

**Six Degree of Freedom Navigation Using X-Ray Pulsar
Signals**

**A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY**

Joel Thomas Runnels

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY**

Advisor: Demoz Gebre-Egziabher

August, 2019

© Joel Thomas Runnels 2019
ALL RIGHTS RESERVED

Acknowledgements

Isaac Newton once famously said that “If I have seen further than others, it is by standing upon the shoulders of giants.” I certainly cannot claim to have seen further than others, but for what I have accomplished, I am most certainly indebted to countless “giants” in my life who have supported me throughout this work. This section is far too brief to fully acknowledge everyone who has helped me in this way, but I hope to at least mention some of the individuals who have played key roles in my success.

I would like to begin by acknowledging organizations which have contributed to this work by providing funding and/or data which has helped to make this research possible. This research has made use of data obtained through the High Energy Astrophysics Science Archive Research Center Online Service, provided by the NASA/Goddard Space Flight Center. The scientific results reported in this article are based in part on data obtained from the Chandra Data Archive. This research has made use of data obtained from the Chandra Source Catalog, provided by the Chandra X-ray Center (CXC) as part of the Chandra Data Archive. I also acknowledge the NASA/Minnesota Space Grant Consortium, the Air Force Research Lab University Nanosat Program, the University of Minnesota Office of the Vice President for Research (OVPR), and Minnesota’s Discovery, Research, and InnoVation Economy (MnDRIVE) for providing funding to support this work. While I gratefully acknowledge the aforementioned individuals and organizations, the views and conclusions expressed in this dissertation are those of the author alone and should not be interpreted as necessarily representing the official policies, either expressed or implied, of any organization.

I would like next to thank my advisor, Dr. Demoz Gebre-Egziabher. It has truly been a pleasure to study under someone with his expertise and experience. I have learned so much from him over the course of my studies, and could not have asked for a

better advisor. I have truly appreciated his help and guidance, as well as his continued optimism, especially when mine started to fade.

I also would like to gratefully acknowledge several professors at University of Minnesota for their help and support during my doctoral studies. I would particularly like to thank Dr. Emad Ebbini for helping me begin to understand the world of digital signal processing through his excellent instruction and countless office hours, and whose course in adaptive signal processing provided me with the initial inspiration for my dissertation. I would also like to thank Dr. Lindsay Glesener for her help and guidance in my use of high energy astrophysics data. Finally, I would like to thank Dr. Peter Seiler for his continued feedback on my work, particularly during seminar presentations.

I am also grateful to many students both within and outside of the Aerospace Department with whom I have become friends over the years, and who have provided much encouragement, support and camaraderie. I would like to especially thank: Raghu Venkataraman, for making me feel at home when I first visited UMN, Inchara Lakshminarayan, for always having something kind to say about my presentations, Jordan Larson, for the many fun conversations and lots of insights, Jordan Hoyt, for all the fun conversations and trips to McDonald's, Kerry Sun for organizing so many lab socials, Guy Squillace, for the fun times our families have shared (and for your help on many projects), and Rafael Camata for being an amazing friend and study-partner. I have also had the great pleasure of making many friends outside of my studies at UMN, and I have greatly appreciated their support. Particularly, I would like to thank Rosie and Earl Klein for befriending me and my family. I would also like to thank our small-group friends: the Fultzes, the Paulsons, the Holleys, and the Ertsgaards, along with the Beards and the rest of my friends at First Evangelical Free Church in Maplewood. Finally, I would like to especially thank George, the 7:00 AM Campus Connector driver, for your many words of encouragement over the years.

It is no exaggeration to say that without my parents, Scott and Joy Runnels, I would not be where I am today. Their tireless efforts, not only in raising me to be responsible and hard-working, but also to give me an education through twelve excellent years of home-schooling, are without a doubt the single most important contributor to my success both academically and personally. My parents have loved me, disciplined me, and taught me how to be a better person, and I am forever grateful to them for the

countless sacrifices that they have made for me. I would also like to acknowledge my siblings and fellow *Runnels Mountain Academy* alumni, Brandon Runnels, Rebecca Smith, and Catherine Runnels, for their continued support, encouragement, and all the good times we have shared together as a family.

I don't know where to begin to thank my wonderful wife, Amanda, for everything she has done to help me throughout my studies. She has listened to countless iterations of the same presentations, each time offering valuable feedback. She has read my papers and dissertation, and helped me find ways to communicate my ideas more clearly. She has supported me when I needed it most, encouraged me when my research seemed hopeless, and has always been willing to listen to me ramble on about my latest problems. I can't thank you enough for sharing this chapter of my life with me and helping me in so many ways. I also would like to thank our son, Teddy. Thank you for all the smiles and laughter which have brought me so much joy. I have enjoyed getting to spend so much time with you as I've worked on research, and I will miss so much getting to take care of you as a stay-at-home dad.

