \hat{\underline{p}}_{m,s} \end{bmatrix} \\ &= \frac{1}{2} \begin{bmatrix} [\delta \hat{\underline{p}}^{(i)}]^T & [\delta \hat{\underline{p}}_{m,s}]^T \end{bmatrix} \begin{bmatrix} \Psi_s^{(i)} & -\Psi_s^{(i)} \\ Sym & \Psi_s^{(i)} \end{bmatrix} \begin{bmatrix} \delta \hat{\underline{p}}^{(i)} \\ \delta \hat{\underline{p}}_{m,s} \end{bmatrix} \end{aligned} \quad (4.17)$$

$$\Psi_s^{(i)} \equiv \frac{\partial^2 \hat{r}_s^{(i)}}{[\partial \hat{\underline{p}}^{(i)}]^2} = \frac{\partial^2 \hat{r}_s^{(i)}}{[\partial \hat{\underline{p}}_{m,s}]^2} = -\frac{\partial^2 \hat{\underline{r}}_s^{(i)}}{\partial \hat{\underline{p}}^{(i)} \partial \hat{\underline{p}}_{m,s}} = \frac{1}{[\hat{r}_s^{(i)}]^3} \begin{bmatrix} (\Delta y_s^{(i)})^2 & -\Delta x_s^{(i)} \Delta y_s^{(i)} \\ Sym & (\Delta x_s^{(i)})^2 \end{bmatrix} \quad (4.18)$$

where *Sym* denotes repetition of an element in these symmetric matrices. Note that this expression for $\Psi_s^{(i)}$ must be modified slightly for a 3-dimensional array.

For a single sample point and given iteration, the aggregate set of range perturbation equations can be written as

$$\delta \underline{r}_{s,k} = G_{s,k} \cdot \delta \underline{z}_{s,k} + Z(\delta \underline{z}_{s,k}) \cdot H_{s,k} \cdot \delta \underline{z}_{s,k} \quad (4.19)$$

$$Z(\delta \widehat{\underline{z}}_{s,k}) \equiv \begin{bmatrix} [\delta \underline{p}^{(1)}]^T & [\delta \underline{p}_{m,s}]^T & \cdots & 0_{1,\mathcal{D}} & 0_{1,\mathcal{D}} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0_{1,\mathcal{D}} & 0_{1,\mathcal{D}} & \cdots & [\delta \underline{p}^{(N)}]^T & [\delta \underline{p}_{m,s}]^T \end{bmatrix} \quad (4.20)$$

$$H_{s,k} \equiv \begin{bmatrix} & \Psi_s^{(1)} & \cdots & 0_{\mathcal{D}} & -\Psi_s^{(1)} \\ & -\Psi_s^{(1)} & \cdots & 0_{\mathcal{D}} & \Psi_s^{(1)} \\ 0_{\mathcal{D} \cdot (N+1), N} & \vdots & \ddots & \vdots & \vdots \\ & 0_{\mathcal{D}} & \cdots & \Psi_s^{(N)} & -\Psi_s^{(N)} \\ & 0_{\mathcal{D}} & \cdots & -\Psi_s^{(N)} & \Psi_s^{(N)} \end{bmatrix} \quad (4.21)$$

where \mathcal{D} is the dimension of the array (either 2-D or 3-D). $H_{s,k}$ is a roughly block-diagonal matrix containing the local system Hessian matrices as the diagonal and right-band terms, and an additional set of zeros on the left band to eliminate the biases in $\delta \underline{z}_{s,k}$, which do not appear in the 2nd-order expansion. $Z(\delta \underline{z}_{s,k})$ is a block-diagonal matrix that contains the transceiver positions: the desired observable of the solution process. Note the multiple instances of the location of the mobile transceiver.

In a similar fashion to the linear case, the gradient equations for each sample point can be collected into a global gradient equation. Once the coordinate constraints have been removed the resulting relation is written as

$$\delta \underline{r}_k = \mathcal{G}_k \cdot \delta \underline{z}_k + \mathcal{Z}(\delta \underline{z}_k) \cdot \mathcal{H}_k \cdot \delta \underline{z}_k \quad (4.22)$$

This equation contains the locations of the transceivers in quadratic form through the presence of $\mathcal{Z}(\delta \underline{z}_k)$ and so presents the same solution difficulties as the true, nonlinear system description. In fact, avoiding the quadratic nonlinearities was the purpose for developing the previous linear solution technique. Fortunately, because Equation 4.22 involves perturbations around a nominal estimate, it is possible to solve it in an approximate manner using a new 2-step cascaded solution method called Quadratic Iterative Least Squares (QILS).¹ Although the quadratic term is not represented exactly using QILS, it does exert a strong

¹Details of the QILS algorithm have also been presented in ([37][38]).

influence on the local gradient and greatly improves the performance of the overall iteration process in situations with large initial estimate errors.

The QILS solution process is as follows: At each iteration stage, the first step is to solve the linearized system equations as before. This gives

$$\delta \underline{z}'_k = \left(\mathcal{G}_k^T \mathcal{G}_k \right)^{-1} \mathcal{G}_k^T \cdot \delta \underline{r}_k \quad (4.23)$$

The resulting state perturbation estimate $\delta \underline{z}'_k$ is used to construct the matrix $\mathcal{Z}(\delta \underline{z}'_k)$, which is then substituted into Equation 4.22. This results in a new linear equation of the form

$$\delta \underline{r}_k = \mathcal{J}_k \cdot \delta \underline{z}_k \quad (4.24)$$

with

$$\mathcal{J}_k \equiv \mathcal{G}_k + \mathcal{Z}(\delta \underline{z}'_k) \cdot \mathcal{H}_k \quad (4.25)$$

Equation 4.24 can then be solved for $\delta \underline{z}_k$ using the left pseudo-inverse of \mathcal{J}_k

$$\delta \underline{z}_k = \left(\mathcal{J}_k^T \mathcal{J}_k \right)^{-1} \mathcal{J}_k^T \cdot \delta \underline{r}_k \quad (4.26)$$

This process is repeated during every iteration step, as is illustrated in Figure 4.5.

The importance of the quadratic term can be seen by examining the relative contribution of the linear and the quadratic terms in Equation 4.24. This is best done through a comparison of the matrix norms. Figure 4.6 presents the ratio of the norms of \mathcal{G}_k and $\mathcal{Z}(\delta \underline{z}'_k) \cdot \mathcal{H}_k$. It is apparent that for very small errors in the initial state estimate, the linearized solution is close to the actual array geometry and the inclusion of a quadratic correction term is unnecessary. Once the estimate errors reach approximately 30% of the array size the quadratic term is of equal importance, and for larger estimate errors the quadratic term is clearly dominant.

Figure 4.7 illustrates the same conclusion using the same Monte-Carlo simulation as was employed for Figure 4.4. For small biases/initial error estimates there is only marginal improvement with the QILS algorithm over the standard linear ILS algorithm. At the crucial intermediate bias level, however, the quadratic solution exhibits a marked improvement over its linear counterpart, raising the successful self-calibration rate from approximately 20%

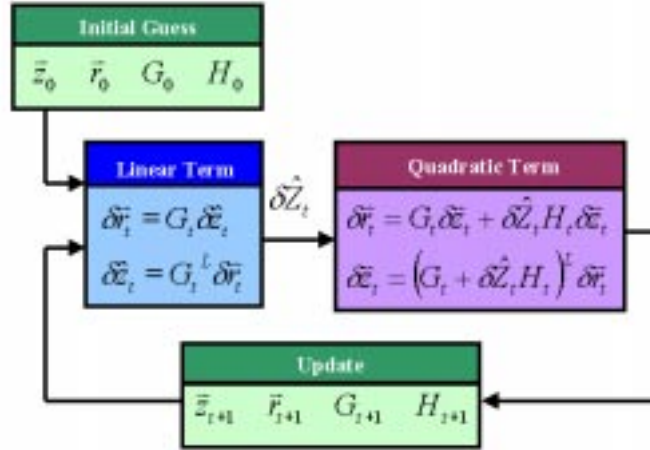


Figure 4.5: Quadratic Iterative Least Squares Process

The basic flow of the new quadratic Iterative Least Squares (QILS) algorithm. The initial estimate of the system state is represented by \underline{z}_0 , and the corresponding ranges between the transceivers by \underline{r}_0 . The system is first linearized about this estimate, yielding the linearized geometry matrix G . The matrix is then inverted using the left pseudo-inverse and multiplied by the difference between the estimated and measured range to give a linear state correction term $\delta \underline{z}$. This linear term is used to populate the block diagonal $\delta \hat{Z}$ matrix from the second-order expansion of the system, resulting in a new 'linearized' geometry matrix $(G + \delta \hat{Z} \cdot H)$ that incorporates some aspects of the actual nonlinear geometry.

Like the linear ILS algorithm shown in Figure 4.2, the QILS algorithm minimizes the RMS measurement error residual $\delta \underline{r}^T \delta \underline{r}$. It presents greatly improved convergence properties, however, especially when the initial state-estimate is poor.

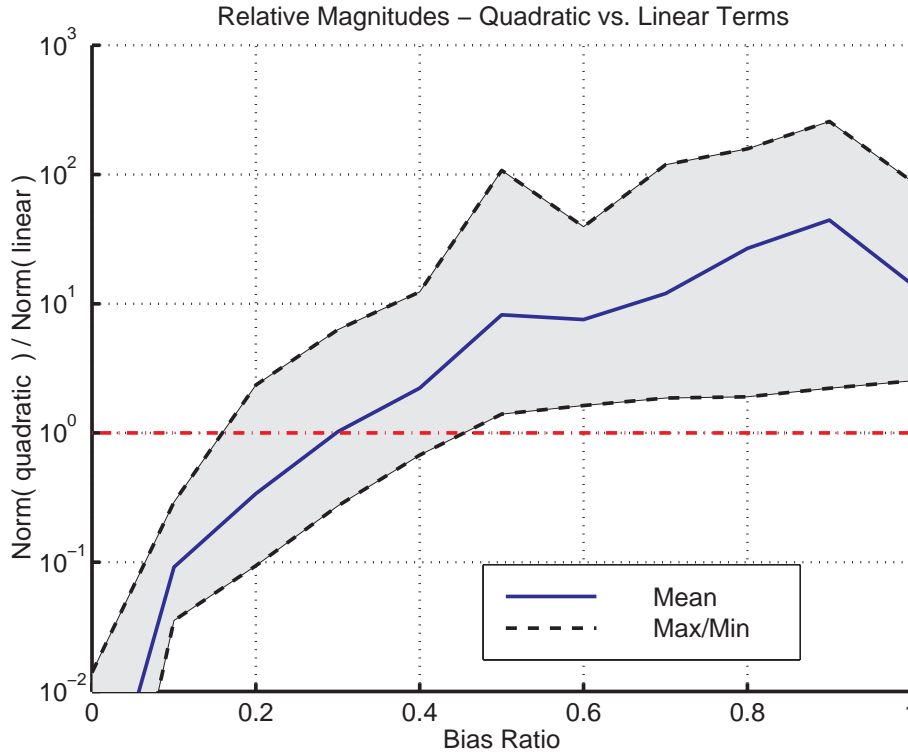


Figure 4.6: Linear Versus Quadratic Contributions

A comparison of the matrix norms of the linear and quadratic terms in the QILS algorithm. The values presented are mean results from 37,000 Monte-Carlo simulations with the array in the nominal triangular configuration. The mobile transceiver makes circular loops around the center of the array, with radii from zero up to 1.5 times the array size.

With bias ratios greater than approximately 30% of array size, the quadratic term dominates the solution process. At very high bias ratios the quadratic term likely has an even greater maximum effect than is shown, because simulations wherein the linear step completely diverged and $\mathcal{Z}(\delta \underline{z}'_k) \cdot \mathcal{H}_k$ could not be computed do not appear in this plot.

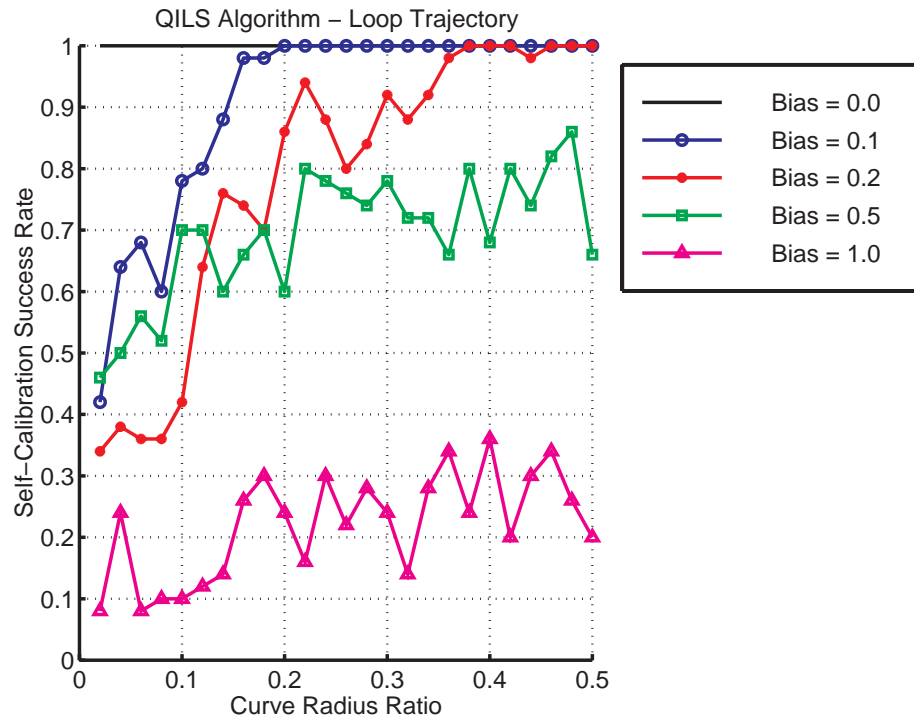


Figure 4.7: Quadratic ILS (QILS) Algorithm Success Rate

The success rate of the Quadratic ILS algorithm as a function of the curve radius and the range-measurement bias magnitude. The array configuration and trajectory are as shown in Figure 4.3, with transceiver #3 located at point (0.5,1.0). Compare with Figure 4.4.

to greater than 70%. With initial biases as large as the array itself, the QILS algorithm is able to successfully self-calibrate the array about 25% of the time.

4.6 Multiple-Estimate Solution

Although the QILS algorithm is very successful at raising the self-calibration success rate for intermediate bias values, neither it nor the standard linear ILS algorithm alone guarantees a successful self-calibration even with relatively small biases. For missions to Mars or other costly applications failure rates as high as 5-10% are simply unacceptable, and in fact success rates would likely have to approach 100% for any such technology to be considered. This section discusses the algorithm modifications necessary to achieve this final level of

self-calibration reliability.

4.6.1 ILS/QILS Failure Modes

The base performance of the linear and especially the quadratic ILS formulations may be substantially improved through an understanding of the failure modes for these algorithms. These failures are caused by three primary factors: local singularities, local minima, and parameter sensitivity.

Singularities

Local singularities occur under two primary conditions. The first is when two of the transceivers move very close to each other, thereby making it difficult to distinguish between them by ranging measurements alone. This may occur when the mobile transceiver passes right next to a stationary transceiver during the calibration trajectory, essentially collocating the devices for a brief instant in time. If it is recognized that this has occurred it is possible to remove that point from the batch process. Singularities also occur when the stationary transceivers are not arranged as a triangle, but rather in a straight line. This situation clearly results in a degenerate geometry for the entire array, and should be avoided by proper placement of the transceivers.

Local Minima

As with many other iterative solution techniques involving approximation of a nonlinear equation, both the ILS and QILS algorithms are subject to falling into local minima. In fact, the majority of cases in which the algorithms fail to provide a successful self-calibration result from convergence to local minima. This situation is mitigated somewhat by the fact that most of these local minima are far removed from the actual array geometry, and examination of the solution results can detect most of these incorrect solutions with little difficulty. More robust solution methods employing techniques such as simulated annealing to avoid these minima might also offer some improvement without requiring more extensive logic.

Parameter sensitivity

Closely coupled with the matter of local minima is that of parameter sensitivity. The successful convergence of the algorithms to the proper solution is sometimes extremely sensitive to the values of certain of the array parameters, and small shifts in those parameters may result in convergence to false local minima. For example, it was discovered that in certain circumstances placing the stationary transceivers in a perfect equilateral triangle yields an extremely favorable geometry resulting in nearly 100% success even with large range biases. A small shift in the location of one of these transceivers, however, would drop the success rate down to 50% because of the nearby presence of a local minimum. This sensitivity is especially strong in the previous simulations, which contain many round numbers because of the uniform bias magnitudes and therefore are more likely to hit solution singularities. Although randomizing all of the simulation parameters would make isolating the effect of specific parameters somewhat more difficult, it would also mitigate such situations and present results more closely aligned with an actual physical deployment of an SCPA.

4.6.2 Multiple Seeding

The technique used in this research for improving the basic algorithm effectiveness involves taking advantage of the parameter sensitivity to avoid the local minima through the use of multiple-seeding. The self-calibration algorithm is run several times on the same data but with a random initial state estimate, generally within 10-20% of the position estimate derived from the code-phase range measurements. Assuming that enough attempts are made, some of these will converge to the correct array geometry. Others will converge to one or more local minima. A comparison of the measurement error residuals $\delta \underline{r}_k^T \delta \underline{r}_k$ from the many resulting solutions then reveals the correct solution. In the presence of additional unmodeled error sources, more sophisticated discriminator heuristics could be used.

Figure 4.8 shows how the QILS algorithm improves greatly under the application of this multiple-estimate solution technique. Even for initial range biases as large as the array itself, successful self-calibration is nearly 100%.