Finally, and most importantly of all, I thank my Lord and Savior Jesus Christ. It is through Him that I live and am saved. He gives my life meaning, direction, and purpose. He is the author and creator of logic, reason, and the universe which we have the joy of living in and exploring.

*When I look at your heavens, the work of your fingers,
the moon and the stars, which you have set in place,
what is man that you are mindful of him,
and the son of man that you care for him?*

Psalm 8:3-4, ESV

Of making many books there is no end, and much study is wearisome to the flesh.

Ecclesiastes 12:12, NKJV

Dedication

To my family, who have always been there for me

Abstract

Navigation in deep space, far away from Earth, is an ongoing challenge and research topic. While spacecraft near Earth have a number of readily available methods for navigation (including GPS and radio ranging), far away from Earth it is more challenging for spacecraft to determine their position. In the absence of external reference objects that can be used to estimate position (for instance nearby planetary objects), the current state-of-the-art for navigation in space relies on NASA's Deep Space Network to provide Earth-based position measurements of the spacecraft. This means of navigation suffers from limitations, including limited availability, high cost, and decreased accuracy far from Earth. Consequently, alternative means of navigation are of interest.

X-ray navigation, or XNAV is a proposed means by which spacecraft can navigate using signals generated by astrophysical signal sources. In particular, x-ray pulsars have been proposed as a naturally occurring signal source which could be used to generate a position, navigation and timing (PNT) solution in space. The basic concept in XNAV is that a spacecraft can compute a PNT solution based on time of arrival (TOA) measurement of signals from x-ray pulsars. Some x-ray pulsars, in particular millisecond pulsars, have extremely precise timing characteristics, with timing stability comparable to modern atomic clocks. If the TOA of signals from several millisecond pulsars could be measured, these TOAs could be used to compute a PNT solution for the spacecraft. The basic concept of XNAV is somewhat analogous to GPS, in that the position of the user is determined by measuring multiple signal TOAs generated by sources with precisely known timing characteristics.

While this technique has been proposed numerous times in literature, there are still several implementation challenges which must be overcome in order for XNAV to become a viable navigation technology. In this dissertation, we address some of the major challenges associated with implementation of XNAV.

The first challenge addressed in this dissertation is the development of a method of determining the signal time of arrival based on measurements of x-ray photon arrival times. This challenge is at the heart of any XNAV implementation, because in order to use pulsar signals as PNT signals, the time-difference of arrival of the signal must

be measured. The estimation of time-difference of arrival from pulsar signals is complicated by the fact that pulsar signals are incredibly weak, resulting in a signal-to-noise ratio near zero. In this dissertation, we develop a recursive algorithm which estimates the time-difference of arrival of a pulsar signal which is based upon adaptive filtering techniques.

The second challenge addressed in this dissertation is the problem of data association. Photons measured by an x-ray detector in space have no way of knowing with certainty the origin of the photons. The presence of the uniform x-ray background results in background photons diluting an already extremely weak signal. If the detector's attitude is known, then the attitude may be used to determine which photons are likely to have originated from a signal source of interest. However, the reliance upon attitude to correctly associate the photons with the correct signal source causes the position and attitude estimates to be coupled. In this dissertation, we present an algorithm which addresses this coupling of the attitude and PNT solutions for the XNAV problem. A joint six degree-of-freedom position and attitude estimator is developed based on the joint probabilistic data association filter. We further demonstrate the effects of attitude uncertainty on the accuracy of the PNT solution using Monte Carlo simulations.

Contents

Acknowledgements	i
Dedication	iv
Abstract	v
List of Tables	x
List of Figures	xi
1 Introduction	1
1.1 Deep Space Navigation: Current Techniques	3
1.2 Opportunistic Navigation	4
1.3 Prior Work	8
1.4 Problem Statement and Contributions	9
1.5 Dissertation Organization	10
2 Mathematical Models of Astrophysical Signals of Opportunity	12
2.1 Signal Sources	12
2.1.1 Gamma ray bursts	13
2.1.2 Asynchronous x-ray sources	14
2.1.3 Pulsars	16
2.2 Signal properties and statistics	16
2.2.1 Methods of determining time difference of arrival	16
2.2.2 Poisson Processes: A Practical Example	18