4.6.3 Algorithm Comparison

Extensive Monte-Carlo simulations have shown that, together with the magnitude of the actual range biases, the success of the self-calibration process is most sensitive to the shape

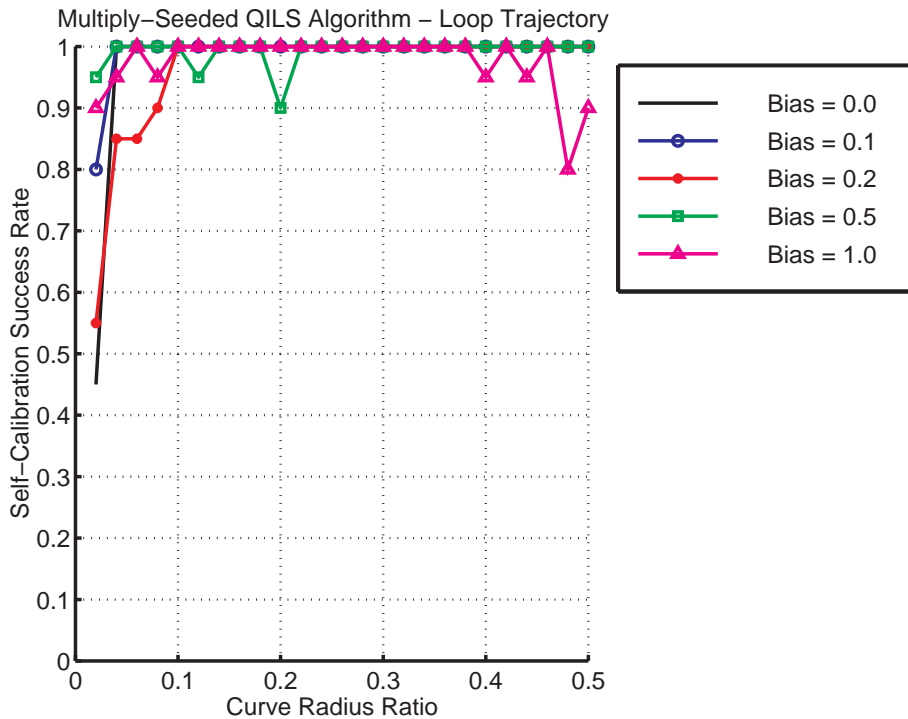


Figure 4.8: Multiply-Seeded Quadratic ILS Algorithm Success Rate

The success rate of the multiply-seeded Quadratic ILS algorithm as a function of the curve radius and the range-measurement bias magnitude. The array configuration and trajectory are as shown in Figure 4.3, with transceiver #3 located at point $(0.5, 1.0)$.

For each test, 20 different runs of the Quadratic ILS algorithm are used. Each run starts from the nominal code-based solution with an additional random Gaussian variation ($\sigma = 0.2$ units) of the locations of the stationary transceivers. The run with the lowest norm of the residuals (2-norm) is chosen as the best result. Compare with Figure 4.7.

of the triangular array space formed by the three stationary transceivers. Because the SCPA is designed for autonomous deployment, the resulting array shape is also a factor that is highly uncertain. The following data show how the multiply-seeded QILS algorithm provides for successful array self-calibration throughout the full range of likely array configurations, even with extremely large bias values.

Figures 4.9 and 4.10 present the results of 30,000 Monte-Carlo simulations performed with respect to the shape of the array. Although the nominal loop trajectory is employed, the radius of the curve is varied at random between very tight trajectories and large, looping ones. The initial range-measurement biases are kept below 20% of the size of the array. This corresponds to a moderate multipath environment, which would impose an average of 10 meters of error within a 100 meter array. This is somewhat greater than the level of multipath that was typically encountered during the field testing of the experimental system — the results of which are described in Chapter 6 — and is therefore considered the normal operating condition of the array. The circle in the plots represents the maximum limit on the variation in location for the third stationary transceiver that is deemed likely to result from a very poor autonomous deployment, given that the nominal array configuration is close to an equilateral triangle. Successful self-calibration for all locations within this boundary is strongly desired.

As the plots reveal, at these lower bias values ($< 20\%$) there is very little difference in success rate between the linear and quadratic ILS solutions. Both yield over 90% success within the critical region with a single application, and with the multiply-seeded approach success is nearly universal: 99.98% for linear ILS and 100.00% for quadratic. The two failure points for the ILS algorithm are right at the boundary of the critical region near the known singularity along the x-axis, and do not represent a serious threat to the system.

When the biases become larger, the difference between the algorithms becomes apparent. Figure 4.11 shows the results of another 30,000 Monte-Carlo simulations, this time with biases up to 100% of the array size. This represents extreme multipath for a normally sized array, and is still very large even for arrays as small as 20 meters across. In this case the multiply-seeded linear algorithm cannot cope with the large error in its initial state estimate, and self-calibration success drops to only 65%. The multiply-seeded QILS algorithm, however, can easily handle even these extreme conditions. Although a few configurations — generally near the boundary — result in algorithm failure, the overall success rate is still 99.80%. This shows the value of the quadratic algorithm in providing algorithm success

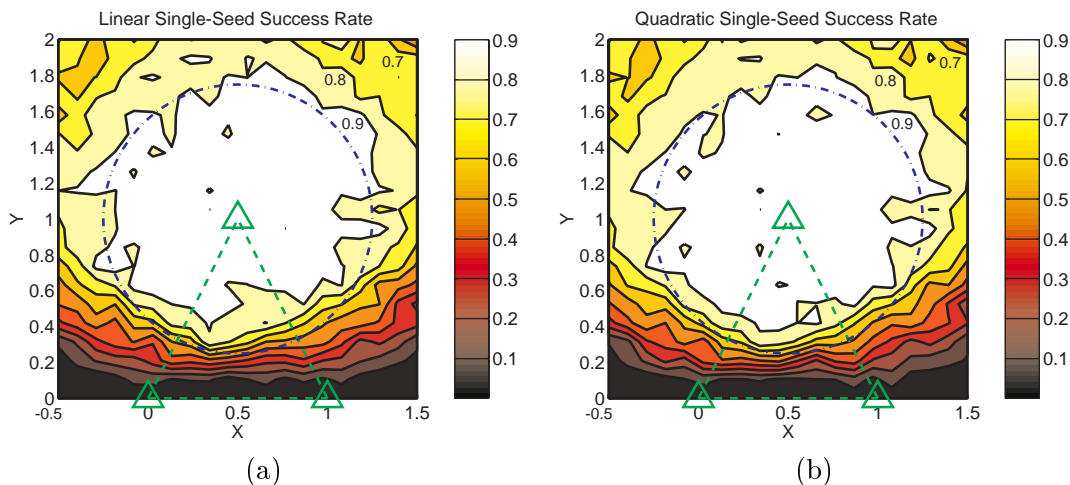


Figure 4.9: Algorithm Success Rates WRT Transceiver 3 Location — Small Biases, Singly-Seeded

The success rates of the singly-seeded self-calibration algorithms as a function of transceiver #3 location, derived from 30,000 Monte-Carlo simulations. Transceiver #1 is located at the origin, and transceiver #2 is at point $(1.0, 0.0)$. Transceiver #3 can be located anywhere within the plot area, with a uniform distribution centered around its nominal location at point $(0.5, 1.0)$. The dotted triangle shows this nominal array geometry.

The rover performs a looping trajectory outside the three stationary transceivers as shown in Figure 4.3. The radius of the trajectory curve is uniformly distributed with $0.05 < R < 1$. The carrier-phase integer bias values are also uniformly distributed, with magnitudes ranging from $0 < Bias < 0.2$.

Plot (a) represents success for the basic linear ILS algorithm, and plot (b) represents success for the QILS algorithm. The same raw data are used to generate each plot: only the applied algorithms are different. The contour lines are drawn at 10% intervals.

The circle in the plots is centered around the nominal transceiver #3 location and has a radius of 0.75 units. It represents a conservative bound on the expected range of locations for transceiver #3 from a hypothetical autonomous deployment, and therefore the region in which 100% algorithm success is desired.

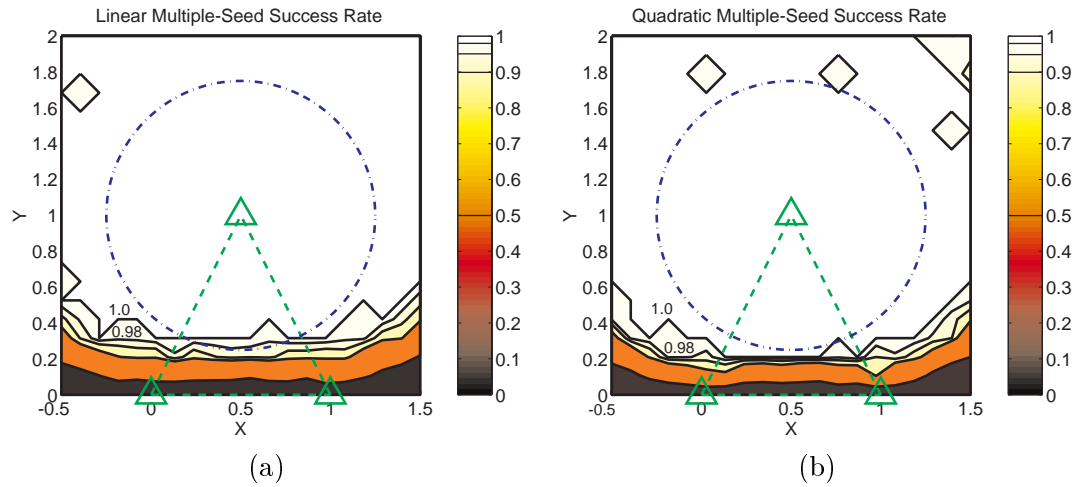


Figure 4.10: Algorithm Success Rates WRT Transceiver 3 Location — Small Biases, Multiply-Seeded

The success rates of the multiply-seeded self-calibration algorithms as a function of transceiver #3 location. The simulation parameters are the same as for Figure 4.9, and exactly the same set of raw measurement data is used. The only difference is that both the linear ILS and the QILS algorithm are each allowed up to 50 seedings with different initial state estimates.

Contour lines are drawn at 100%, 98%, 95%, 90%, 50%, and 0% success rates. The ‘diamond’ shapes are contour lines enclosing a region of 98% success caused by a single failure point within that region. The diamond shape itself is an artifact of the coarseness of the grid used for tallying the self-calibration success rates.

Algorithm	Maximum Bias Value		
	0.2	0.5	1.0
Linear ILS, Single Application	0.9085	0.5300	0.1359
Quadratic ILS, Single Application	0.9130	0.7208	0.5791
Linear ILS, Multiply-Seeded	0.9998	0.9582	0.6502
Quadratic ILS, Multiply-Seeded	1.0000	0.9978	0.9980
Number of Points	12,808	12,829	12,571

Table 4.2: Algorithm Success Comparison

Success rates of the different self-calibration algorithms as a function of the maximum bias errors in the initial state estimate, within the circular bounding region described in Figure 4.9. The location of transceiver #3 is evenly distributed throughout the bounding region, and the total number of Monte-Carlo simulations (common to each algorithm) that fall within that region is given. A successful self-calibration is one that correctly determines all of the system states of interest: the locations of the stationary transceivers, the trajectory of the mobile transceiver, and the carrier-phase integer biases.

under worst-case situations.

Table 4.2 lists the algorithm success rates within the critical regions for Figures 4.9, 4.10, and 4.11, in addition to a more intermediate case with biases up to 50% of the array size. It is clear that the success rate of the linear ILS falls off the fastest as the biases grow. A single application of the QILS algorithm is almost as effective as the multiply-seeded linear algorithm at large bias values; but at only 60% success neither of these methods is reliable enough for critical applications. By comparison, the multiply-seeded QILS algorithm maintains its high success rate for *all* bias values.

4.7 Additional Considerations

Even the multiply-seeded QILS algorithm fails consistently when presented with degenerate array configurations. An obvious one occurs when the stationary transceivers become almost collinear, resulting in a lack of state observability. An understanding of the effects of the different array parameters greatly aids the system designer in the creation of an SCPA configuration that will maximize the likelihood of success and minimize the need for additional human analysis and intervention. This section examines some of the more critical system parameters, and how they affect system performance.




Figure 4.11: Algorithm Success Rates WRT Transceiver 3 Location - Large Biases

The success rates of the various self-calibration algorithms as a function of transceiver #3 location. All details are the same as for Figures 4.9 and 4.10, except that the biases range from $0 < \text{Bias} < 1$. In addition, the contours for plots (a)-(c) are all now drawn at 10% intervals. (Plot (d) retains the uneven contour spacing.)

4.7.1 Geometric Factors

Array Geometry

As the previous simulations show, one of the most important consideration is the array geometry, consisting of the locations of the stationary transceivers. Although both field tests and simulations show that an equilateral triangle gives a good geometry for successful convergence, it is highly unlikely that an autonomously-deployed system would in fact be equilateral. This could be because of measurement and estimation errors during the deployment and before the self-calibration, or it could be because terrain or other factors preclude a more regular geometry.

One measure of the suitability of the array geometry is the DOP, which was discussed earlier. Although it does not accurately predict algorithm convergence it does give an indication of potential geometries to avoid. Figure 4.12 shows how the location of transceiver #3 affects the DOP of the array, both for an individual unknown parameter and for the array as a whole. Although the indicated DOPs are reasonably small (< 10) over many of the possible locations for transceiver #3, it rises dramatically near the singularities that exist as transceiver #3 approaches the location of either of the other two stationary transceivers. As has been previously noted, a similar singularity exists along the $y_3 = 0$ axis as the array degenerates to a straight line.

Although DOP is a useful metric for basic configuration studies because it shows the best possible post-calibration accuracy for differing geometries, it does not provide any direct knowledge as to whether the self-calibration process itself will be successful from an arbitrary initial configuration estimate. While there is often a strong correlation between a low overall system DOP and successful self-calibration, final validation of all proposed geometries should be done through Monte-Carlo simulations to account for the nonlinear convergence process.

Trajectory Shape

Changes in the trajectory shape are more difficult to evaluate because the large variety of possible trajectories cannot be easily characterized by a simple set of parameters. Rather than examining all possible trajectories, this section merely summarizes the results of extensive simulation and experimental testing with the hope that this knowledge will help guide the design of trajectories for future applications, especially those for differing array

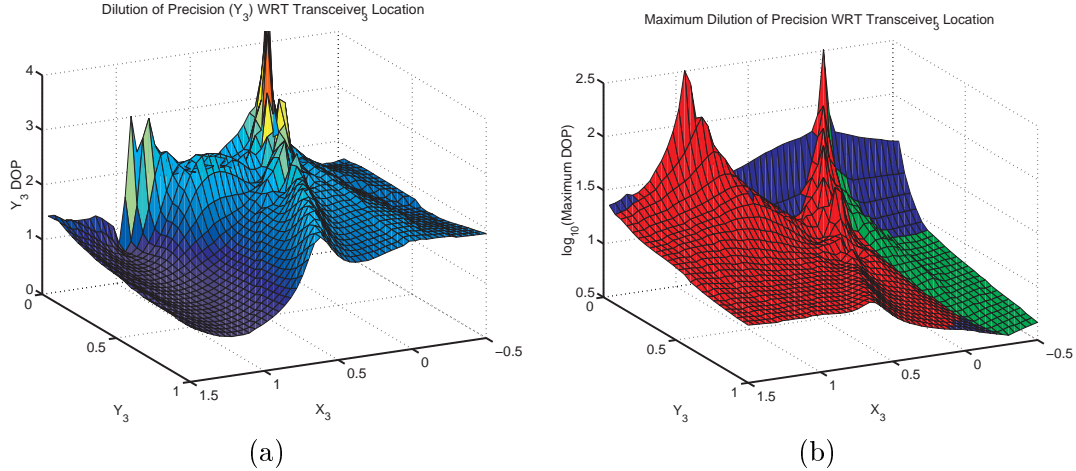


Figure 4.12: DOP Variation With Respect To Transceiver 3 Location

The linear Dilution of Precision (DOP) of the array as the position of transceiver #3 is varied. Transceiver #1 is located at point $(0,0)$ and transceiver #2 at point $(1,0)$. (Note the unusual orientation of these plots for clarity.) The rover performs a looping trajectory outside the three stationary transceivers as shown in Figure 4.3, with a turn radius around each transceiver of 0.25 units. Sample points are taken approximately every 0.2 units traveled.

Plot (a) shows the DOP of the estimate of the y -location of transceiver #3, and Plot (b) shows the maximum DOP for all states: the locations of the stationary transceivers, the range biases, and the locations of the sample points along the trajectory. As per Equation 4.14, the greater the DOP, the less observable the state. The different shades in Plot (b) indicate where different states are least observable. Note that Plot (b) uses a log scale.

Noticeable singularities ($DOP \rightarrow \infty$) exist when transceiver #3 is located close to either of the other two stationary transceivers. Another very strong singularity exists when $y_3 \approx 0$ and the array geometry degenerates to a straight line. For clarity, points close to this singularity are omitted from the plot.

The slightly rough nature of the plot comes from quantization issues associated with the sample spacing as the location of transceiver #3 changes. Finer quantization results in a smoother DOP variation — as well as a slight overall reduction in DOP — at the expense of increased computational burdens. The asymmetry over the $x_3 = 0.5$ axis is the result of asymmetries in the sample points along the trajectory as well as the choice of coordinate frame.

geometries.

For any trajectory, the key to success relies on relative transceiver motion (a) of large magnitude and (b) that gives independent changes in the range measurements to the different stationary transceivers. This also tends to correspond to situations where the mobile transceiver — during different sections of the trajectory — will sweep out large angles with respect to one or more of the stationary transceivers while maintaining more constant bearings to the others.

The loop trajectory around the outside of the array has been shown to be very effective for self-calibration. The radius of this loop is not a critical factor, providing that the rover spends only a small fraction of the trajectory in the immediate vicinity the singularities associated with the locations of the stationary transceivers. These singularities are considerably smaller than those associated with collocation of two of the stationary transceivers, but they can still be problematic, especially if the rover cuts close to all of the stationary transceivers rather than just one. This can be seen in Figure 4.7 by the drop in the rate of successful self-calibration when the curve radius drops below 10% of the overall array size. Larger loops generally work as well, with the effectiveness very gradually dropping off as the curve radius gets very large. This reduction in effectiveness occurs because the relative range changes become small and the relative angular displacements with respect to the stationary transceivers become very similar when the mobile transceiver is far away, and hence the overall observability of the actual array shape decreases.

Other trajectory shapes have also met with success. Circular trajectories work reasonably well, again assuming that reasonable separation from the stationary transceivers is maintained. Both circular and loop trajectories also work when the mobile transceiver moves inside the array instead of outside. The tradeoff that must be made in this case is to keep the trajectory as long as possible while still avoiding the singularities associated with the stationary transceivers. More-complicated trajectories such as figure-eight/cloverleaf patterns also give good results, although improvement over the simpler trajectories appears to be only marginal.

Number of Transceivers

The previous discussion has focused on the minimal triangular array consisting of three stationary transceivers. Similar analyses may be made for N-transceiver arrays with various configurations. Although this work has not examined expanded arrays for their potential

benefits, it is clear that adding more stationary transceivers to the array cannot hurt the array effectiveness, since at worst each sub-triangle in a larger array could be solved for independently and then combined after the fact. It is expected that adding additional devices to the array will eliminate or reduce some of the singularities through the presence of redundant range measurements, thus boosting overall array effectiveness. Further studies must be performed to determine if this is in fact the case.

4.7.2 Non-Geometric Factors

The success of both the ILS and QILS algorithms presented may be affected by many additional factors that are not related to the geometrical arrangement of the transceivers themselves. Some of these factors, such as vertical approximation error, multipath error, and the problems of cycle slips and occlusions have been discussed previously in Chapter 3, and can affect the self-calibration process in the same manner in which they affect the raw positioning process. Others affect only the self-calibration process. Some of the more critical of these are discussed below.

Sampling frequency/quantization

Table 4.1 presents the minimum number of trajectory sample points required for successful self-calibration. In order to maximize success, however, more sample points are recommended. The exact number and spacing of these points is a matter of engineering judgment. Continuous sampling gives the maximum chance of success, but is naturally computationally prohibitive. Simulations have shown that taking samples every 0.2 units traveled is adequate for most situations. Closer spacing provides only small advantages, while greater spacing does increase the likelihood of missing a sample in a geometrically significant region.

Coordinate frame

Several different potential choices for coordinate frame exist. A floating frame attached to the center of the array is in many ways optimal, since it does not bias the measurement errors in any particular direction. Such a frame does present both increased complexity and computational burden, however, due to the incorporation of additional constraint equations into the solution. The constraint method used for this research is to fix a frame arbitrarily with respect to two of the stationary transceivers. Although this skews the errors towards

the one side of the array, it has the benefit of being both simpler to analyze and code — and thus be less prone to error — and also of providing a stable reference frame that does not shift during the self-calibration process. This latter factor is especially useful when multiple rovers are navigating with respect to the array, and consistency in coordinate frame is essential.

Computational considerations

Significant improvements in algorithm speed can be achieved by noting that the linearized system matrices are very sparse. For example, the FLOPs (floating-point operations) count for the linear ILS algorithm has been demonstrably reduced by greater than an order of magnitude when this sparsity is taken into account and wasteful operations are avoided. This is the recommended course if these algorithms are encoded using C/C++ or another mid-to low-level language. Implementation in Matlab does not benefit from this consideration. First, Matlab appears to already have some built-in sparsity detection capability. Second, the logic associated with applying the sparsity knowledge is so expensive computationally that most the benefits of the raw FLOPs reduction are eliminated.

4.8 Summary

The techniques presented in this chapter allow for successful self-calibration of an array of pseudolite transceivers — i.e. the determination of the positions of the stationary pseudolites to centimeter-level accuracy — through utilization of the motion of a single mobile transceiver. The trajectory of that mobile transceiver and the resulting carrier-phase integers and line biases are simultaneously determined through a batch process.

The new QILS algorithm developed in this research for this application provides significant improvement in success rate over the more conventional linear ILS algorithm, especially under conditions with greater error in the initial estimate of the array configuration. Combining QILS with the stochastic multiple-seeding technique yields a self-calibration method that is effective, accurate, and robust: Monte-Carlo simulations of the self-calibration process demonstrate 100% successful self-calibration within the expected operating parameters of an SCPA, and 99.80% even in worst-case scenarios with initial estimate errors as large as the size of the array itself. Experimental verification of the effectiveness of the QILS algorithm for array self-calibration appears in Chapter 6.

Chapter 5

Experimental System

An operational experimental system has been created in order to demonstrate and validate both navigation using the Self-Calibrating Pseudolite Array (SCPA) invented in this research, as well the new self-calibration process itself. This prototype, which follows the overall system architecture presented in Section 1.3, consists of four GPS transceivers, a wireless data-collection system, and a central base-station computer running a custom software application for data processing. One of these transceivers has been specially modified for mounting onboard the NASA Ames K9 rover, enabling joint operations of the SCPA with other Mars-precursor technologies. This chapter describes some of the details of the experimental apparatus, focusing on the GPS transceivers themselves and the GPSPlexer software application.

In addition to the basic functionality requirements, the challenges of low-budget and low-manpower fieldwork have generated several auxiliary system requirements. These considerations dictate that the system be...

- Highly portable
- Simple to use
- Operable by a single person
- Easy to maintain and repair
- Rapidly reconfigurable to respond to changing operational needs

The current version of the prototype system meets each of these design goals.

5.1 GPS Transceivers

The primary physical component in the experimental system is the GPS transceiver. It is a custom design utilizing many commercial off-the-shelf (COTS) components in order to speed development time and reduce costs. Because very few GPS transceivers have ever been constructed, and because there are so many different possible architectures (Chapter 2), there is no generally accepted methodology towards their construction. Trial and error is an important part of the design process.

Many of the individual components perform critical functions in the overall transceiver design. Therefore, each of the transceiver components will be discussed in detail below. It is hoped that this document will aid future transceiver designers in their efforts.

5.1.1 Transceiver Architecture

The transceiver architecture used for this project is illustrated in Figure 5.1. The transceiver is self-differencing, so the receiver monitors the transmitter output at radio-frequency (RF) rather than by directly synchronizing to it via a common oscillator. Because this particular receiver has two front ends it can track the pseudolite signal either through the use of a direct line to a dedicated front end or through the normal airwave transmission. The former method works well at low signal power and may help reduce outside interference to the direct signal. At high power, however, leakage in the RF connectors tends to eliminate this benefit.

Both transmission methods have been used in the course of this research. Because hardware failures in the receivers left several with only one working front end, the current configuration utilizes only the airwave-propagated signal for the self-differencing. The high signal power and the pulsing method employed to overcome the near/far problem enable this strategy to work very effectively.

The GPS transceiver is designed to be both easily portable and to operate completely independently of any external infrastructure. Besides the receiver and pseudolite themselves the transceiver equipment includes the RF antenna system, amplifiers and attenuators to modify the received signal strength, a wireless communications system for data collection and remote command and control, and the unified power bus. An optional pulse generator allows for external pulse synchronization of the pseudolite output.

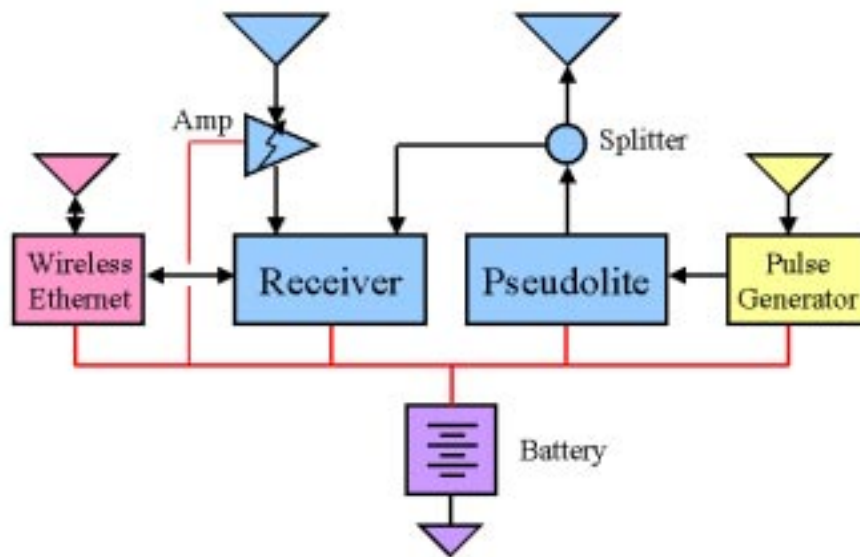


Figure 5.1: Transceiver Architecture

The basic transceiver architecture developed for this research. The receiver and pseudolite operate as independent devices, with the receiver monitoring the pseudolite output either through a direct input path or by airwave transmission through the transmit and receive antennas (the splitter is therefore optional). A wireless Ethernet collects the raw receiver data for processing at the base station computer. The pulse generator can be used to synchronize the pseudolite pulsing pattern to an external clock, although this is not necessary for all operating modes.




Figure 5.2: Transceiver Components

The two primary components in the self-differencing transceiver. Photo (a) shows an IntegriNautics IN200C pseudolite, with its SMA output connector on the front left face. Photo (b) shows the Stanford-modified Orion receiver. The two RF front ends are located on the left end of the board, with their SMA connectors just visible on the underside.

5.1.2 Pseudolites and Receivers

The pseudolite employed in the transceiver is an IntegriNautics IN200C signal generator, shown in Figure 5.2a. It features an adjustable signal power, variable output frequency, variable C/A code data rate, and a variety of pulsing schemes [13]. To overcome the near/far problem, the pseudolites were pulsed with a 31-bit (3% duty cycle) RTCM-derivative pulsing pattern.

The receiver is a 12-channel Mitel Orion receiver, which was modified at Stanford to include two separate RF front ends [46]. The Orion utilizes the GP2000 chipset, which has demonstrated some potential for space applications by successfully surviving over one year of on-orbit operations [54]. The greatest benefit of this receiver is that it is user programmable to enable it to track the pseudolite signals, which GPS receivers would normally reject due to their non-standard data messages. The receiver outputs raw code- and carrier-phase tracking data for all 12 channels at either 5 or 10 Hz over a standard 38.4 kbps RS-232 serial interface. Because of the limited communications bandwidth, the 5 Hz data rate was utilized for most of this project. The same interface can be used to send commands to the receiver in order to alter its tracking characteristics during operation.

5.1.3 Antennas

The antennas used by the transceivers for both the broadcast and reception of the pseudolite signals were of special concern because of the unusual geometry of the pseudolite array. There are two primary factors to consider when choosing the antenna design: gain pattern and polarization.

Most GPS antennas have a hemispherical gain pattern with roughly uniform gain from zenith to about 10 degrees above the horizon, after which the gain drops off dramatically. This pattern allows the user to adequately receive satellite signals originating from most of the sky, while reducing potentially interfering multipath from low elevation angles. Unfortunately this pattern is ill-suited for use with a ground-based pseudolite array because the desired signals all originate from low elevations, precisely the region conventional GPS antennas are designed to suppress. What is desired instead is an omnidirectional pattern with high gain at low elevations. Although this offers no multipath rejection, it is an unavoidable result of the array geometry.

In addition, conventional GPS antennas are right-hand circularly polarized (RHCP) in order to match the polarization of the GPS satellite signals, which are RHCP in order to afford users additional rejection of ground-based multipath. Such multipath tends to reverse polarization upon reflection, and when these reflected signals are received through a RHCP antenna this results in a roughly 6 dB reduction in strength of the multipath signals. This would be a highly desirable feature for SCPAs because multipath is not reducible through the choice of the gain pattern. Problems arise when the transceivers in the array make looping motions around each other, however, or when one rotates with respect the others. This causes phase-windup of the GPS signal, giving a cumulative error of one wavelength per rotation [2]. Because an SCPA relies upon looping motions of a mobile transceiver around the array for the self-calibration process, it is therefore desirable to have a non-polarizing antenna for this application.

The antenna design finally chosen for both transmit and receive antennas is the half-wave dipole shown in Figure 5.3a, which has the desired non-polarized omni-directional transmission pattern. Consisting of a female bulkhead SMA connector and two pieces of bus wire, it is very inexpensive and trivial to construct. The measured radiative efficiency as a function of frequency for one such antenna is presented in Figure 5.3b. Although the transmission pattern is not as sharp as with a commercial patch antenna, the wide pattern is actually desirable because the wire transmission elements tend to be prone to accidental

bending and are otherwise difficult to fine tune. The broad pattern therefore provides a level of robustness at the expense of a slight increase in the received noise level.

Other antenna designs are of course possible. A full-wave dipole, for example, would provide higher directional gain. The added length of the elements, however, and the associated susceptibility to bending make this an undesirable tradeoff. Also note that unlike commercial GPS antennas, the custom antennas do not include a low-noise amplifier (LNA). While this is desirable for the transmission antenna in each transceiver, it necessitates the addition of an external amplifier in the receive path.

5.1.4 Secondary Transceiver Components

Antenna Tripods

The antennas are mounted roughly one meter above the ground on top of tripods, as is shown in Figure 5.4, and attach to them using custom mounting plates. These 11 cm by 25 cm plates are constructed of 3 mm-thick Plexiglas, and have through-hole mountings for the antenna SMA connectors. Cutouts above and below each antenna, which are vertically positioned 12 cm apart, allow 360-degree horizontal transmission around the tripod.

Amplifiers

Because of FCC (Federal Communications Commission) broadcast-power limitations on the pseudolites themselves [8], the reliable operational range of the array is only about 10-15 meters without external amplification. Beyond this range tracking tends to be poor, with many signal dropouts and losses of lock. Adding an amplifier to the receive path increases the operational range to roughly 30-50 meters, as is demonstrated in Chapter 6.

The amplifier used for the receive path is a MiniCircuits ZLJ-3G broadband amplifier which provides about 19 dB gain. The gain was chosen so that even if the pseudolite were to broadcast directly through the amplifier into the receiver front end it would still be below the rated power limits. The ZLJ-3G is not an especially low-noise amplifier and its wide bandwidth (20-3000 MHz) greatly adds to the receiver noise figure. The fact that it is not collocated with the antenna but is further downstream at the end of a coax cable — and will thus amplify the cable noise as well — further adds to its inefficiency. Experimental tests showed no significant performance degradation in comparison with much more expensive LNAs, however, at least for ranges under 30 meters.

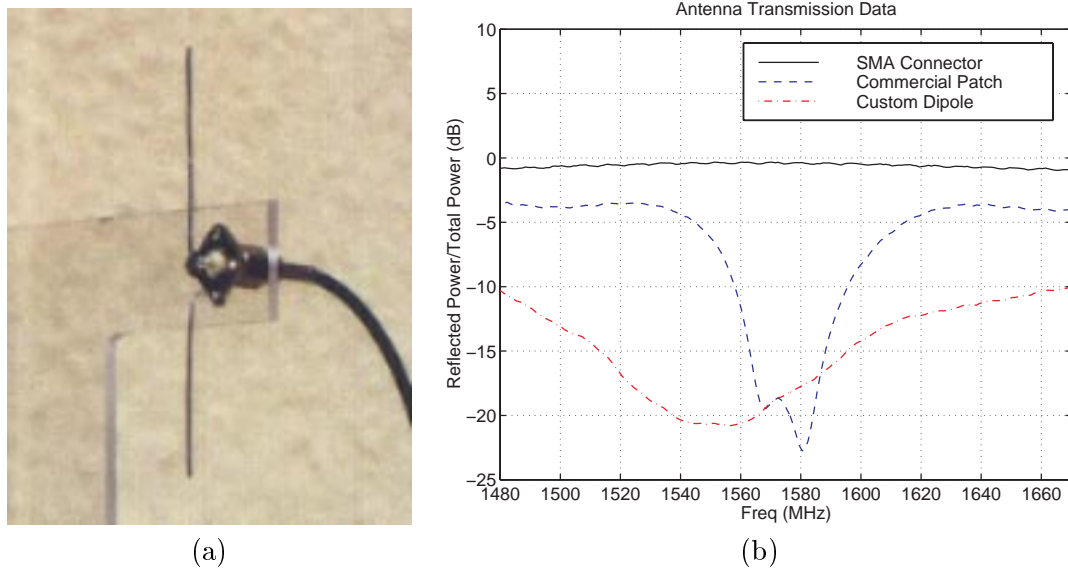


Figure 5.3: Custom Dipole Antenna

Photo (a) shows a custom dipole antenna constructed for the SCPA prototype, mounted upon a Plexiglas antenna plate. The total length of the antenna is 9.0 cm, slightly shorter than an ideal half-wave dipole (9.5 cm). Each antenna is individually tuned by connecting it to an HP 4396 Network/Spectrum Analyzer and then gradually reducing its length with a wire cutter until its radiant peak is near $L1$ (1575.42 MHz).

Plot (b) shows the spectrum of one such antenna, compared with a commercial patch antenna and an unconnected SMA connector. The plot shows the power reflected back into the analyzer (i.e. not transmitted), so a low value represents high transmission gain. The custom antenna gives a somewhat wider transmission spectrum than the commercial antenna, but is otherwise a very effective radiator ($> 99\%$ near $L1$).



Figure 5.4: Antenna Tripod

One of the antenna tripods in its deployed configuration. The antennas are located on a Plexiglas mounting plate at the top of the tripod, one above the other and in-line with the pre-surveyed 'X' reference point marked on the ground. The mounting plate is labeled with the transceiver number and the associated pseudolite PRN code number. The transceiver is underneath the tripod, and is connected to the antennas with RG 174 A/U flexible cable.

Attenuators

Although the external amplifier is necessary for reliable tracking, the high signal power occasionally makes it difficult for the receiver to initiate tracking on the pseudolite signal. It is therefore useful to have a variable attenuator after the amplifier in order to reduce the signal power and facilitate initial signal acquisition. The model used is an Alan Industries 50SV33-1773, a custom unit designed for previous Stanford GPS projects utilizing a carrier frequency at L1. The attenuation setting ranges from 0-33 dB in 3 dB increments.

The addition of attenuation to the receiver to aid in signal acquisition is rather counter-intuitive, because the problem normally encountered is low received signal power. Although the added amplifier yields its primary benefit once lock is obtained, however, experiment operation has shown that attenuating the signal can sometimes reduce the time required to obtain signal lock. With the current receiver hardware and software this tends to be necessary about 10% of the time, when the receiver cannot lock onto the signal despite knowing the correct frequency bin in which to search. This behavior is somewhat unexpected because the AGC still keeps the ADC in saturation even with the additional attenuation - as is evidenced by the constant post-correlation signal-to-interference (S/I) ratio throughout the range of attenuation values - and is not adequately described by the theory relating to receiver tracking of pulsed signals.

Wireless Communications

A wireless communications system is necessary in order to collect the receiver data at a common point for processing, in this case the RangeLAN2 series from Proxim. Each transceiver uses a Model 7910 Serial Adaptor, which takes in the raw receiver RS-232 output, wraps it in TCP/IP packets, and rebroadcasts it using a 2.4 GHz frequency-hopping spread-spectrum (FHSS) signal. This signal is received by a corresponding Model 7510 Access Point at the central base station, where it is accessed through the Ethernet port using standard socket interfaces. A collision-detection/rebroadcast algorithm ensures that packet collisions do not result in any data loss, although additional data latency may result. The range of this system is at least 150 meters, and no interference has been observed between the RangeLAN devices and either the GPS pseudolites or other wireless units.

Pulse Synchronization

In order to eliminate interference when using certain (generally wide-pulse) pseudolite pulsing schemes, it is necessary for the pseudolites to be synchronized to some external timing source. This allows the pseudolites to broadcast in turn in a TDMA manner, with each one broadcasting its pulse in the temporal nulls of the other pseudolite signals. For this purpose the IN200C's are synchronizable off of a 1 pulse-per-second (PPS) external RS-422 signal. For distributed operations this signal must of course be wireless.