2.2.3	Statistical Models	20
2.2.4	Methods For Estimating λ and Correlation	25
3	Recursive Range Estimation	32
3.1	Correlation Vector Kalman Filter	35
3.1.1	Correlation Vector Time Update	36
3.1.2	Alternative implementation of time update	42
3.1.3	Correlation Vector Measurement Update	44
3.1.4	Alternate Formulation of Measurement Update Equations	47
3.2	Estimation of TDOA from Correlation Vector	49
3.2.1	Linearized Transform	50
3.2.2	Unscented Transform	52
3.3	Monte Carlo Simulations	53
3.3.1	Theoretical Accuracy	55
3.3.2	ML Estimator	56
3.3.3	TDOA Kalman Filter	57
3.3.4	Monte Carlo Results	58
3.4	Validation Experiment	62
3.5	Implementation Details	65
3.5.1	Peak Centering	65
3.5.2	Integrated Navigation Filter	67
4	Attitude Estimation	77
4.1	Quaternions and Spatial Rotation	78
4.1.1	Composition of Spatial Rotations	80
4.2	Attitude State Vector	81
4.3	Attitude Measurement Update Equations	82
4.3.1	Unit Vector Matching	82
4.3.2	Angle Measurements	86
4.4	Overview of Attitude Estimation Algorithm	90
5	Data Association and Joint Range and Attitude Estimation	96
5.1	Naive Data Association	96

5.2	Joint Probabilistic Data Association Filter	98
5.2.1	Association Probabilities	99
5.2.2	Joint Measurement Update	106
5.3	Results	107
5.3.1	Monte Carlo Simulation Results	107
5.3.2	<i>Chandra</i> Validation Experiment	112
6	Conclusions and Future Work	121
6.1	Conclusions	121
6.2	Filter Initialization	121
6.2.1	Algorithm Development	123
6.2.2	Monte Carlo Results	131
6.2.3	<i>Chandra</i> Validation Experiment	133
6.3	Derivation of Lower Bounds on Error	137
6.4	Detector Hardware Development	138
6.5	Use of Alternative Signal Sources	138
	References	140
	Appendix A. Alternate Derivation of NHEP	150

List of Tables

3.1	Monte Carlo Simulation Parameters	53
3.2	Pulsar Parameters	54
5.1	Pulsar Parameters	112
5.2	Monte Carlo Input Parameters	112
5.3	HRC-S Parameters (timing mode)	113
5.4	Standard deviations on simulation inputs	115
5.5	Estimate error standard deviations from <i>Chandra</i> validation experiments	116
6.1	Signal Parameters	132
6.2	Monte Carlo Parameters	133

List of Figures

1.1	Ancient navigational techniques included accurate measurement of the position of stars.	2
1.2	Absolute NAVSOP concept	6
1.3	Relative NAVSOP concept	6
2.1	Pulsar signal profiles from four x-ray pulsars (J0437-4715 , B1937+21 , B1957+20 , B1821-24). As shown here, each pulsar has a unique pulse profile, period, photon flux, and unpulsed signal component.	15
2.2	Correlation (top panel) is a measure of how similar two signals are as a function of time delay (bottom panels)	17
3.1	X-ray navigation filter block diagram. The filter outputs are the complete PNT state vector consisting of the attitude quaternion \mathbf{q} , the angular velocity rate $\boldsymbol{\omega}$, the position vector \mathbf{r} , and the velocity vector \mathbf{v}	33
3.2	Correlation function as state vector	35
3.3	Correlation peak shifted by δ	37
3.4	TDOA axis shifted by δ	43
3.5	Monte Carlo Error Analysis, FOV=0.1°	59
3.6	Monte Carlo Error Analysis, FOV=0.5°	60
3.7	<i>Suzaku</i> 's position relative to the origin	62
3.8	Range errors for four different time segments of the <i>Suzaku</i> mission. Observation IDs: ae100023020 ($\sigma = 80km$), ae100023020 ($\sigma = 64km$), ae101003010 ($\sigma = 54km$)	64
3.9	Good and bad quadratic fits	65
3.10	External one-dimensional navigation filter	67

4.1	Rotation from navigation frame N to spacecraft body frame B by axis-angle rotation	79
4.2	AOA and unit vector	83
4.3	Photon angle of arrival measurement as viewed from the detector focal plane. It will be noted that the actual detector coordinates of the photon arrival will be a “mirror image” of the photon’s sky coordinates.	84
4.4	Estimated and true measurements with pitch and yaw error only	87
4.5	Estimated and true measurements with roll, pitch and yaw error	89
5.1	A naive approach to data association	97
5.2	A measurement update in the presence of source uncertainty	98
5.3	Uncertainty in the spacecraft’s position results in uncertainty in the estimated flux at \mathcal{I}	102
5.4	TDOA estimate standard deviations, 100 cm ² detector area	108
5.5	TDOA estimate standard deviations, 1000 cm ² detector area	109
5.6	Attitude estimate standard deviations, 100 cm ² detector area	110
5.7	Attitude estimate standard deviations, 1000 cm ² detector area	111
5.8	In both <i>Chandra</i> observations, at least one static x-ray star was visible in the detector’s field of view in addition to the pulsar being observed. This allowed for a full three dimensional attitude solution to be estimated.	113
5.9	Range estimate errors	119
5.10	Attitude estimate errors	120
6.1	Integer ambiguity. The signal phase ϕ is estimated by algorithms of Chapters 3 and 5. However, N is unobservable in these algorithms.	122
6.2	Acceleration magnitudes near the Earth-moon system. Orbits shown in dotted lines. Notable in this figure is the location between the Earth and the Sun at which the acceleration vector magnitude approaches zero, i.e. where the gravitational force of Earth is equal to and opposite that of the Sun.	124
6.3	Acceleration magnitudes in the vicinity of the inner planets. Orbits shown in dotted lines.	125
6.4	Acceleration magnitudes in the vicinity of the outer planets. Orbits shown in dotted lines.	126