For the prototype SCPA this signal - when used - is generated by a standard radio-control (R/C) transmitter such as those used for model aircraft. The R/C receiver located with each transceiver outputs a set of pulses at 50 Hz, which is then subsampled via digital counters to generate 1 PPS. This works very well for static pulsing schemes in which the pseudolite pulse is at a fixed offset from the trigger pulse, since the pseudolite pulse repeats every epoch (1 ms). It does not work well for sweeping patterns in which a 20 ms offset makes a difference, however, because each digital counter will generally start at a different location on the raw 50 Hz R/C receiver output. Efforts to bypass this ambiguity by turning on the R/C transmitter *after* all of the R/C receivers so that they will all see the same initial pulse met with only limited success, because in the absence of a transmitter signal the receivers will occasionally generate spurious noise that appears like a pulse to the counters. A more sophisticated decoding system on the R/C receivers could solve the problem and eliminate the ambiguity by using one of the extra channels on the R/C transmitter as a switch to reset the pulse count on all the receivers simultaneously. This proposed technique has not been implemented in this experimental system, however.

Note that pulse synchronization does not eliminate interference between pseudolites for pseudo-random pulsing schemes such as RTCM, although it may reduce the interference probability. Because the present pulsing scheme is a narrow-width RTCM pattern that generates only minimal interference, performance is satisfactory without using pulse synchronization for the transceivers.

Transceiver Totes

All of the transceiver components except for the antenna tripod are mounted in a convenient tote-bucket for portability. The permanent components are securely mounted to prevent accidental damage. A 12 V, 4.4 mAh NiCd battery pack provides all of the necessary power



Figure 5.5: Transceiver Tote

These tote-buckets carry all of the components for the stationary transceivers with the exception of the antenna tripods. The major components visible in this photo are (clockwise from the upper left) the wireless serial adaptor, the GPS receiver, the GPS pseudolite, and the battery pack.

and gives a continuous operational endurance of about four hours. The total transceiver mass (excluding the tripod) is just over 5 kg.

5.2 Base Station

The second major component of the SCPA is the base station, which is used to collect raw data from the transceivers for processing. These processing tasks include bidirectional ranging between transceiver pairs, triangulation for determination of the transceiver locations, full array self-calibration, and raw data storage for later replay and analysis. All of these tasks must be done in near real-time.

The computer used is a 133 MHz Dell Latitude laptop running the Windows NT operating system, which communicates with the transceivers through its Ethernet port using the RangeLAN2 Access Point. The Access Point runs off of a separate NiCd battery pack. The main program used for SCPA operation is described in detail in Section 5.3.

The base station computer, wireless ethernet, and the transceivers and tripods are all transportable on a single cart.

5.3 GPSMixer Software

Because an SCPA utilizes non-standard operations and algorithms it was necessary to develop a new software application to manage the data processing. The resulting application is named the ‘GPSMixer,’ and is inspired in part by the mixing consoles used in professional sound systems. The software ‘mixes’ together the raw data from the receivers and combines them in various permutations to form bidirectional ranging pairs. These pairs are then processed to generate the overall array geometry. The highest program-level manages the self-calibration process, while the user is also able to interface with the receivers directly at a low level. In addition the software is capable of storing and retrieving experimental data for later analysis, a feature that is especially useful when developing and debugging positioning algorithms. Interaction with the program is performed through a graphical user interface (GUI).

There were many different technical requirements for this software, the most important of which are listed below.

- Operate at near real-time
- Interface with the receivers over the wireless ethernet at 10 Hz
- Provide a low-level interface to the receivers, including command ability
- Be robust to data delays or dropouts
- Save all collected raw data
- Be able to replay any past data for analysis at both normal and fast speeds
- Allow user reconfiguration on the fly
- Allow user aiding for integer initialization or array configuration solutions
- Save configuration data for replay

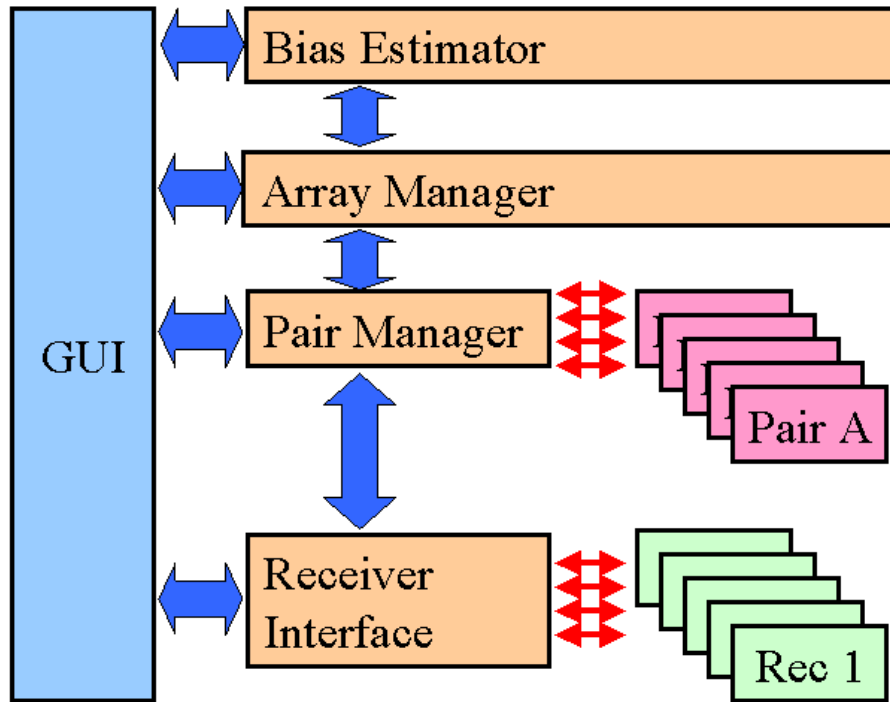


Figure 5.6: GPSTMixer Software Architecture

The overall software architecture for the GPSTMixer application. The array is built from the simplest components up, starting with the GPS receivers themselves, proceeding to bidirectional pairs of transceivers, and finishing with the array as a whole. Self-calibration is an additional layer above the basic positioning functionality. The user can interact with the program at any of these levels through the custom GUI.

The overall architecture addresses many of these requirements through its multi-tiered structure, which is presented in Figure 5.6. The user has full control and is able to interface with the application at any of the levels. A comprehensive view of most of the GUI windows, all open simultaneously, appears in Figure 5.7. In order to help illustrate the entire positioning and self-calibration process, each of the primary levels of the GPSTMixer is summarized below.

The GPSTMixer application is programmed using Visual C++ in the Microsoft Foundation Class (MFC) environment on a computer running the Windows NT 4.0 operating system. This choice has both advantages and disadvantages. On the positive side a large

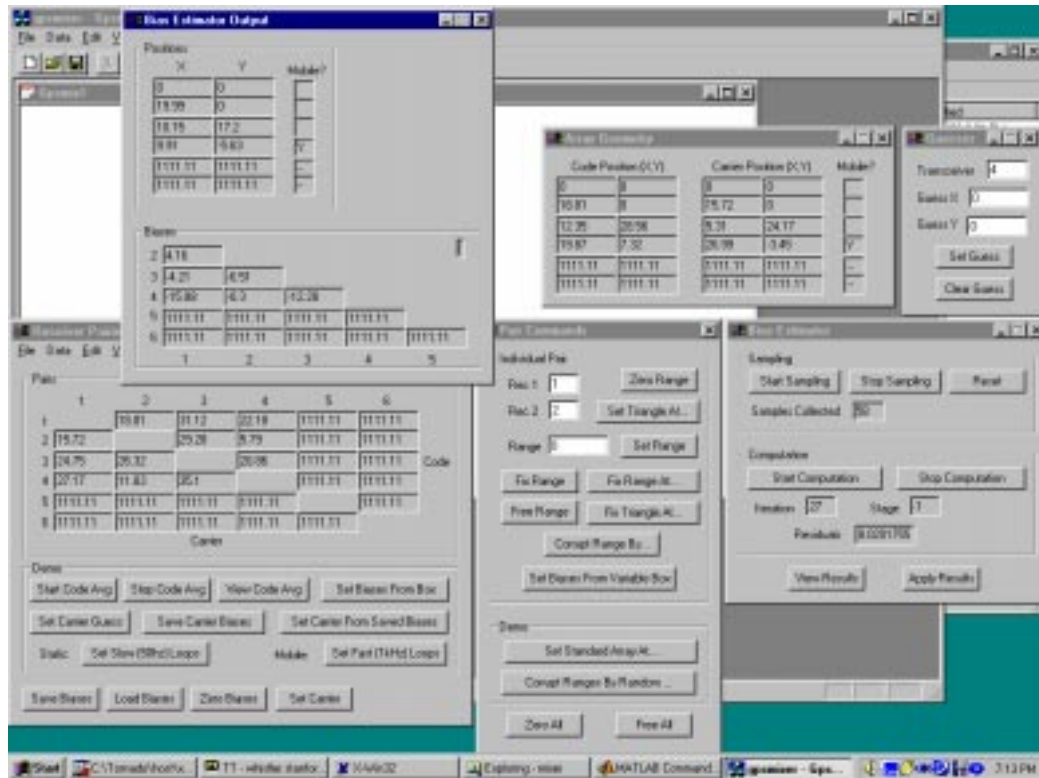


Figure 5.7: GPSMixer Software Interface

The primary input and output windows for the GPSMixer interface, displayed together on a single screen. Clockwise from the upper left are the final Bias Estimator output, the Array Manager output, a window where the operator can apply a priori knowledge to aid in ambiguity removal during positioning, the Bias Estimator controls, and the Pair Manager control and output windows.

number of libraries exist for common tasks such as socket interfacing; and because of the extensive auto-code framework, generating the GUI windows themselves is almost trivial. In addition, this choice leveraged off of the software heritage already existing within the ARL from other complex, field-deployable systems such as the HUMMINGBIRD autonomous helicopter ([67][56]). On the negative side the application must be structured to fit the MFC Document/View framework, which is well-suited to word processing and spreadsheets, but is less well-suited to robotics or other engineering applications. In addition, Windows NT is not a real-time operating system and so is not designed for time-critical data-collection processes.

5.3.1 Receiver Interface

The lowest level in the GPSPMixer architecture is the Receiver Interface, the GUI for which is shown in Figure 5.8. At this level the user can examine each channel on the receiver to view the current numerically-controlled oscillator (NCO) frequency offset, the received signal-to-interference (S/I) ratio, and information on tracking-loop synchronization. Commands can be sent to each receiver channel to change the tracked pseudo-random noise (PRN) numbers, switch between receiver RF front ends, and search for the signal at different NCO offsets, among other options.

5.3.2 Pair Manager

The next level up is the Pair Manager, which is pictured in Figure 5.9. At this level the ranges between all possible transceiver pairs are calculated from the the raw code- and carrier-phase data. These ranges (both code- and carrier-based) are displayed in a table as the off-diagonal terms of a Euclidean Distance Matrix (EDM). Biases can be added to these ranges, averaging filters can be utilized to reduce code-noise, and the array as a whole can be initialized to standard patterns of various sizes in order to examine raw ranging stability in the absence of biases. The Pair Manager also incorporates an error-correction filter to attempt to detect and fix cycle slips and other signal dropouts.

5.3.3 Array Manager

Above the Pair Manager is the Array Manager, which takes the range measurements and computes the corresponding array configuration and the x-y locations of the transceivers

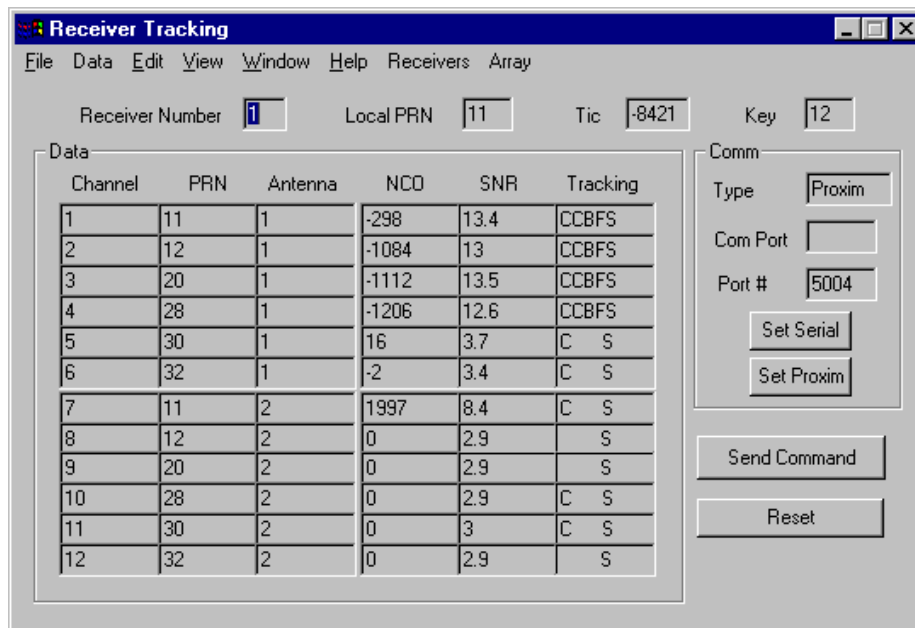


Figure 5.8: Receiver Interface

The interface for a single receiver. The main table presents tracking information for each of the 12 channels including the PRN number, NCO offset, signal-to-noise/interference ratio (SNR) and raw tracking flags for code-, carrier-, bit-, and frame-synchronization. The ‘Tic’ and ‘Key’ fields present time-synchronization data, while the box on the right controls the data communications system.

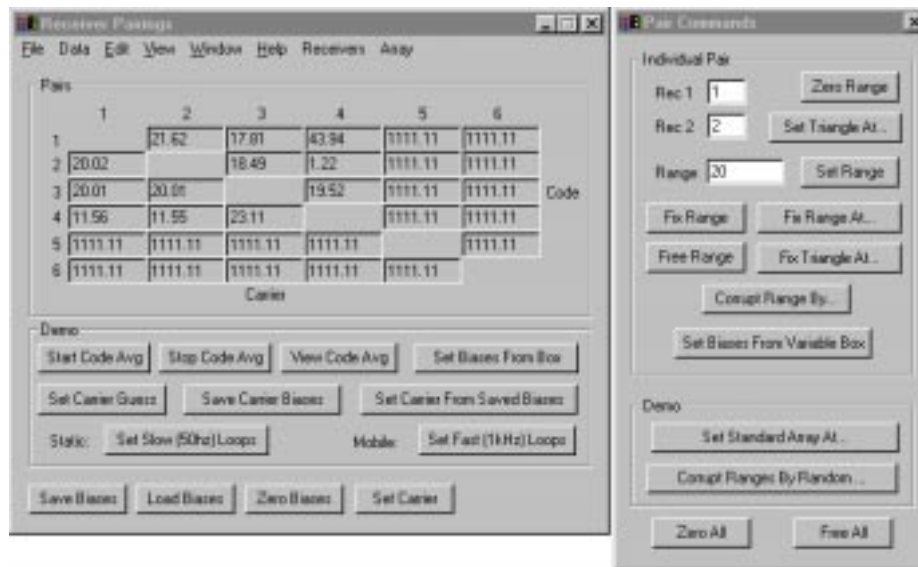
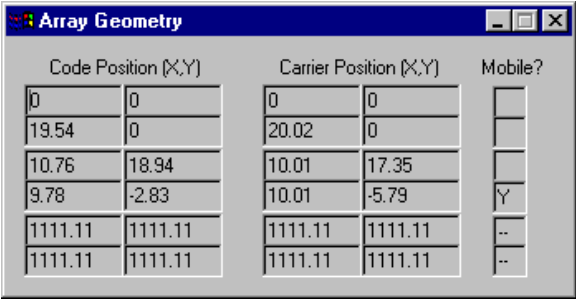


Figure 5.9: Pair Manager

The main table in the upper left corner presents the Euclidean Distance Matrix (EDM) of the ranges between the transceivers (in meters), code-range above and carrier-range below. Underneath the EDM are several array-specific commands related to the self-calibration procedure. The command box on the far right allows the user to individually set the bias on the range between any two transceivers.



Code Position (X,Y)		Carrier Position (X,Y)		Mobile?
0	0	0	0	<input type="checkbox"/>
19.54	0	20.02	0	<input type="checkbox"/>
10.76	18.94	10.01	17.35	<input type="checkbox"/>
9.78	-2.83	10.01	-5.79	<input checked="" type="checkbox"/>
1111.11	1111.11	1111.11	1111.11	--
1111.11	1111.11	1111.11	1111.11	--

Figure 5.10: Array Manager Output

This table presents the positions of the transceivers in the array (in meters), as determined by using either code- or carrier-ranging. Each row is for a different transceiver. Rows 1-3 represents the stationary transceivers (transceiver #1 is defined to be at the origin, transceiver #2 on the positive x-axis), and row 4 represents the mobile transceiver on the K9 rover.

using the methods presented in Chapter 3. Output from the Pair Manager, which is generated using both code- and carrier- measurements, is displayed to the user in the window shown in Figure 5.10. Note that for both programming and user convenience, the user commands to the Array Manager are handled through the Pair Manager GUI.

5.3.4 Bias Estimator

The final and highest level of the GPSMixer is the Bias Estimator, which incorporates the functionality of the array self-calibration algorithms. At any time the user may choose to begin the self-calibration process by starting the sampling process using the GUI shown in Figure 5.11. Position and carrier-range data samples are then taken by the program every 1-2 meters traveled by the mobile transceiver until the user tells the program to stop sampling and start the computation process. The Bias Estimator then applies the Quadratic Iterated Least Squares (QILS) algorithm described in Chapter 4 in a succession of stages, kicking out automatically once convergence is reached. The user may then view the resulting calculated biases and transceiver positions, as displayed in Figure 5.12. These bias corrections are applied to the raw range measurements only if they are approved by the user, at which point the self-calibration process is complete.

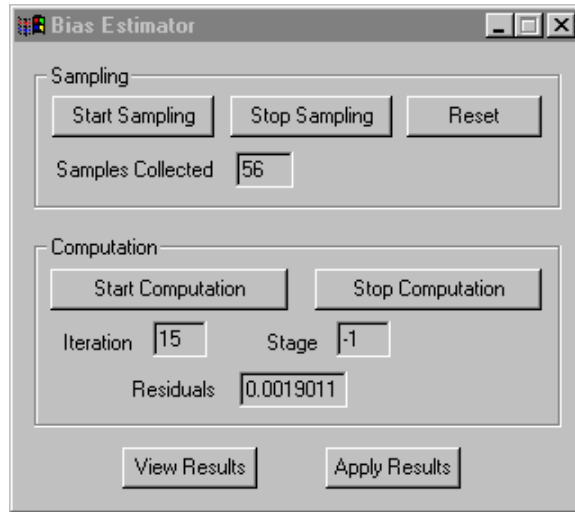


Figure 5.11: Bias Estimator Controls

Controls for the Bias Estimator provide the ability to start and stop both sampling and computation at will. Feedback for the operator consists of information on the iteration stage and the RMS range-error residual.

5.3.5 Real-Time Considerations

Several important features are incorporated into the underlying GPSMixer framework in order to make it emulate a real-time system.

First, all raw receiver data are timestamped, and these timestamps are compared whenever data from different receivers are combined in a bidirectional-ranging double-difference calculation. This is done to account for hardware latencies, delays from the wireless communications devices, and software delays due to process scheduling in Windows NT; and is especially critical because using non-simultaneous measurements at the double-difference level can result in large range errors (Section 3.4.4).

Second, the breakdown of the positioning solution into bidirectional ranging followed by array localization is less prone to time-delay errors than a more conventional solution wherein all of the raw measurements from all the receivers are combined and solved for simultaneously. Individual ranges may be computed whenever the data arrive, and if it is known that the two transceivers in the range measurement are both stationary then that measurement can be applied to array localization at any time, regardless of delay. Range

The screenshot shows a window titled "Bias Estimator Output" with two main sections: "Positions" and "Biases".

Positions Table:

X	Y	Mobile?
0	0	<input type="checkbox"/>
19.99	0	<input type="checkbox"/>
10.15	17.2	<input type="checkbox"/>
9.91	-5.63	<input checked="" type="checkbox"/>
1111.11	1111.11	--
1111.11	1111.11	--

Biases Table:

2	4.16				
3	-4.21	-6.51			
4	-15.88	-0.3	-12.28		
5	1111.11	1111.11	1111.11	1111.11	
6	1111.11	1111.11	1111.11	1111.11	1111.11
	1	2	3	4	5

Figure 5.12: Bias Estimator Output

The Bias Estimator, once it has completed computation, provides the operator with both the calculated positions of all of the transceivers (top table) and the corresponding carrier-phase biases in the bidirectional range measurements between the transceivers. The operator has the choice of whether to apply these results to the array.

measurements in which one or both of the transceivers are mobile are of course timestamped and only applied with other concurrent ranges.

Third, all computationally-intensive calculations — especially the bias-estimation and array self-calibration process — are broken up into many substeps. Each of these substeps returns control to main program after execution. Although this slows down the calculations considerably, it ensures that the algorithms do not appropriate all of the processor resources for large periods of time and that the continuous incoming raw data are still processed in a timely manner.

5.3.6 Future Improvements

The GPSMixer user interface reliably performs all of the required functions outlined in this section. Because it was developed as this research progressed, however, it has not been optimized for operator usability. Future versions of the software should include a better-organized and more comprehensive set of configuration setup commands and shortcuts, as opposed to the minimal ones necessary for the current hardware system. Of great benefit would be real-time pictorial output of the locations of the stationary transceivers and the rover trajectory — both before and after self-calibration — similar to the post-processed plots presented in Chapter 6. This would ease operator workload, reduce operator error, and facilitate system debugging.

Other potential software improvements include expanding the positioning and self-calibration algorithms to accommodate SCPAs with more than four transceivers, or those with out-of-plane devices. Chapter 7 discusses some of the implications of these two potential system extensions.

5.4 K9 Rover

Because the intended purpose of the Self-Calibrating Pseudolite Array is future Mars exploration, joint experimental tests were conducted at NASA Ames Research Center using its K9 Mars rover testbed (Figure 5.13). K9 is a variant on the JPL FIDO rover chassis that is used by NASA Ames to test new sensing and control techniques for future Mars exploration missions. These tests, the results of which are presented in Chapter 6, were used to verify that the SCPA is able to provide usable navigation data and that array self-calibration is possible using rover motion.



Figure 5.13: K9 Rover

The NASA Ames K9 rover that was used for validation testing of the SCPA. The mobile transceiver fits within K9's equipment bay, while the antennas are mounted on the Plexiglas plate on the right rear corner of the rover. The Proxim wireless ethernet is mounted halfway up the main mast.

The K9 rover is 1.1 meters long, 1.7 meters high including the sensor mast, and has a mass of approximately 65 kg. It has a six-wheel rocker-bogey suspension system suitable for rough terrain, and each wheel is independently steerable to allow for spot turns or lateral motion. K9's top speed is roughly 10 cm/s. The main sensor mast is topped with a pair of stereo cameras used for science tasks and area mapping, and a scanning laser rangefinder mounted on the front looks for nearby obstacles. The primary non-visual navigation sensors are reasonably primitive, consisting of odometry and a magnetic compass. More information about K9, especially its command and control architecture, can be found in [5].

A special GPS transceiver design was necessary for mounting integration with K9. Figure 5.14a shows this transceiver, which is mounted in an aluminum frame that fills one half of K9's internal bay next to its onboard computer.




Figure 5.14: K9 Transceiver

The interface frame used to mount one of the transceivers within the K9 electronics bay appears in photo (a). The receiver is the smaller box on the left. The right side of the frame supports the NiCd battery pack (above) and the pseudolite (below). Photo (b) shows the transceiver mounted within K9: The transceiver frame is in the upper half of this photograph.

5.5 Test Location

Field testing at NASA Ames was conducted in an open lot near the intake nozzle of the 80'x120' cross-section subsonic wind tunnel. The lot itself is roughly 30 meters wide by 60 meters long, and features slightly rolling terrain. The test area is shown in Figure 5.15. A small portable trailer at the back of the lot serves as the command center for K9 and also provides 120V AC power for the computers, cameras, and other equipment.

The presence of the wind tunnel inlet — which is about the size of a football field — only 30 meters from the edge of the test area presents a significant difficulty because of the multipath it generates: The horizontal omnidirectional transmission pattern of the transceivers leaves them very vulnerable to large vertical reflectors. The chain link fence surrounding the test area is a very good reflector near the pseudolite broadcast frequency as well, and the test data seem to indicate some code-range multipath errors greater than 10 meters (see Chapter 6). Carrier-phase operations seem to be largely unaffected once the integers are determined, however, because multipath errors are generally limited to one-half of the tracked wavelength, in this case less than 10 cm [53].

Figure 5.16 shows one of the field tests in progress, wherein K9 is traversing through the middle of the triangular array formed by the three stationary transceivers. Because of the



Figure 5.15: NASA Ames Test Site

The empty lot used for testing the SCPA with the K9 rover. The area is mostly flat, although landscaping efforts subsequent to this photograph have added several small (< 0.5 meter high) hills. The test field is surrounded by a 2 meter high chain link fence, and is located next to the 80'x120' subsonic wind tunnel and across the street from a large power distribution complex (not shown). These man-made features are the primary multipath sources affecting the system.



Figure 5.16: NASA Ames K9 Field Test

The SCPA in operation at NASA Ames. All three of the stationary transceivers are visible as K9 traverses through the middle of the array. Several pre-surveyed reference points are barely visible in the center of the photograph.

limited size of the test area, the largest usable triangular pseudolite array is 20 meters on a side. This allows K9 several meters of clearance around the outside of the array. The small array size somewhat limits the effectiveness of code-based ranging for array localization because the code errors are relatively large compared with the array dimensions. Despite this small size, however, it still takes K9 about 15 to 20 minutes to make a single 100 meter circuit around the outside of the array.

5.6 Summary

This chapter has presented the actual hardware and software used for the Self-Calibrating Pseudolite Array prototype developed in this research. This prototype includes all of the basic functionality required in an SCPA, and is both reasonably portable and easily operable by a single user. Although COTS components were used as much as possible, the unique nature of the SCPA has required a fair amount of hardware design, and has necessitated

a completely new software application — the GPSMixer — to manage the data from the GPS transceivers. The next chapter presents results from the testing of this experimental prototype: results which validate the full navigation and self-calibration capabilities of the SCPA.

Chapter 6

Experimental Results

This chapter presents results from field testing of the new Self-Calibrating Pseudolite Array (SCPA) concept performed using the experimental system described in Chapter 5. These tests cover the full functionality of the SCPA architecture, proceeding from more fundamental underlying measurements such as a raw code- and carrier-tracking, through positioning of the mobile transceiver in the array, and finally concluding with a full self-calibration of the array. Although these experimental studies cannot explore every conceivable SCPA configuration, they do cover the essential technologies at all levels and both prove the viability of the SCPA concept and verify its potential accuracy.¹

6.1 Bidirectional Ranging

The pseudolite pulsing schemes presented in Chapter 2 are designed to allow the receivers to track the pseudolite signals in the presence of extreme near/far ratios. It is important to verify that these techniques do in fact mitigate the near/far problem and allow the transceivers to maintain tracking lock under the range of conditions expected for SCPA operations. This is most easily done by examining code- and carrier-tracking performance at both close and long ranges.

¹Subsets of these experimental results have been published in ([35][37][38]).




Figure 6.1: Measured Range Data

Measured code- and carrier-range data from bidirectional ranging between stationary transceivers (with pulsing). The three transceivers are configured in a 10 meter triangular array, and the carrier-phase integers are assumed to be known. Data were collected continuously for 30 minutes. A brief period of signal loss was experienced around $t=800$ seconds; the resulting cycle slips have been removed.

6.1.1 Stationary Ranging

Figure 6.1 presents measured code- and carrier-range data from a single long-duration test of three stationary transceivers. The code data are reasonably noisy, with a standard deviation of approximately 3.7 meters. This is considerably greater than the standard satellite error model without Selective Availability (SA) of 1.4 meters presented in [53], both because of the relatively poor oscillator in the pseudolite when compared with the satellite clocks and also because of tracking degradation due to the low signal-to-interference ratio (S/I) and the short pulse duty cycle. Table 6.1, which compares code- and carrier-range tracking precisions both with and without pulsing, clearly shows a strong code-tracking degradation of about a factor of three in the presence of pulsing. The benefits of extreme near/far operations with pulsing, however, outweigh this drawback. Greater code-precision can be simply obtained through long-period averaging of the raw code-range measurements. Code noise may also be reducible through the integration of the receive amplifier into the antenna and through a cleaner design of the radio-frequency (RF) chain leading up to the receiver.

	Code (m)	Carrier (cm)
Pulsing	3.66	0.45
No Pulsing	1.23	0.76

Table 6.1: Range Measurement Stability

The standard deviations of the code- and carrier-range data from bidirectional ranging between transceivers, both with pseudolite pulsing (3% duty cycle) and without. Data with pulsing correspond to the plots in Figure 6.1. Data without pulsing come from a previous 12-minute test using 4 transceivers (6 pairings) [36].

Code biases are more difficult to measure accurately, because it requires a long averaging process to remove the random variation. Biases in the current experimental system can be as large as 5 meters, although they typically are observed to be less than two meters. Much of this bias — and the uncertain quantification — are a result of the physical construction of the prototype system. For example, in order to maintain portability many of the RF cables must be routinely disconnected, causing the line biases to change with every experiment. A more integrated transceiver architecture could largely eliminate this variability and improve the ability to accurately characterize the biases. A large part of the bias may also be multipath related. The current level of range bias is not a great problem for the Quadratic Iterative Least Squares (QILS) algorithm presented in Chapter 4, provided that the array is of large size (i.e. $\gtrsim 50$ meters). For smaller arrays, such as those used in these experiments due to space considerations, it becomes more important to use the multi-estimate solution method to guarantee accurate calibration.

The carrier-range data, which are used for the entire self-calibration process except for the initial estimate, are extremely stable even in the presence of pulsing. The standard deviation is less than one centimeter over extended time periods, and most of the biases can be eliminated during the array self-calibration process. The remaining potential error sources include non-uniformities in the antenna phase pattern, multipath, and undetected cycle slips.

The data presented in Figure 6.1 have been corrected to account for a brief signal dropout on most channels close to halfway through the test. Although signal dropouts and cycle slips are not frequent, they do occur intermittently, especially during longer tests. Some of these — such as the one in the data shown — are a result of interference when the

free-drifting pseudolite oscillators cause the separate pulses to briefly overlap. Others may be due to outside interference, rapid temperature changes in the RF electronics due to wind gusts, or are otherwise unexplained. Regardless, it is necessary to have software filters to detect when these dropouts occur and to maintain the correct carrier integers during the outage. The availability of redundant range measurements or other sensors can greatly aid in this task. Without these, some knowledge of the transceiver dynamics (i.e. stationary or slowly moving) must be used.

6.1.2 Dynamic Ranging

Figure 6.2 demonstrates the transceiver ranging capability under dynamic conditions during an array circumnavigation by the mobile transceiver, which was mounted on the K9 rover. The clockwise trajectory was driven open-loop by a human operator and attempts to bring K9 back to its starting location. Plot (a) shows the path traveled by K9, and plot (b) shows ranging data between the mobile transceiver and one of the stationary transceivers. As expected, the code-range measurements are somewhat noisy and show strong multipath signatures when the mobile transceiver is in line with both the stationary transceivers and the surrounding fence. The general motion trends are readily apparent, however. The carrier-range data are excellent, and show K9 returning to within centimeters of its initial starting point. Cycle slips during the trajectory were automatically corrected through the use of redundant range measurements to the stationary transceivers.

6.1.3 Maximum Range

The maximum achievable range of the transceivers is of special importance because it limits the operational size of the array. Figure 6.3 shows code- and carrier-range data between a pair of transceivers during one test of the system. Signal lock was maintained both on the collocated pseudolites and also on the distant pseudolites out to a range of slightly over 150 meters, yielding a dynamic range of well over 40 dB. This is much greater than the inherent GPS signal separation of 21.6 dB, and shows the ability of short, high-power pulses to alleviate near/far signal interference. Although this is an atypical example, other experiments exhibit reliable tracking to approximately 30-50 meters on a regular basis, which is approximately the size of the field test area at NASA Ames. The limiting range factors are the broadcast power and the noise figure of the RF chain, both of which can be improved upon through modifications to the transceiver design.

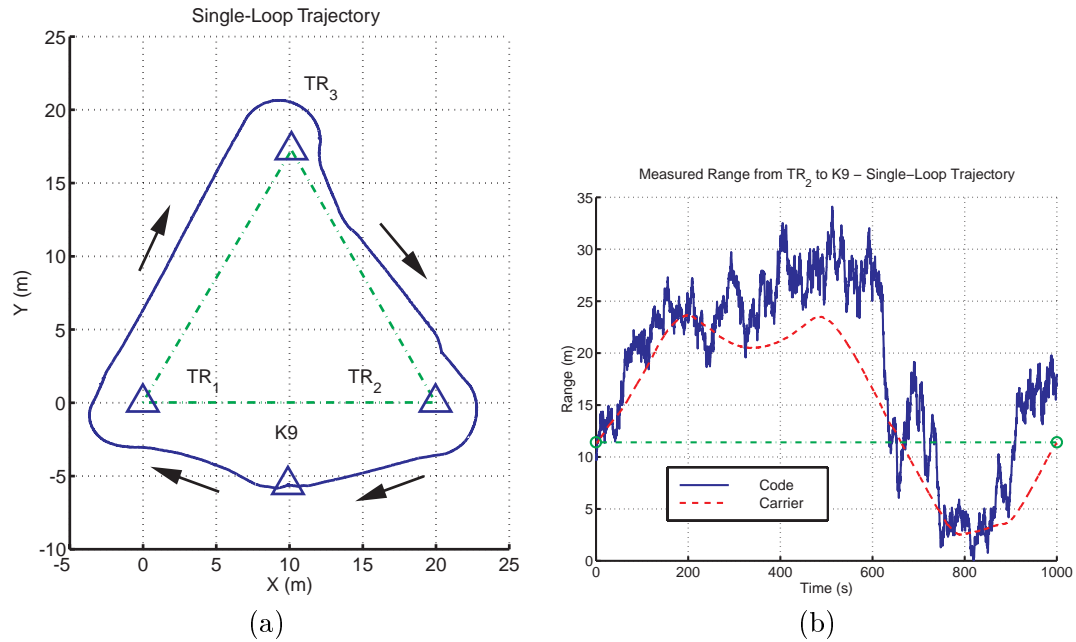


Figure 6.2: Dynamic Range Data

Plot (a) shows the configuration for this test, which was conducted using the K9 rover at NASA Ames. The mobile transceiver onboard K9 starts at the location marked at the bottom of the array and then moves clockwise around the three stationary transceivers in the trajectory shown. The locations of the stationary transceivers are known, as are the trajectory starting point and the associated carrier-phase integers. The trajectory was driven by a human operator; hence the concave path from oversteering on the right side of the plot.

Plot (b) shows raw code- and carrier-range data between the mobile transceiver and stationary transceiver #2 during this experiment. The code-range data show significant random noise in addition to large jumps at approximately 650, 750, and 900 seconds. These jumps — which occur at $\sim 90^\circ$ bearing intervals — are probably due to multipath reflections from the chain-link fence (which acts as a reflector near 1.6 GHz) surrounding the test area. The carrier-range data are very smooth, are mostly immune to the multipath, and correctly show K9 returning to its starting location at the conclusion of the trajectory.

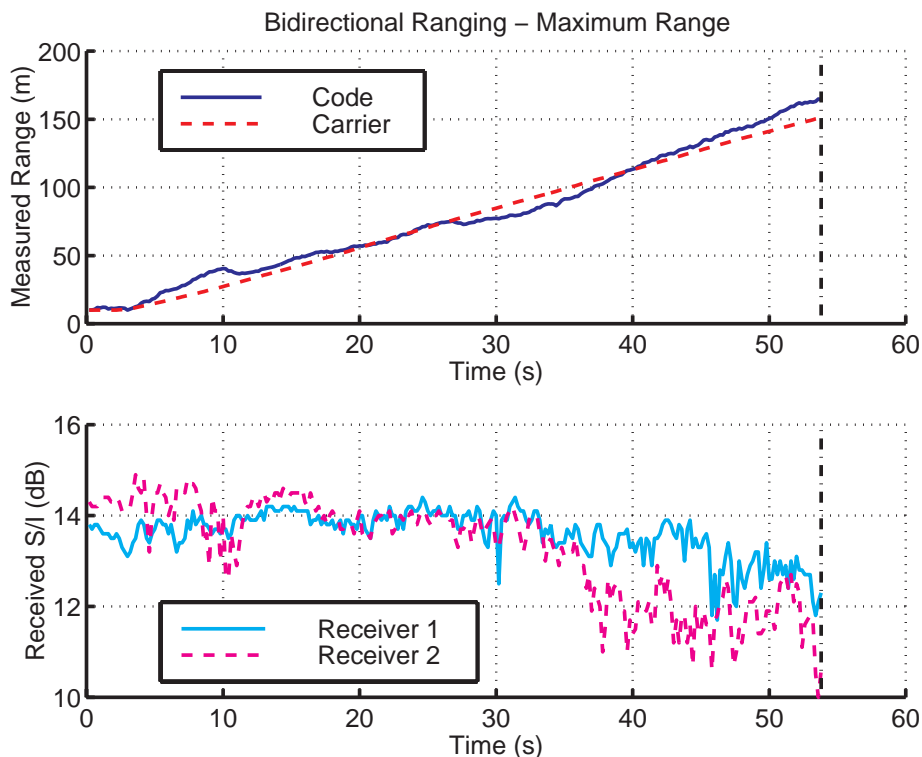


Figure 6.3: Maximum Range

Results from a maximum-range test conducted on Roble Field at Stanford. One transceiver is stationary and a second is started at 10 meters separation. This transceiver is then moved away from the first transceiver at a speed of roughly 3 m/s until it suffers an instantaneous loss of tracking lock at approximately $t = 54$ seconds (vertical dashed line). The mean broadcast signal levels were approximately $0.25\mu\text{W}$ per pseudolite.