6.5	Map of acceleration required measurement accuracy near the Earth-moon system. Orbits shown in dotted lines.	128
6.6	Map of acceleration required measurement accuracy in the vicinity of the inner planets. Orbits shown in dotted lines.	129
6.7	Map of acceleration required measurement accuracy in the vicinity of the outer planets. Orbits shown in dotted lines.	130
6.8	Standard deviation of acceleration estimate as a function of runtime . .	134
6.9	Acceleration estimate errors for three <i>Chandra</i> observations	136
A.1	Generation of NHPP by thinning	150

Chapter 1

Introduction

Navigation is among the oldest problems faced in engineering. While the Egyptians were building pyramids in Giza, Polynesians were learning how to use stars to navigate between islands and how to determine if they were near land by studying the waves and observing the patterns of the birds. Meanwhile on the island of Crete, the Minoans were building ships that sailed to Egypt, using the stars as navigational aides. In ancient civilizations around the world, people were learning to use natural landmarks as tools to determine their location.

As the centuries went by, civilization became more skilled at the art of navigation. Tools such as the cross-staff (a precursor to the modern sextant) and the compass, along with basic nautical charts began to enable exploration further and further away from land, and eventually allowed for the great explorers to make voyages all the way around the world. A major breakthrough came with the invention of the pocket-watch in 1761. It had long been known that accurate knowledge of time was necessary to accurately determine longitude. The problem was that clocks tend to drift. Before the invention of the pocket-watch the best available technology for time-keeping was the hour glass, which tended to be less accurate aboard ships that pitch and roll on the waves. The pocket-watch solved this problem and allowed sailors to accurately determine their longitude as well as latitude by observing the position of the sun. This development underscores an interesting fact about many navigation problems: In order to know *where* you are, you also must have an idea of what *time* it is.

In the 20th century, vast strides were made in the art of navigation thanks to the

use of electronics and radio. By generating our own navigational signals rather than relying on those which occur naturally, it became possible to determine position to a much higher degree of accuracy. These man-made navigational signals ultimately gave birth to the first Global Navigation Satellite System: GPS.

In modern navigation we tend to rely heavily on man-made “beacons” or reference points to determine our position. We use street signs to determine where we are on a map, we use GPS signals to determine where we are anywhere on Earth, and airplanes use a vast network of radio navigational beacons to determine their positions. We tend to use artificial signals because they’re convenient; we can design the signals to be optimized for navigational purposes. However, before the advent of modern navigational technology, explorers had to rely on natural signals, such as the sun, the stars, and the Earth’s magnetic field. These navigational beacons do not always offer the same ease and convenience as their man-made counterparts, but they do offer a certain universal availability and reliability which man-made signals lack.

The dawn of the space-age required navigation in a new domain. Traditionally, navigation in space has relied on man-made signals, just as most other modern navigation does. However, the vastness of space cannot be overstated, and this poses a problem for man-made signal based navigation. On Earth, it is feasible (if expensive) to broadcast navigational signals over the entire planet using satellites. This simply is not possible in space due to its vastness. The best we can do is to broadcast highly targeted navigational signals directly at the spacecraft that need to navigate, but even this becomes challenging as spacecraft explore further and further




Figure 1.1: Ancient navigational techniques included accurate measurement of the position of stars. [1, 2]

away from Earth. As we continue to explore the outer boundaries of the solar system and beyond, it may well be that our ability to broadcast man-made navigational signals is not able to “keep up” with our pace of exploration. Consequently, we again turn to the ancient means of navigation: Using the signals and beacons which are already there, rather than trying to create our own. This is the basic idea of this work. We are seeking ways to use naturally occurring signals in outer space to determine a position, navigation, and timing (PNT) solution.

1.1 Deep Space Navigation: Current Techniques

Navigation in deep space, or regions beyond geosynchronous orbit (GEO), presents an ongoing research challenge [3]. While spacecraft between low Earth orbit (LEO) and GEO have a wide variety of navigational aides available to them, outside of GEO, navigation is much more challenging. This is because most of the current techniques used for navigation in space rely on some form of Earth-based measurements to fix the position of the spacecraft. This becomes more and more difficult to do as the spacecraft travels farther from Earth. In particular, while it is fairly simple to measure a spacecraft’s distance from Earth in the radial direction, measuring its crosstrack (tangential directions) is more difficult because it relies on accurate measurements of the spacecraft’s angular position from Earth. The effect of errors in this measurement increases linearly with distance between Earth and the spacecraft. For instance, to determine the position of a spacecraft near Pluto, the angular position of the spacecraft would have to be measured to an accuracy of 1 nanoradian in order to achieve an accuracy of 1 kilometer in the calculated crosstrack.