Plot (a) shows the raw code- and carrier-range data. Signal tracking was maintained until the range was greater than 150 meters. Plot (b) displays the calculated signal-to-interference ratio (S/I) of each pseudolite as detected by the opposite receiver. The pulsed signals result in a nearly constant S/I throughout most of the range of operation. Once outside of the effective range, the S/I drops quickly with a corresponding loss of lock.

Experiment has shown that the receivers sometimes have difficulty acquiring — but not maintaining — tracking lock on high-strength pulsed signals when several pseudolites are in very close proximity ($\lesssim 2$ meters). Although this effect is not completely understood, it is easily avoided through the use of a variable attenuator in the RF receive path: The gain is lowered until tracking is achieved, after which the gain is increased again to allow long-range operations.

6.2 Array Localization

Once ranging between the transceivers is successfully achieved it is possible to recreate the array geometry using the methods of Chapter 3, assuming that the carrier-phase integers/biases and the locations of the stationary transceivers are known. The trajectory shown in Figure 6.2a, for example, was recreated solely from the carrier-range range measurements to the stationary transceivers. Although an independent quantitative verification of this trajectory at cm-level accuracy is not available, the computed trajectory agrees in character with that observed in person and on video of the test, as well as with the K9 onboard odometry. Moreover, the final calculated position of K9 after the trajectory is within 5 cm of the starting position. Some of this error is due to placement error by the human operator, while most of the rest is caused by the clock mis-synchronization described in Section 6.4.2.

Figure 6.4 shows an earlier test involving 12 continuous loops around the array by the mobile transceiver and the resulting computed ground track. To provide a truth metric, the mobile transceiver is placed at each of three pre-surveyed reference points for thirty second soaks during each loop. Before a brief loss of lock midway through the test, placement accuracy at these reference points was 4.5 cm RMS, roughly the expected placement error by the human operator. The current experimental system achieves increased signal tracking robustness and a resulting reduction in accumulated positioning errors beyond that shown in this test through the addition of power amplifiers in the receive paths. The improved system was used for the tests in Figures 6.2 and 6.3. These and similar tests demonstrate that the high accuracy and repeatability associated with CDGPS is fully applicable to transceiver systems, and that they provide a viable method of near-field positioning.

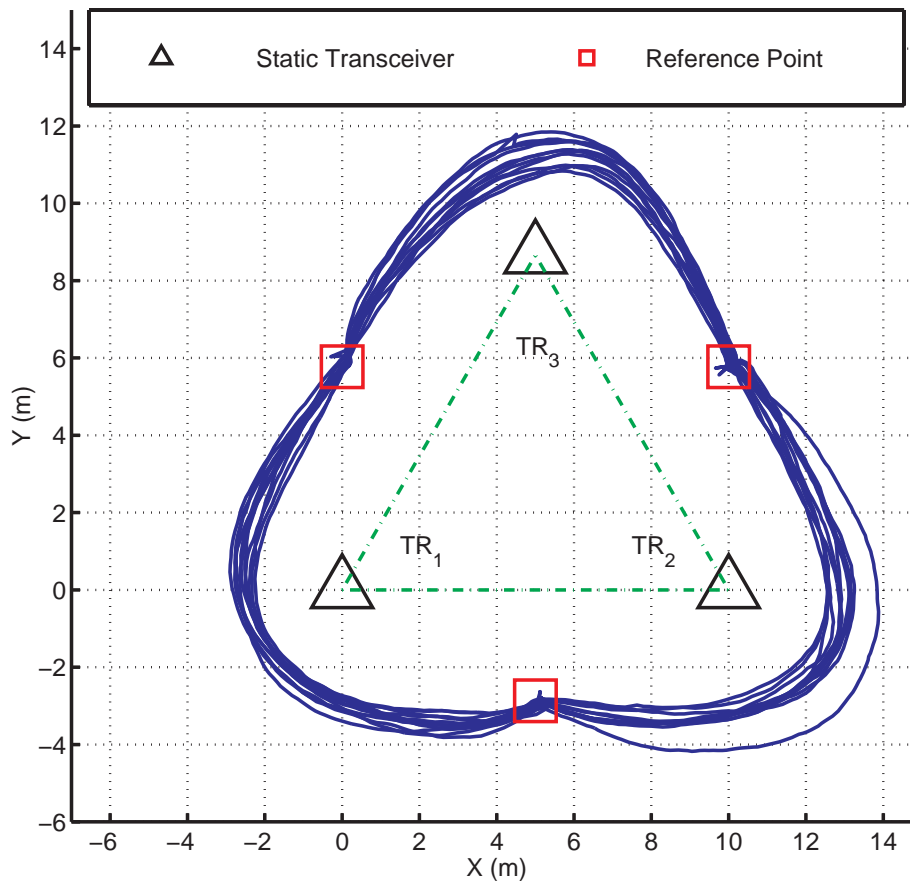


Figure 6.4: Positioning of Mobile Transceiver

Ground track of the mobile transceiver around the stationary transceiver array (hand-carried through 12 counterclockwise loops). The mobile transceiver stops briefly at each of the three reference points as a truth measurement. These points are pre-surveyed and marked on the ground. Prior to an uncorrected dropout on the 6th loop, transceiver placement accuracy at the reference points is 4.5 cm RMS, consistent with the expected accuracy for human placement.

6.3 Self-Calibration

The previous sections of this chapter have demonstrated the successful operation of a GPS transceiver system and its ability to successfully determine relative position under dynamic conditions, assuming that the carrier-phase integers and other biases are known. The methods developed by this research to determine these biases and the locations of the stationary transceivers are presented in Chapter 4. This section demonstrates the successful self-calibration of a completely unknown array using these methods.

The experiment described here corresponds with the configuration and provides the data shown in Figure 6.2a, with the stationary transceivers in the 20-meter array shown and K9 starting and finishing its trajectory at the indicated point. The trajectory is driven open-loop. Although the actual positions of the stationary transceivers are known as a truth reference, this information is not used during the self-calibration process. Odometry from wheel encoders onboard K9 provides an additional point of comparison for the trajectory itself.

6.3.1 Coarse Calibration

The first stage of the calibration process is utilization of the raw code-range measurements between the transceivers to generate a rough estimate of the starting location of the array. Because of the high noise level, these code data are averaged for roughly 15 minutes before being applied to the array. This removes most of the random noise, but retains any biases due to static multipath or other factors. Figure 6.5 shows the initial positions of the four transceivers as derived from the code-range measurements. The actual positions of the stationary transceivers are at the corners of the triangle shown, while K9's actual starting position is near the coordinate (10.0,-5.6). Table 6.2 lists the corresponding positioning errors. From this data, the locations of the stationary transceivers are estimated to be within roughly 3-4 meters of their true locations, or 20% of the size of the array. This is approximately the value expected, and would generally be solvable using the ILS algorithm alone. Larger arrays would be even less effected by biases of this size. The initial estimate of K9's location, however, is off by greater than 20 meters, most likely due to strong local multipath. Unlike in satellite-based GPS, such multipath does not average out because there is currently no relative motion between the transceivers. The QILS algorithm was developed to converge even in the presence of large biases of this magnitude.




Figure 6.5: Pre-Calibration Position Estimates

Positions of the stationary transceivers and the ground track of the mobile transceiver on K9 as determined by the SCPA prior to carrier-phase calibration. The locations of the stationary transceivers are determined via average code-range measurements: The actual positions are at the corners of the triangle shown.

The dotted line shows the trajectory followed by K9 as determined by its onboard odometry, departing from the point (10.0,-5.6). The actual trajectory returns to its starting point.

The solid line shows the trajectory as determined by (uncalibrated) carrier-phase ranging to the stationary transceivers. Although K9 (correctly) returns to its perceived starting location, the large biases resulting from the poor initial position estimate make the estimated trajectory itself nearly useless.

TR #	X (m)	Y (m)
1	3.38	-2.13
2	-0.87	-2.13
3	-2.52	4.27
K9	20.42	0.60

Table 6.2: Initial Position Errors (Code-Phase)

Positioning errors of the three stationary transceivers and the K9 rover start/finish location before array self-calibration, as measured with respect to the pre-surveyed reference points at the test area. RMS positioning errors for the stationary transceivers are 2.76 m.

These data correspond to Figure 6.5.

Once the code-range array estimate is obtained, the initial carrier-phase integers are set based upon the estimated transceiver locations. K9 then makes its trajectory around the array, taking carrier-phase measurements as it moves. This is analogous to using carrier-smoothed code for ranging during the trajectory. K9's ground track is computed based upon these carrier-ranges, and is also shown in Figure 6.5 (solid line). Because the initial estimate of K9's starting location was so poor — yielding very poor integer estimates and large range biases — the calculated trajectory shows almost no resemblance to the actual path traversed. Large jumps are evident where triangulation is simply impossible. It is this remarkably poor estimate of the rover trajectory that is passed on to the QILS algorithm.

The rover trajectory as derived from the onboard odometry is also shown in the same plot. Although it shows the characteristic drift, it does provide at least a recognizable estimate of the actual trajectory. This estimate could also have been used as the seed for the QILS algorithm. An operational Mars rover would have additional sensors such as an inertial measurement unit (IMU), sun-sensors, and computer vision with which to generate a trajectory estimate, and when combined with the drift-free carrier-range measurements this integrated navigation system could provide a very good initial estimate. The current example utilizes only the GPS transceiver ranging itself, however, in order to show system performance under worst-case conditions.

TR #	X (m)	Y (m)
1	-0.01	0.03
2	-0.05	0.03
3	0.06	-0.05
K9	-0.16	0.09

Table 6.3: Final Position Errors

The final positioning errors of the three stationary transceivers and the K9 rover start/finish location following the array self-calibration, as measured with respect to the pre-surveyed reference points at the test area. (See Figure 6.6.)

The RMS position error for the stationary transceivers is 4 cm, and the K9 start-location error is indicated to be 18 cm. Much of this error is due to the clock-synchronization problem described in Section 6.4.2. In addition, measurement of these reference points and technical placement error of the transceivers, together with uncertain knowledge of the precise antenna phase centers, may account for up to 5 cm of total system error.

6.3.2 Fine Calibration

The QILS self-calibration algorithm utilizes the initial estimates of the stationary transceiver positions and the rover trajectory and attempts to determine the actual positions, trajectory, and the corresponding range biases. Figure 6.6 shows the resulting position and trajectory estimates after multiple applications of the algorithm. The multiple applications enable the system to correct cycle slips even in the presence of large measurement biases, a technique which is described more fully in Section 6.4.1. Table 6.3 presents the corresponding position errors after self-calibration, while Table 6.4 lists the range biases removed during the process. It is apparent that the QILS algorithm correctly determines the locations of the stationary transceivers to centimeter-level accuracy, and that the final trajectory estimate agrees well with the character of the onboard odometry. The algorithm was able to remove range biases greater than 75% of the array size, and converged despite position errors greater than 100% of the array size. It was also able to overcome the triangulation problems during the coarse estimation process that were noted previously.

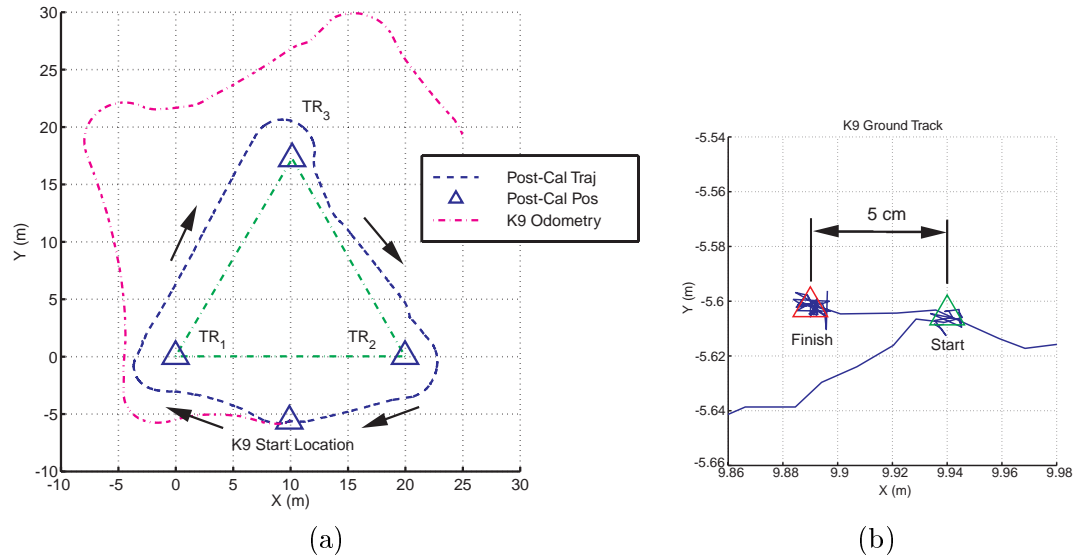


Figure 6.6: Post-Calibration Positioning

Positions of the stationary transceivers and the ground track of the mobile transceiver on K9 as determined by the SCPA following carrier-phase calibration. Compare plot (a) with Figure 6.5. The new dashed line shows the post-calibration trajectory, and the small triangles indicate the perceived transceiver locations, all of which are accurate to a few centimeters.

Plot (b) zooms in on the starting/finishing location of the rover trajectory. The starting point is on the right, and the finishing point is on the left. These points are separated by only 5 cm, roughly the accuracy that the K9 operator can achieve when driving the rover to a specified point. (This technical error is discussed in Section 6.4.3.) The small separation of the two points shows the repeatability of the CDGPS measurements over long periods of dynamic operation.

Removed Ranging Biases (m)			
TR #	1	2	3
2	4.16		
3	-4.18	-6.50	
K9	-15.88	-0.29	-11.42

Table 6.4: Final (Removed) Ranging Biases

The final estimated range biases between the four transceivers after self-calibration. The values represent the difference between the range estimates determined by using averaged code measurements and those derived from carrier-phase self-calibration, and include both integers and line biases. These estimated biases were subsequently removed from the range data, resulting in the final position errors listed in Table 6.3.

6.4 Error Sources

Despite the great success of the self-calibration algorithms, examination of the results above shows that some centimeter-level errors remain even after completion of the self-calibration process. These are caused by the particular hardware configuration utilized by the experimental system rather than by fundamental limitations of the self-calibration algorithms themselves. Three error sources are of particular interest: cycle slips, clock drift, and truth-system error. These sources and their implications are described below.

6.4.1 Cycle Slips/Signal Dropouts

Cycle slips were mentioned previously in Section 3.4.3 with respect to the general array-positioning problem and are of great importance during the self-calibration process. The self-calibration algorithms rely upon the integers remaining constant during the entire trajectory. If a cycle slip occurs, an attempt is made to recover the integer through the use of redundant range measurements and the last measured position and bias values. If the signal loss is more than instantaneous, however, additional small range biases are likely to be introduced in mid-trajectory. Because these bias changes are not included in the system model they will adversely affect the final array estimate. Moreover, large positioning errors at the time of the slip will generally cause larger errors in the slip correction, making the self-calibration process more vulnerable to cycle slips than basic transceiver positioning alone.

In the experimental results previously presented, the effects of cycle slips have been greatly mitigated through repeated applications of the QILS algorithm on the raw measurement data. During each application the state estimates greatly improve, allowing for more effective correction of the cycle slips during the next application. This process gradually eliminates the errors associated with cycle slips in an iterative fashion as the trajectory estimate approaches its true value. Because the cycle-slip filters act upon the raw range measurements themselves, each iteration requires re-examination of the entire collected data set and therefore takes several minutes to complete. Self-calibration is normally a one-time event that is performed off-line, however, so this processing time is generally not a concern. Figure 6.7 shows how the positioning errors for the experiment described in Section 6.3 were reduced through multiple applications of the QILS algorithm. In this case the slip errors are not entirely eliminated due to the presence of the clock-drift errors that are described in detail in the next section.

Although this current method of iterative cycle-slip correction is reasonably effective, it would be unwise to rely upon repeated self-calibrations to eliminate cycle-slip errors from SCPAs deployed in the future. First of all, the presence of additional range-measurement error sources can corrupt the process, regardless of the number of iterations performed. Second, multiple slips or dropouts occurring simultaneously — which may occur if the pseudolite pulses overlap briefly and cause widespread local interference — can be beyond the ability of this method to correct because of the inadequate number of redundant range measurements with which to estimate the integer values. It is therefore highly desirable to have an integrated navigation system containing inertial and/or other sensors that can coast through brief GPS dropouts.

6.4.2 Clock Drift

One of the dominant error sources in the experimental data presented above is carrier drift. Ideally a GPS system exhibits no drift, since time-dependant clock errors are eliminated in the double-difference procedure. As was mentioned in Section 3.4.4, however, this is only true if the raw carrier- and code-phase measurements at each receiver are taken simultaneously. During the testing at NASA Ames a hardware failure in the primary RF front end of receiver #3 forced a switch to the backup front end, which was not correctly programmed to synchronize to the master pseudolite signal.

Figure 6.8a shows the effects that resulted from this improper synchronization. The

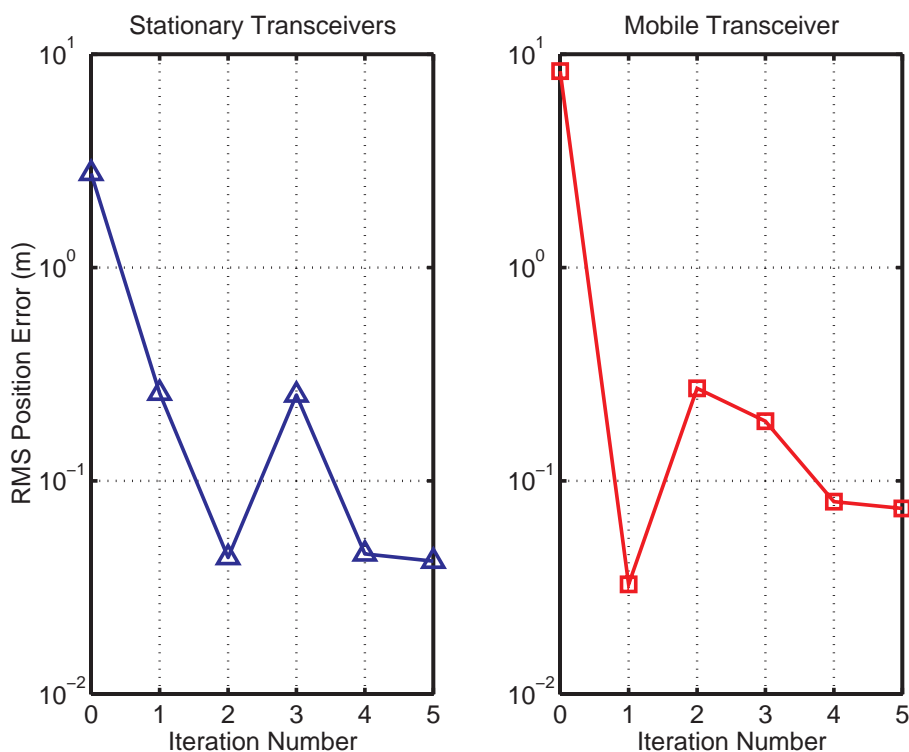


Figure 6.7: Cycle-Slip Removal Through Multiple Self-Calibrations

RMS position errors for (a) the stationary transceivers and (b) for the mobile transceiver start location during the experiment in Section 6.3 as a function of the number of iterations through the self-calibration process. Each iteration is performed on the same range data, with the output of the previous iteration as the new initial state estimate.

Cycle-slip correction depends upon accurate knowledge of the true transceiver locations. As the state estimates improve, the overall system-wide cycle-slip correction generally improves as well, although the error in individual states may increase or decrease from iteration to iteration.

plot shows carrier-ranging measurements between the stationary transceivers in the array after the integers/biases have been removed via the self-calibration process. The measured range from transceiver #1 to transceiver #2 remains constant as expected for the duration of the 15-minute test, with the exception of a partly-corrected cycle slip midway through. The measured ranges from transceiver #3 to both transceiver #1 and transceiver #2, however, exhibit a very slow drift, accumulating up to 10 cm of error over the course of the test. This occurred because receiver #3 was not synchronizing its sample time to the master pseudolite and was actually freely drifting with respect to the other receivers. Because of the slow nature of the drift, this error was not detected until the data were more closely examined during postprocessing. It is this drift in the measured ranges that prevents complete correction of the cycle slip, and that ultimately causes most of the final positioning error presented in Table 6.3. Note that natural averaging in the self-calibration algorithms causes these range errors to have a greater effect on the estimated location of K9 compared with the locations of the stationary transceivers.

The magnitude of the expected drift when the clocks are improperly synchronized can be approximated using Equation 3.15, which is repeated below.

$$\Delta R \lesssim \left(\frac{\delta f}{f} \right) \cdot \Delta T \cdot c \quad (6.1)$$

In this case

$$\Delta T = \Delta T_{drift} = \left(\frac{\delta f}{f} \right) \cdot (T_{final} - T_{initial}) \quad (6.2)$$

so that the maximum expected difference in sampling times increases roughly linearly with time in proportion to the oscillator stability. This is in contrast with the synchronized system, in which ΔT is a constant value that depends on the range from the receiver to the master pseudolite. This results in the bounding expression for the range error due to drift presented in Equation 6.3, which also increases in a linear manner with time.

$$\Delta R_{drift} \lesssim \left(\frac{\delta f}{f} \right)^2 \cdot (T_{final} - T_{initial}) \cdot c \quad (6.3)$$

For the 1000-second duration of the test presented here, and with oscillator stabilities of approximately 10^{-6} , the maximum expected error is 30 cm. Thus the actual observed drift of 5-10 cm agrees well with expectations.

Figure 6.8b shows test results using the same experimental hardware and configuration

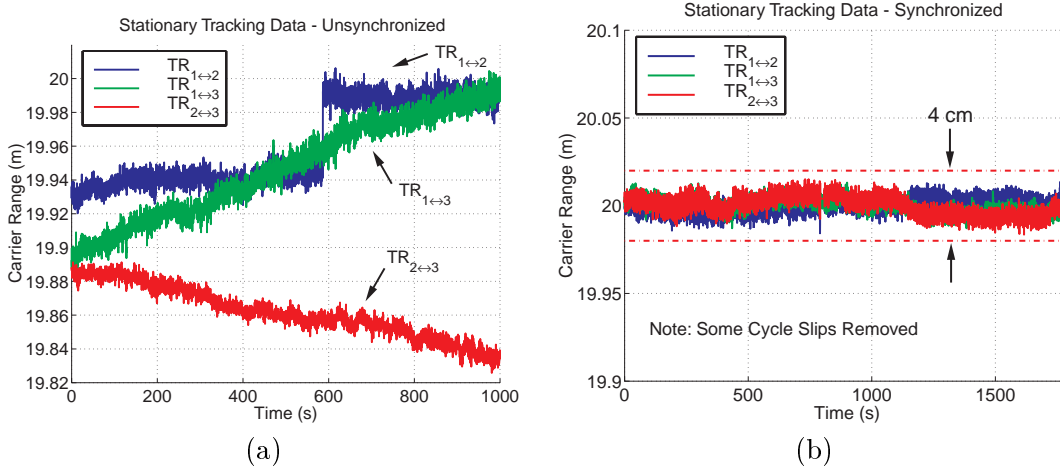


Figure 6.8: Receiver Synchronization Errors

Plot (a) shows carrier-phase ranging between three stationary transceivers. Receiver #3 is improperly synchronized to the master pseudolite, resulting in a slow drift in all of its resulting range measurements. Plot (b) shows similar data — taken during a different experiment — after receiver #3 has been properly synchronized: The drift has been eliminated.

after receiver #3 was upgraded with new code so that all of the channels were synchronizing to the master pseudolite. In this case, as expected, no carrier drift is observed.

6.4.3 Truth-System Error

The final known error source in these experiments does not relate to the accuracy of the GPS system at all, but rather to the truth system used to measure its accuracy. Continuous, centimeter-level, drift-free position determination over long distances is still limited to a few special technologies; this need is of course the motivation for the development of the SCPA. Among those technologies deserving consideration as a truth system, laser metrology would potentially have been acceptable if it were available. Because the tests were conducted outdoors on Earth, satellite-based CDGPS is also a very attractive possibility. The receivers employed in this research were not configured to receive satellite signals, however, and the current dipole antenna pattern gives poor reception of sky-based signals. This necessitates the use of another metrology system.

Because the raw positioning capabilities of carrier-phase GPS are a well-documented, it

is only necessary to have truth data at a few dispersed locations; once positioning accuracy is established at these points, it can reasonably be interpolated to be of similar character and accuracy at the intervening locations. Pre-surveyed reference points were therefore utilized to provide the system truth reference. These reference points include the locations of the stationary transceivers, the starting/finishing location of the rover trajectory, and a few other points along the trajectory. The points are surveyed using a simple tape measure, resulting in a reference accuracy of approximately 2-3 cm depending somewhat upon the flatness of the terrain.

Besides this uncertainty in the true locations of the reference points, additional uncertainty comes from the placement of the transceivers themselves. The antennas are located on tripods approximately 1.5 meters above the ground, and the use of plumb-bobs for accurate placement above the reference points was hampered by the high winds at the test site. Additionally, the actual phase centers of the antennas themselves are in doubt. These factors result in another 2-3 cm of potential placement error.

Although these technical error sources are individually small, together they mean that truth cannot be verified to better than approximately 3-5 cm. Since this is worse than the accepted accuracy of CDGPS systems and is on the same order as most of the observed error, there is a strong potential that the self-calibration process itself — when the clock synchronization is handled correctly — is in fact even more accurate than the final positioning-error data from these experiments would indicate.

6.5 Summary

The experimental data presented in this chapter prove the feasibility of the SCPA concept invented in this research by verifying the successful operation of each level of the prototype system: bidirectional ranging, positioning of the mobile transceiver, and full array self-calibration. Even with very low broadcast power, successful tracking at ranges of up to 150 meters have been observed, overcoming near/far ratios greater than 40 dB: Even greater ranges should be achievable through the use of greater broadcast power and a cleaner and more integrated transceiver design, which would reduce the overall RF noise figure. Repeatable relative positioning has been achieved with accuracies of better than 5 cm, the same order of magnitude as the error in the truth system. Finally, complete array self-calibration has been accomplished through an array circumnavigation by a mobile transceiver, as is

shown in Figure 6.6 and Table 6.3. Again, the resulting placement accuracies are better than 5 cm RMS, even in the presence of known clock synchronization errors.

Chapter 7

Conclusions

Prior to the research presented in this dissertation, there was no known method to determine accurately the location of rovers moving among stationary GPS pseudolites in uncertain positions. Surveying the locations of these pseudolites required either precise knowledge of the location of the rover along a special calibration trajectory, or a completely independent measurement system.

The new concepts and techniques developed in this dissertation show that it is in fact possible to survey the locations of stationary pseudolites and determine the associated carrier-phase integers by using only ground-based GPS transceivers together with imprecisely-known relative rover motion. The resulting Self-Calibrating Pseudolite Array (SCPA) is well suited to a variety of robotic applications that cannot use the GPS satellite constellation, Mars exploration among them. This chapter summarizes the main contributions of this research, and then proceeds to discuss additional factors that must be considered in order to successfully apply this research to Mars exploration. Finally, it presents potential system enhancements that could greatly improve SCPA utility in the future.

7.1 Contributions and Results

Although the specific contributions of this research were presented in detail in Section 1.5, it is instructive to reiterate the most significant points in light of the information presented in Chapters 2-6:

- Invented the concept of a Self-Calibrating Pseudolite Array: an array of GPS pseudolite transceivers that is able to self-survey the locations of all devices without a priori

position knowledge or other sensor information. Demonstrated its feasibility through simulations with a wide range of configurations and starting estimates, and achieved complete experimental verification of the system concept through the construction and testing of an operational prototype.

- Identified the major factors affecting SCPA operation and self-calibration, such as array geometry and multipath errors. Developed strategies including positioning heuristics and trajectory design to overcome these factors.
- Developed nonlinear and stochastic self-calibration algorithms (multiply-seeded Quadratic Iterative Least Squares) that use limited transceiver motion to provide 100% calibration success under nominal operating conditions, and greater than 99% calibration success under worst-case conditions. These worst-case conditions include multipath-induced code-phase positioning errors as large as the array dimensions themselves. Demonstrated the effectiveness of these algorithms on the experimental system.
- Advanced the theory relating to GPS pseudolite transceivers, and to the use of pseudolite pulsing as a solution to the ‘near/far’ problem.
- Using the operational prototypes, demonstrated the viability of GPS transceivers for local positioning. Developed a simple self-differencing transceiver architecture for use with commercial off-the-shelf (COTS) components, and showed the ability of pseudolite pulsing to overcome the extreme ‘near/far’ problems associated with transceiver operations. Proposed the use of bidirectional ranging as the fundamental building block for transceiver operations, and experimentally demonstrated its feasibility.
- Experimentally demonstrated positioning of a mobile transceiver within a stationary transceiver array to centimeter-level accuracy using bidirectional ranging, and successfully self-calibrated the entire system to similar accuracy through the use of that mobile transceiver. No *a priori* knowledge of the array configuration beyond that provided by this GPS-based transceiver system was used in the self-calibration process. This is the first time the operation and self-calibration of a stand-alone GPS transceiver system has been accomplished.

7.2 Mars Considerations

Although the prototype system and algorithms prove the feasibility of an SCPA, there are a number of complementary details that must be considered before such a system will be suitable for a Martian environment, with its very challenging conditions. Some of the most critical ones are listed below.

- **Power:** Although the transceivers use relatively little energy, power is a special concern for long-duration missions. This is because solar cells decrease in efficiency over time due to coating with atmospheric dust. Either methods must be developed to keep solar panels free of dust, or other power sources such as radiothermal generators might be required.
- **Thermal control:** This is an issue when dealing with precision timing equipment. Although the bidirectional ranging process inherently corrects for first-order clock drift, extreme temperature swings — especially between different transceivers — could potentially cause the receivers to lose tracking lock on the pseudolite signals. More likely, frequency changes due to temperature differences would merely cause the signal acquisition process to become more intensive, and appropriate wide-band frequency search methods would have to be employed.
- **Multipath:** Multipath from natural terrain is very difficult both to predict and to avoid. Although ground-based arrays can get strong multipath reflections from the surface, these are mostly parallel with the direct line-of-sight path and so contribute little ranging error. Placement of the array near cliffs or other vertical obstructions, however, should be avoided if possible because of the large associated signal delays. This is primarily important for code-based ranging. Also, care should be taken with the design of the rovers and other hardware to limit the amount of multipath from their solar arrays or other structures. Note that a wide range of multipath-mitigation algorithms exist, some of which may prove effective under these conditions.
- **Terrain flatness:** Although the analysis in Chapter 3 shows that positioning errors are small even in the presence of vertically uneven terrain, in some situations this error may become significant. In such cases the positioning and self-calibration algorithms could be made more accurate by incorporating estimates of the component altitudes

in the calculations, although there would likely be insufficient observability to actually calculate the altitudes using the SCPA alone.

- Obstructions and dropouts: If the array has been deployed in uneven or rocky terrain, there is a high likelihood that the signal paths between the rover and the stationary transceivers may be obstructed at various points. After calibration this is not a great problem if there are redundant range measurements: the integers can be relatively easily recovered. During calibration, however, the current algorithms cannot function with lengthy signal dropouts. It is probable that an array with a greater number of transceivers for extra redundancy would be able to coast through intermittent dropouts even during the self-calibration process, especially with aiding from inertial sensors. The algorithms with which to accomplish this have yet to be developed, however, and present an area of future research.

7.3 Future Directions

The Self-Calibrating Pseudolite Array employs a very flexible, modular architecture. This makes the basic idea well-suited to expansion beyond the basic implementation presented in this dissertation. This section presents four such possible extensions. While nominally intended for Mars exploration, these variations could be applied to many other potential applications.

7.3.1 Signal Structure

This research has utilized a standard GPS pseudolite signal structure consisting of the 1 Mbps C/A code modulated upon the 1575.42 MHz L1 carrier. This signal was used to maintain compatibility with the hardware from earlier research projects and to minimize the system development time. Because GPS chipsets and other RF components tuned to L1 are readily available, this also reduces system cost. The utilization of such widely-used — and in some cases space-tested — COTS components would also ease incorporation into an actual flight mission by minimizing development time.

Some advantageous changes to signal structure could still be made, however, with relatively little effort. Increasing the code chipping rate from 1 to 10 Mbps (similar to P-code) would reduce the raw code-ranging errors by an order of magnitude, greatly improving the

ability of the system to operate before — or even in the complete absence of — carrier-phase self-calibration. It would also lower the susceptibility of the array to multipath.

Another modification that would be extremely beneficial would be the inclusion of additional frequencies in the pseudolite signal. This would enable Cascaded Integer Resolution (CIR) for the determination of the carrier-phase integers, as is described in Chapter 4. Although the addition of these extra frequencies complicates the radio-frequency (RF) design of the transceivers, the potential benefits are very great. Especially when combined with a higher chipping rate, it would be possible to self-calibrate the array without the necessity of transceiver motion. This is in fact the strongly-recommended course of action for an actual Mars-based array. In this case the self-calibration methods presented in this dissertation could still be used to remove residual line biases due to temperature and other effects and improve the overall accuracy of the system.

Other signal changes could be made as well. For example, some of the more recent Ultra-Wideband (UWB) techniques might be employed as part of the system. UWB-based ranging potentially approaches the accuracy of carrier-phase differential GPS (CDGPS), and exhibits excellent multipath rejection. One Stanford proposal by Opshaug and Enge for integrating UWB with GPS appears in [50], and several companies are currently working on developing UWB-only positioning systems. In fact, because of the absence of FCC (Federal Communications Commission) or similar regulations on Mars and the small number of potentially conflicting users, the signal structure of SCPAs could be modified almost arbitrarily to suit mission requirements. Maintaining some compatibility with the current signal structures, however, may still be desirable in order to reduce costs.

7.3.2 3-Dimensional Arrays

The array configuration used for these studies is planar. Although small vertical displacements of the transceivers may introduce small horizontal positioning errors, the vertical displacements are themselves essentially unobservable. This is not the case, however, if one of the transceivers is moved considerably out-of-plane. In such a situation the SCPA becomes a full 3-dimensional positioning system. Most likely this would be accomplished through at least one extra stationary transceiver, either positioned on a high hill or cliff overlooking the main array or perhaps suspended from a tethered balloon such as those proposed for Mars use by Jones [25]. Adding this extra capability would not only reduce positioning errors due to out-of-plane effects, it could also improve signal coverage in the

presence of ground-based obstructions and obstacles such as boulders, hills, and canyons.

The positioning and self-calibration principles developed in this dissertation would still be largely applicable to the 3-dimensional geometry, although the system would be somewhat more complicated. Preliminary simulations involving an extra (possibly mobile) transceiver situated above the 2-dimensional array used for this dissertation indicate that a similar motion of a ground-based rover is sufficient to determine the integers in the bidirectional ranging measurements to this aerial transceiver, and hence determine its relative position [40]. It is not clear at this point, however, what the optimal 3-dimensional array configuration would be. This is a subject for further study.

Another interesting extension involves the incorporation of mobile aerial vehicles into the array. One significant example would be active positioning of landers and other descent vehicles on their arrival to Mars. Such a capability would be greatly beneficial when building up large-scale infrastructure at a given site on the surface. Rather than landing each component at an uncertain location and then transporting the components overland to the central site, each lander could be guided down by the transceivers in the array, potentially landing within centimeters of its target location. This would create a Mars analog of the Integrity Beacon Landing System/Local Area Augmentation System (IBLS/LAAS) under development for autoland aircraft on Earth. Although it is assumed that the ground-based transceivers would be pre-surveyed for these types of operations, more research must be done into the ability of such a system to determine the associated carrier-phase integers during the short duration of the descent.

7.3.3 Large Arrays

It has been previously noted that having additional stationary transceivers beyond the minimum of three is strongly desired in order to reduce error, to provide redundancy, and to add robustness in the case of signal dropouts. Larger and denser arrays can also be greatly beneficial by improving coverage in very uneven terrain by broadcasting behind hills and into canyons. It is still necessary to develop effective schemes for the integration of additional redundant transceivers into the array to take advantage of this potential.