NASA’s Deep Space Network (DSN) is capable of measuring the angular position of spacecraft to approximately 0.33 nrad [4]. However, relying on the DSN to measure position of spacecraft in deep space has multiple limitations. The fundamental limitation is that even with extremely high accuracy angular measurements, the error in position measurement will always increase linearly with distance away from Earth. Consequently, for a given angular accuracy, the resulting errors in position may become unacceptably large at great enough distances from Earth. In addition to this fundamental limitation, navigation via the DSN is limited practically by time availability.

The number of spacecraft that rely on the DSN for navigation has steadily increased over the last several decades, in part due to new planned missions, and in part due to old missions outliving their life expectancy (e.g. Voyager I and II). As more and more spacecraft rely on the DSN for positioning, it will become increasingly difficult to schedule navigational updates. In order to measure a spacecraft’s angular position to the accuracy of 0.33 nrad, the DSN uses Very Long Baseline Interferometry, which requires approximately 1.25 hours of observation time to perform the measurement [4]. This long observation time makes navigation via the DSN difficult to schedule, and increasingly impractical.

A recent mission which exemplifies the current state-of-the-art capability of spacecraft for navigation in deep space is the *New Horizons* mission and the recent Pluto and Ultima-Thule flybys. During the spacecraft’s approach to the systems, the spacecraft’s position error was determined by tracking data from the DSN. In order for a successful acquisition of the targets to occur, the spacecraft’s ephemeris error had to be on or below the order of several thousand kilometers. The tracking data provided by the DSN was accurate to about one thousand kilometers [5].

While the DSN produces accurate navigational solutions for missions such as New Horizons, alternative navigation techniques which reduce or eliminate dependency upon the DSN would enable more accurate navigation very far from Earth, and more readily available navigation solutions for spacecraft. Autonomous (self-contained) navigation systems which rely on naturally occurring astrophysical signals have been proposed as an alternative navigation technique which could augment or even replace navigation via the DSN. This technique is often referred to as “signal of opportunity” or “opportunistic” navigation.

1.2 Opportunistic Navigation

Navigation via Signal of Opportunity (NAVSOP) generally refers to a navigation technique which uses signals native to the operating environment to estimate a position, navigation and timing (PNT) solution. A wide variety of signals have been proposed for use with NAVSOP, including signals from WiFi [6], cell-phones [7], pulsars [8, 9], and γ -ray bursts [10], to name a few. The primary advantage of NAVSOP techniques is that

they do not rely on an external infrastructure designed for navigation measurements, but rather opportunistically use signals which are already present in the environment. The advantage of being able to operate independently of an external infrastructure is especially apparent in the case of deep-space navigation.

The fundamental idea behind NAVSOP is to use signals native to the environment to estimate the position of the user. This can be achieved by measuring the time-difference of arrival, or TDOA of the signal, at two separate points in space (the user and the origin of the navigation coordinate frame). Mathematically, TDOA is defined as follows:

$$\tau^{O \rightarrow A(S)} = t_A - t_O \quad (1.1)$$

where t_A is the time-of-arrival of a specific feature of the signal at the spacecraft A , and t_O is the time-of-arrival of that same feature of the signal at the origin.

NAVSOP techniques can be broadly broken down into two classes; absolute navigation and relative navigation. Absolute NAVSOP produces an absolute navigation solution for the user, and requires that the signal is known *a priori*. In this case, the signal values at the origin may be stored onboard the spacecraft before departure. Then, no matter how far away the spacecraft is from the origin, it will always have the value of the signal at the origin available for computing the TDOA. This scenario is illustrated in Figure 1.2.

In the case where the signal is not known *a priori*, as is the case with quasi-periodic signals, absolute NAVSOP may still be possible. Here, the limiting factor would be the bandwidth required to transmit a time history of the signal from the origin to the spacecraft, especially if the spacecraft is far away from the origin.

However, even if the signal cannot be known *a priori*, it may still be useful for relative NAVSOP, or computing the relative position between two vehicles that are relatively close together. The bandwidth limitation is only a real problem if the distance between the origin and the spacecraft is very large. If multiple spacecraft are close to each other, then the signals could easily be passed from user to user, allowing the cooperating spacecraft to compute their positions relative to each other. As long as one




Figure 1.2: Absolute NAVSOP concept




Figure 1.3: Relative NAVSOP concept

of the spacecraft knows its absolute position, the other cooperating spacecraft could use that absolute position combined with the relative navigation solutions to compute their respective positions. One such scenario is illustrated in Figure 1.3.

It is apparent from examination of Figures 1.2 and 1.3 that even though the signals and applications of the two methods may be different, both depend on being able to generate a measurement of TDOA. If, without loss of generality, one considers the absolute NAVSOP case only, then the key equation relating TDOA to the position vector is given by:

$$\mathbf{y}(\mathbf{x}) = \begin{bmatrix} \tau_O^{A(1)} \\ \vdots \\ \tau_O^{A(N)} \end{bmatrix} = \begin{bmatrix} \frac{1}{c} (\boldsymbol{\ell}^{(1)})^T & 1 \\ \vdots & \\ \frac{1}{c} (\boldsymbol{\ell}^{(N)})^T & 1 \end{bmatrix} \begin{bmatrix} \mathbf{r}_O^A \\ \delta t_A \end{bmatrix} \quad (1.2)$$

where $\mathbf{y}(\mathbf{x})$ is a vector of measured values, related to the PNT state vector \mathbf{x} . The variable $\tau_O^{A(1)}$ is the time-difference of arrival of the signal generated by signal source 1 between the spacecraft A and the navigation frame origin O ; c is the speed of light; $\boldsymbol{\ell}^{(1)}$ is the line-of-sight unit vector to source 1; \mathbf{r}_O^A is the position vector to the spacecraft and δt_A is the spacecraft's clock error. Given four TDOA measurements of signals arriving from different directions, the position \mathbf{r}_O^A and clock offset δt_A of the user can be fixed.

¹

At first glance, this problem bears many similarities to the differential GNSS multilateration problem. However, by the time they reach our solar system, most astrophysical signals are extremely weak. The signals are so weak, in fact, that instead of a continuous signal corrupted by noise, the signal is composed of individual, discrete photons of electromagnetic energy.

The extremely weak signal characteristics present multiple challenges which must be overcome in order for this method of navigation to be functional. Three of these challenges are addressed in this dissertation: Estimation of TDOA from very weak signal sources (Chapter 3), estimation of the spacecraft's attitude from photon arrivals (Chapter 4) and the problem of data association, or correctly determining the origins of measured photons (Chapter 5).

¹A complete table of notation used in each chapter is given at the end of the chapter.

1.3 Prior Work

In [11], the idea of navigation using pulsars as navigational beacons is proposed. Since then, numerous papers examining the technical issues associated with navigation via pulsars have been written [9, 8, 12, 13, 14, 15, 16, 17]. As discussed above, it is of fundamental importance to be able to estimate the TDOA of the signal being used. If the signal is strong enough, then the TDOA may be estimated with reasonable accuracy simply by taking the cross-correlation between the signal observed at the spacecraft, and the signal at the navigation frame origin, and finding the time offset at which the cross-correlation is maximized. However, with weak signals, this method is not viable. Much of the work done on the subject of pulsar navigation is specifically focused on the task of estimating the TDOA from a very weak signal.

Epoch folding refers to the technique used by astronomers to build the pulse profile for a pulsar. Because the signals produced by pulsars are very weak, for any one period of observation of the pulsar, the observed signal will not resemble the actual pulse profile, (as discussed further in Section 2.2). In order to build an accurate pulse profile of the pulsar, many observations of the pulsar are added on top of each other so that eventually, the true pulse profile of the pulsar is obtained. In order to perform this operation, two things must be known *a priori*: The trajectory of the observer during the measurement (so that the effect of the motion may be removed from the photon arrival time), and the period of the pulsar being observed (so that the measurements are folded over using the right time period).

In [15] this technique of “folding” the signal on top of itself is used to determine the pulse TOA, or more specifically the pulse phase difference between the user and another location for PNT purposes. This method suffers from two primary limitations. First, it requires that the trajectory of the user be known before the measurements are made. This means that the user must have at least some idea of their initial velocity in order to effectively use epoch folding. While this is not a completely unreasonable requirement, it does preclude its use in a “cold-start” situation (i.e., when a spacecraft is navigating with no prior trajectory knowledge). Secondly, this method requires precise knowledge of the frequency of the pulsar being observed. This constraint makes the use of epoch folding on anything other than a pulsar (e.g. an asynchronous x-ray source) challenging.

Another approach for estimating pulse phase applies a maximum likelihood estimator (MLE) to estimate the phase offset and Doppler shift of the pulsar between the user and some other location [14, 13]. The pulsar navigation demonstration mission SEXTANT used such an MLE approach for determining pulse phase [17, 18, 19]. This method offers the distinct advantage over epoch folding in that it does not require that the trajectory of the spacecraft be known *a priori*. However, similar to the epoch folding method, MLE estimation of phase difference requires that the pulse profile be known to a high degree of accuracy beforehand. Another possible disadvantage of MLE estimation is that, as pointed out in [14], the complexity of calculations increases significantly as the observation time grows.