The prototype transceiver used for this research could not be simply replicated and used en masse, because the pseudolite pulses would soon overlap significantly and prevent successful tracking. There are three solutions to this problem. The first is to use a pulse synchronization signal such as is described in Section 5.1.4 to give each pseudolite its own

independent timeslot. This simple time-division multiple-access (TDMA) scheme allows continuous operation of all the transceivers, but fundamentally limits the total number of devices in the array. A 5% pulse duty cycle, for example, permits at most 20 pseudolites in the array.

Larger arrays demand a slightly more sophisticated approach. Rather than having the entire array broadcasting simultaneously, smaller portions or cells may be used in analogous fashion to cell-phone systems on earth. Mobile transceivers would operate within a given cell, in which all of the transceivers are pulsing continuously. As the rover transitions to another cell, handoff would occur and the new cell would begin broadcasting. Integer handoff could be difficult in such a system, however, unless the cells somewhat overlap.

A final TDMA scheme that might be more effective — although somewhat more difficult to implement — utilizes a rover-centric approach. Rather than dividing the array into cells, positioning can occur on a rolling basis with only the closest — or alternatively the most effective in terms of Dilution of Precision (DOP) — set of transceivers broadcasting. Multiple rovers could then operate very effectively in a large distributed array, each surrounded by a floating ‘bubble’ of transmitting devices.

Handoff in these last two schemes is greatly facilitated by the bidirectional ranging scheme used instead of conventional double differencing: Each individual range measurement is independent and has its own associated integer, and so may be included or dropped in the positioning process without affecting other measurements. These TDMA methods are also more efficient in the use of power, but do have the disadvantage of forcing tracking re-acquisition with every broadcast transition.

In each of these scenarios the specific implementation details have not been addressed, nor has an in-depth analysis to determine the optimal TDMA scheme been attempted. In addition, efficient self-calibration methods for such large arrays must also be developed. If only code-level precision is desired, of course, the more difficult of the self-calibration issues are easily avoided.

7.3.4 Over-The-Horizon Arrays

The previous discussion of large arrays may be easily extended to include arrays stretching over the horizon. As long as each transceiver maintains contact with at least two other transceivers that are themselves in contact with the array, the overall continuity of the array would be maintained. Airborn elements would be especially effective at increasing the

coverage area in this manner with a minimal number of devices, because of their increased line-of-sight capabilities.

7.3.5 Mobile Arrays

Extending the SCPA concept for long traverses through the use of very large arrays quickly becomes impractical. Such capability can be provided, however, by allowing the array as a whole to move over the surface. Although nominally stationary, each transceiver would be mounted on a mobile base. The transceivers would periodically pick up and move themselves to keep the coverage area centered around the rover(s), while using positioning from the remaining stationary transceivers to maintain knowledge of the overall array location with respect to the surface. In this manner the entire array could ‘bootstrap’ itself across the surface to provide continuous relative positioning along an arbitrary trajectory. Although some global error would gradually accumulate with each successive stage of the process, relative positioning accuracy between the array devices could be maintained at all times. This mode of operations would of course be most effective when multiple rovers are cooperating, perhaps for long-distance survey or mapping tasks.

Current work at Stanford by Masayoshi and Opshaug using the experimental system described in this thesis is examining the error propagation resulting from this bootstrapping method of array extension, with application to both GPS and UWB transceiver arrays [41].

7.4 Final Thoughts

The goal of this research has been to adapt existing, commercially-available GPS hardware into a system that is capable of meeting the navigational requirements of future Mars exploration missions, with the aim of making such navigation more cost-effective and easily-implementable in the near future. Having achieved this goal, the Self-Calibrating Pseudolite Array is now a viable option for engineers and scientists seeking to do precise, repeatable positioning in locations where more-conventional navigation and sensing systems are difficult to use, are too costly, or are simply unavailable. The positioning and self-calibration methods developed as part of this research are also directly applicable to a wide range of radio-navigation systems besides GPS, and the Quadratic Iterative Least Squares algorithm in particular has potential for use in a very wide range of nonlinear estimation problems both within and outside of the field of navigation.

Bibliography

- [1] Pathfinder/sojourner: Map of end-of-day image from imp (as of sol 83). Jet Propulsion Laboratory, <http://mars.jpl.nasa.gov/MPF/roversci/site-map-End-of-day-image-from-IMP.html>, 1997.
- [2] John Carl Adams. *Robust GPS Attitude Determination for Spacecraft*. PhD thesis, Stanford University, 1999.
- [3] David J. Bell, Robert Cesarone, Todd Ely, Chad Edwards, and Steve Townes. Mars Network: A Mars Orbiting Communications & Navigation Satellite Constellation. IEEE, 2000.
- [4] Michael S. Braasch. Multipath effects. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. I*, pages 547–568. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [5] John L. Bresina, Maria G. Bualat, Laurence J. Edwards, Richard M. Washington, and Anne R. Wright. K9 Operations in May '00 Dual-Rover Field Experiment. In *Proceedings of the 6th International Symposium on Artificial Intelligence and Robotics & Automation in Space*, Quebec Canada, June 2001. Canadian Space Agency.
- [6] H. Stewart Cobb. *GPS Pseudolites: Theory, Design, and Applications*. PhD thesis, Stanford University, 1997.
- [7] Clark E. Cohen, Boris S. Pervan, David G. Lawrence, H. Stewart Cobb, J. David Powell, and Bradford W. Parkinson. Real-Time Flight Test Evaluation of the GPS Marker Beacon Concept for Category III Kinematic GPS Precision Landing. In *Proceedings of the Institute of Navigation GPS-93 Conference*, volume 2, pages 841–849, Salt Lake City UT, September 1993.

- [8] Federal Communications Commission. Experimental radio station construction permit and license #wc2xlc, 2000.
- [9] Andrew R. Conway. *Autonomous Control of an Unstable Model Helicopter Using Carrier Phase GPS Only*. PhD thesis, Stanford University, Stanford, CA 94305, December 1995. Also published as SUDAAR 664.
- [10] Tobé Corazzini and Jonathan How. Onboard GPS Signal Augmentation for Spacecraft Formation Flying. In *Proceedings of the Institute of Navigation GPS-98 Conference*, Nashville TN, September 1998.
- [11] Tobé Corazzini and Jonathan How. Onboard Pseudolite Augmentation System For Relative Navigation. In *Proceedings of the Institute of Navigation GPS-99 Conference*, Nashville TN, September 1999.
- [12] Tobé Noel Corazzini. *Onboard Pseudolite Augmentation for Formation Flying Spacecraft*. PhD thesis, Stanford University, 2000.
- [13] IntegriNautics Corporation. In200c general-purpose pseudolite signal generator: User's manual. Palo Alto CA, 1998.
- [14] R. J. Cosentino and D. W. Diggle. Differential gps. In Elliot D. Kaplan, editor, *Understanding GPS Principles and Applications*, pages 321–383. Artech House, Norwood MA, 1996.
- [15] Matthew Deans and Martial Hebert. Experimental Comparison of Techniques for Localization and Mapping Using a Bearing-Only Sensor. In *Proceedings of the 2000 International Symposium of Experimental Robotics*, 2000.
- [16] A. J. Van Dierendonck. Gps receivers. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. I*, pages 329–407. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [17] A. J. Van Dierendonck, Pat Felton, and Chris Hegarty. Proposed Airport Pseudolite Signal Specification for GPS Precision Approach Local Area Augmentation Systems. In *Proceedings of the Institute of Navigation GPS-97 Conference*, pages 1603–1612, Kansas City MO, September 1997.

- [18] Todd A. Ely, Rodney Anderson, Yoaz E. Bar-Sever, David Bell, Joseph Guinn, Moriba Jah, Pieter Kallemeyn, Erik Levene, Larry Romans, and Sien-Chong Wu. Mars Network Constellation Design Drivers and Strategies. In *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Girwood AK, August 1999. American Astronautical Society.
- [19] Per K. Enge and A.J. Van Dierendonck. Wide area augmentation system. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. II*, pages 117–142. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [20] Jose Guivant, Eduardo Nebot, and Hugh Durrant Whyte. Simultaneous Localization and Map Building Using Natural Features in Outdoor Environments. In *Proceedings of the Intelligent Autonomous Systems 6 Conference*, volume 1, pages 581–588, Italy, July 2000. IOS.
- [21] Albert Van Helden. *The Telescope*. Rice University, <http://es.rice.edu/ES/humsoc/Galileo/Things/telescope.html>, 1995.
- [22] Tom Holden, John Raquet, Gerald Lachapelle, Weigen Qui, C. Pelletier, Anthony Nash, and Patrick Fenton. Development and Testing of a Mobile Pseudolite Concept for Precise Positioning. In *Proceedings of the Institute of Navigation GPS-95 Conference*, pages 817–826, Palm Springs CA, September 1995. Institute of Navigation.
- [23] Jr. James J. Spilker. Fundamentals of signal tracking theory. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. I*, pages 245–327. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [24] Jr. James J. Spilker. Gps signal structure and theoretical performance. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. I*, pages 57–119. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [25] Jack A. Jones. Inflatable Robotics for Planetary Applications. In *Proceedings of the 6th International Symposium on Artificial Intelligence and Robotics & Automation in Space*, Quebec Canada, June 2001. Canadian Space Agency.

- [26] Jaewoo Jung, Per Enge, and Boris Pervan. Optimization of Cascade Integer Resolution with Three Civil GPS Frequencies. In *Proceedings of the Institute of Navigation GPS-2000 Conference*, pages 2191–2200, Salt Lake City UT, September 2000.
- [27] Paul Karol and David Catling. *Planet Mars Chronology*. NASA Ames Research Center, http://humbabe.arc.nasa.gov/mgcm/fun/mars_chro.html.
- [28] Myron Kayton and Walter R. Fried. *Avionics Navigation Systems*. John Wiley & Sons, Inc., New York NY, 1997.
- [29] Changdon Kee. Wide area differential gps. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. II*, pages 81–115. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [30] Changdon Kee, Doohee Yun, Haeyoung Jun, Bradford W. Parkinson, Thomas Langenstein, and Sam Pullen. Precise Calibration of Pseudolite Positions In Indoor Navigation System. In *Proceedings of the Institute of Navigation GPS-99 Conference*, pages 1499–1507, Nashville NT, September 1999.
- [31] Dale Klein and Bradford W. Parkinson. The Use of Pseudo-Satellites for Improved GPS Performance. In *Journal of Navigation*, volume 31, Alexandria VA, Winter 1984. Institute of Navigation.
- [32] Erwin Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, Inc., New York NY, 7 edition, 1993.
- [33] Kenneth Lau, Stephen Lichten, Lawrence Young, and Bruce Haines. An Innovative Deep Space Application of GPS Technology for Formation Flying Spacecraft. In *Proceedings of the AIAA GN&C Conference*, San Diego CA, July 1996.
- [34] Alfred Leick. *GPS Satellite Surveying*. John Wiley & Sons, Inc., New York NY, 2nd edition, 1995.
- [35] Edward A. LeMaster, Masayoshi Matsuoka, and Stephen M. Rock. Field Demonstration of a Mars Navigation System Utilizing GPS Pseudolite Transceivers. In *Proceedings of the 2002 IEEE Position, Location, and Navigation Symposium*, pages 150–155, Palm Springs CA, April 2002.

- [36] Edward A. LeMaster and Stephen M. Rock. Field Test Results for a Self-Calibrating Pseudolite Array. In *Proceedings of the Institute of Navigation GPS-2000 Conference*, pages 1046–1055, Salt Lake City UT, September 2000.
- [37] Edward A. LeMaster and Stephen M. Rock. A Local-Area GPS Pseudolite-Based Navigation System for Mars Rovers. In *Journal of Autonomous Robots*, 2002.
- [38] Edward A. LeMaster and Stephen M. Rock. An Improved Solution Algorithm for Self-Calibrating Pseudolite Arrays. In *Proceedings of the 2002 Institute of Navigation National Technical Meeting*, San Diego CA, January 2002.
- [39] J. Matijevic. Mars Pathfinder Microrover - Implementing a Low cost Planetary Mission Experiment. In *Proceedings of the 2nd IAA International Conference on Low-Cost Missions*, Laurel MD, April 1996. IAA.
- [40] Masayoshi Matsuoka. Stanford University, 2002.
- [41] Masayoshi Matsuoka and Guttorm Opshaug. Stanford University, 2002.
- [42] P. H. Milne. *Underwater Acoustic Positioning Systems*. Gulf Publishing Company, Houston TX, 1983.
- [43] Paul Y. Montgomery. *Carrier Differential GPS as a Sensor for Automatic Control: Development of a Full-State Estimation and Flight Control System for an autonomous Aircraft Based on the Global Positioning System*. PhD thesis, Stanford University, Stanford, CA 94305, December 1996. Also published as SUDAAR 688.
- [44] Awele Ndili. GPS Pseudolite Signal Design. In *Proceedings of the Institute of Navigation GPS-94 Conference*, pages 1375–1382, Salt Lake City UT, September 1994.
- [45] Michael O’Conner, Thomas Bell, Gabriel Elkaim, and Bradford Parkinson. Real-Time CDGPS Initialization for Land Vehicles Using a Single Pseudolite. In *Proceedings of the Institute of Navigation National Technical Meeting*, San Monica CA, January 1997.
- [46] Eric Olsen. Stanford University.
- [47] Eric A. Olsen. *GPS Sensing for Formation Flying Vehicles*. PhD thesis, Stanford University, 1999.

- [48] Eric A. Olsen, Chan-Woo Park, and Jonathan P. How. Carrier-phase Bias Initialization for Formation Flying Vehicles Using Onboard Pseudolites. In *Proceedings of the Institute of Navigation GPS-99 Conference*, pages 459–468, Nashville TN, September 1999.
- [49] Eric A. Olsen, Patrick A. Stadter, and Mark S. Asher. Long-Baseline Differential GPS-based Relative Navigation for Spacecraft with Crosslink Ranging Measurements. In *Proceedings of the Institute of Navigation GPS-2000 Conference*, pages 1612–1621, Salt Lake City UT, September 2000.
- [50] Guttorm R. Opshaug and Per K. Enge. Integrated GPS and UWB Navigation System. In *Proceedings of the 2002 IEEE Conference on UWB Science and Technology*, Baltimore MD, May 2002.
- [51] Chan-Woo Park, Jonathan P. How, and Larry Capots. Sensing Technologies for Formation Flying Spacecraft in LEO using CDGPS and an Inter-Spacecraft Communications System. In *Proceedings of the Institute of Navigation GPS-2000 Conference*, pages 1595–1607, Salt Lake City UT, September 2000.
- [52] Bradford W. Parkinson and Per K. Enge. Differential gps. In Bradford W. Parkinson and James J. Spilker Jr., editors, *Global Positioning System: Theory and Applications, Vol. II*, pages 3–50. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [53] Bradford W. Parkinson and James J. Spilker Jr., editors. *Global Positioning System: Theory and Applications*. American Institute of Aeronautics and Astronautics, Washington DC, 1996.
- [54] S. Purivigraipong, M.J. Unwin, and Y. Hashida. Demonstrating GPS Attitude Determination from UoSAT-12 Flight Data. In *Proceedings of the Institute of Navigation GPS-2000 Conference*, pages 2625–2633, Salt Lake City UT, September 2000. Institute of Navigation.
- [55] Radio Technical Commission for Aeronautics. *GNSS Based Precision Approach Local Area Augmentation System (LAAS) - Signal-In-Space Interference Control Document (ICD)*, January 2000. RTCA/DO-246A.

- [56] Steve M. Rock, Eric W. Frew, Henry L. Jones, Edward A. LeMaster, and Bruce R. Woodley. Combined CDGPS and Vision-Based Control of a Small Autonomous Helicopter. In *Proceedings of the 1998 American Control Conference*, Philadelphia, June 1998.
- [57] P. S. Schenker, E. T. Baumgartner, P. G. Backes, H. Aghazarian, L. I. Dorsky, J. S. Norris, T. L. Huntsberger, Y. Cheng, A. Trebi-Ollennu, M. S. Garret, B. A. Kennedy, A. J. Ganino, R. E. Arvidson, and S. W. Squyres. FIDO: A Field Integrated Design & Operations Rover for Mars Surface Exploration. In *Proceedings of the 6th International Symposium on Artificial Intelligence and Robotics & Automation in Space*, Quebec Canada, June 2001. Canadian Space Agency.
- [58] GEC Plessey Semiconductors. Gps builder-2 designer's guide. Swindon/Wiltshire UK, 1995.
- [59] Jonathan M. Stone, Edward A. LeMaster, J. David Powell, and Stephen M. Rock. GPS Pseudolite Transceivers and their Applications. In *Proceedings of the Institute of Navigation National Technical Meeting*, pages 415–424, San Diego CA, January 1999.
- [60] Edward Harrison Teague. *Flexible Structure Estimation and Control Using the Global Positioning System*. PhD thesis, Stanford University, 1997.
- [61] Jr. Thomas A. Stansell. RTCM CS-104 Recommended Pseudolite Signal Specification. In *Global Positioning System*, volume III, Washington DC, 1986. Institute of Navigation.
- [62] James Bao-Yen Tsui. *Fundamentals of Global Positioning System Receivers: A Software Approach*. John Wiley & Sons, Inc., New York NY, 2000.
- [63] Phillip Ward. Gps satellite signal characteristics. In Elliot D. Kaplan, editor, *Understanding GPS Principles and Applications*, pages 83–117. Artech House, Norwood MA, 1996.
- [64] Phillip Ward. Satellite signal acquisition and tracking. In Elliot D. Kaplan, editor, *Understanding GPS Principles and Applications*, pages 119–208. Artech House, Norwood MA, 1996.

- [65] Martin Weiser. Development of a Carrier and C/A-Code Based Pseudolite System. In *Proceedings of the Institute of Navigation GPS-98 Conference*, pages 1465–1475, Nashville TN, September 1998.
- [66] Lawrence L. Wells. The New Translated GPS Range System. In *Proceedings of the Institute of Navigation GPS-97 Conference*, pages 1705–1709, Kansas City MO, September 1997.
- [67] Bruce R. Woodley, Henry L. Jones, Edward A. LeMaster, Eric W. Frew, and Stephen M. Rock. Carrier Phase GPS and Computer Vision for the Control of an Autonomous Helicopter. In *Proceedings of the Institute of Navigation GPS-96 Conference*, pages 461–465, Kansas City MO, September 1996.
- [68] Kurt R. Zimmerman. *Experiments in the Use of the Global Positioning System for Space Vehicle Rendezvous*. PhD thesis, Stanford University, Stanford, CA 94305, December 1996. Also published as SUDAAR 692.