Both MLE estimation and epoch folding make use of the kinematic state of the vehicle at the start of the observation time to produce their estimates of TDOA. Additionally, some implementations of the MLE approach, such as [17] also estimate a frequency (Doppler) shift of the signal along with TDOA, which yields an estimate of both position and velocity from the observations. Further derivatives of position could be added to an MLE estimation scheme, but at the cost of increased computational expense.

In [12], adaptive filters are proposed as a means of estimating the TDOA of a pulsar signal. The authors of [12] show that adaptive filters such as LMS, nLMS, and RLS may be used to compute the correlation vector between a measured and desired signal.

1.4 Problem Statement and Contributions

In order for x-ray navigation to function as a viable alternative to navigation via traditional means, there are several challenges which need to be overcome. The two challenges that are addressed in this work are outlined below.

The first problem addressed by this work is the problem of time-difference of arrival estimation using extremely weak signals. We present a recursive algorithm which estimates the TDOA for very weak signals, in particular signals which consist of individual photon arrivals. We show that this algorithm is capable of estimating TDOA with mean error approaching theoretical limits on accuracy derived in previous work.

The second problem addressed in this work is the problem of data association. The

ability of any TDOA estimator to function is dependent upon its ability to correctly associate incoming photons with the correct signal sources. We show that the quality of the time-difference of arrival estimate is inherently coupled to the quality of the spacecraft's attitude solution. We then derive a technique through which attitude may be estimated based on photon measurements. Finally, an estimation technique is presented which uses a probabilistic data association technique to address the problem of data association, and allows the joint estimation of attitude and range.

1.5 Dissertation Organization

The remainder of this dissertation is organized as follows. In Chapter 2, various astrophysical signals of opportunity which could be used for PNT purposes are presented, and the statistical properties of these signals are derived. Chapter 3 presents a recursive range estimation algorithm which allows the time-difference of arrival, and consequently range to be estimated from extremely weak signals. The problem of attitude determination is discussed in Chapter 4, and the equations necessary to estimate a spacecraft's attitude from photon angle-of-arrival measurements are derived. The algorithms from Chapter 3 and Chapter 4 are combined in Chapter 5 to form a joint range and attitude estimator, and it is shown that the problem of position estimation and attitude estimation are in fact coupled. Finally, in Chapter 6, some of the remaining challenges which are not addressed in this dissertation are discussed, and some results are presented from preliminary efforts at addressing some of these challenges.

Nomenclature

A table of nomenclature is included here as well as at the end of each chapter.

Data Association

N	Number of signals being tracked
-----	---------------------------------

Physical Constants

c	Speed of light (km/s)
-----	---------------------------

Signals and Measurements

t_i Time-of-arrival measurement of photon i (*seconds*)

S^i i^{th} signal source

State Estimation

$\rho_t^{A(S)}$ Range to object A, at time t, along line-of-sight vector to source S (*km*)

\mathbf{r}_t^A Position of A, at time t (*km*)

O Denotes the navigation frame origin

$\tau_t^{A(S)}$ Time-difference of arrival of signal from signal source S measured between vehicle A and navigation frame origin (*seconds*)

\mathbf{x} State vector

$\mathbf{y}_{t,a}$ Measurement of scalar quantity a

$[\cdot]^-$ Indicates *a priori* value

$[\cdot]^+$ Indicates *a posteriori* value

$\hat{\cdot}$ Indicates estimated value

\mathbf{a} Indicates vector quantity

a Indicates scalar quantity

Misc

A, B Indices to designate spacecraft(s)

$\ell^{(S^i)}$ Unit vector pointing towards i^{th} signal source

Chapter 2

Mathematical Models of Astrophysical Signals of Opportunity

At the heart of any opportunistic navigation scheme is a naturally occurring signal which can be used for navigational purposes. In this chapter, the types of signals which can be used for signal of opportunity navigation in space are discussed. Then, we will develop mathematical models for the one signal on which we will focus for the remainder of this dissertation, namely x-ray pulsars. These mathematical models are important because they will enable us to develop PNT algorithms that use these signals in later chapters of this dissertation.

2.1 Signal Sources

The first question that must be asked if we wish to formulate a signal of opportunity navigation scheme is: What signals are available? It turns out that naturally occurring signals are relatively plentiful in outer space, though not all of them are useful.

This leads to the next question: What sorts of signals are most useful for navigation? There are two attributes that make a signal particularly useful for navigation. One is that the signal should contain some form of spatial variation, which may be used to

derive information about the position or orientation of the spacecraft. The higher the information content of the signal, the more useful it is for navigational purposes. The other is that the signal must be strong. The signal might have an abundance of information, but if it is too weak to be detected by the spacecraft, then it is of little use for navigational purposes.

In this work, we have restricted our attention to electromagnetic signals of wavelengths significantly shorter than visible light, particularly hard x-rays and soft γ rays, containing photons ranging in energy from 2-100 kiloelectron-volts (keV). While electromagnetic signals of all wavelengths exist in space, the higher frequency/lower wavelength signals are of particular interest because they can generally be detected with smaller detectors than what would be required to detect longer wavelength signals, such as radio signals. Additionally, the hard x-ray/soft γ ray window represents a sort of “sweet spot” in the electromagnetic spectrum. On the low end of this range, the cosmic x-ray background becomes more intense, and tends to drown out signals of interest. On the high end of this range, each individual photon is so energetic that even a relatively high energy signal consists of relatively few photons. As will be shown in Section 2.2, the number of photons measured from a signal directly impacts the signal-to-noise ratio (SNR), and fewer photons leads to a lower SNR, even if the individual photons are very energetic.

Given this energy range, there are a few classes of signals which are available for navigational use. Some of these signals are discussed below.

2.1.1 Gamma ray bursts

Gamma ray bursts (GRBs) are the most energetic events that have been observed by man. It is estimated that a GRB releases more energy in a few seconds than our sun is expected to release during its entire lifetime [20]. GRBs are thought to be caused during supernovae, when massive stars collapse to form neutron stars or black holes. These extremely energetic events produce a large amount of radiation, including gamma rays and x-rays.

From Earth, a GRB appears as a bright spot in the the gamma-ray sky which appears and then fades over time during a period referred to as the “afterglow.” The duration of GRBs is highly variable, ranging from less than one second to over 100 seconds [21].

Additionally, it has been shown that the temporal profiles of GRBs are highly variable, with no two GRBs that are alike [22].

GRBs are of interest particularly because they are so energetic. This makes them easier to detect amid the sources of background noise. However, GRBs would be difficult to use for navigation because of the nature of the signal. While the time-varying nature of the signal contains information which could be used for navigation in theory, this is difficult in practice because it is impossible to predict the time, location, or profile of a GRB before it occurs. An analogous situation would be trying to use a signal from a GPS satellite without knowing where the GPS satellite is, or what the GPS signal time-history is supposed to be.

One way GRBs could theoretically be used for navigation is in a cooperative, relative navigation framework [10]. In such a framework, multiple spacecraft would work together and share information about incoming GRB signals, allowing the spacecraft to compute their positions relative to each other.

2.1.2 Asynchronous x-ray sources

Asynchronous x-ray sources are another class of signals which could be used for navigation purposes. In this class of signals, we include any point source of x-rays in which the time variation of the x-ray flux is either constant or asynchronous, i.e. the changes in flux do not exhibit a strong periodic nature. This definition encompasses a very wide range of sources, including most main-sequence stars [23], x-ray binaries [24], x-ray bursters [25], and quasars [26].

This class of signals has the advantage over GRBs in that they originate from a specific location which may be known *a priori*. Consequently a spacecraft can know ahead of time where to look for these signals. However, like GRBs, the temporal variation of these signals cannot be known in advance. Consequently, to be used for positioning applications, these signals would also have to be used in a cooperative, relative navigation framework.




Figure 2.1: Pulsar signal profiles from four x-ray pulsars (J0437-4715 [27], B1937+21 [28], B1957+20 [28], B1821-24 [29]). As shown here, each pulsar has a unique pulse profile, period, photon flux, and unpulsed signal component.

2.1.3 Pulsars

Pulsars, short for pulsating radio star, are highly magnetized spinning neutron stars [30]. The rotation of the pulsar and its magnetic field cause it to emit a beam of electromagnetic radiation which sweeps across the sky, much like the beam from a lighthouse. If this beam of radiation happens to hit our solar system, the result is a periodic signal which may be observed in several spectra, including radio, visual, x-ray and gamma-ray.

Several pulsar pulse profiles are illustrated in Figure 2.1. Some millisecond pulsars have a period stability comparable to that of modern atomic clocks [31]. The precise period of pulsars make them ideal candidates for NAVSOP because a spacecraft could have *a priori* knowledge of the signal it is trying to use. However, pulsar signals have the disadvantage of being relatively weak, resulting in a relatively low signal to noise ratio (SNR) [16, 8], as will be shown below. In this dissertation, we exclusively consider pulsar signals, although many of the techniques developed here could be extended to other signal sources. Thus, the remainder of this chapter is dedicated to developing mathematical and statistical models for pulsar signals.

2.2 Signal properties and statistics

Regardless of the signal of opportunity used to determine position, the end goal is the same. The spacecraft must have a way to measure or estimate the difference between the signal's arrival time at the spacecraft, and the signal's arrival time at the navigation frame origin. This quantity is referred to as the time-difference of arrival, or TDOA. In order to understand why this quantity is difficult to measure, it is beneficial to discuss the nature of the signal.

2.2.1 Methods of determining time difference of arrival

The simplest way to measure time-difference of arrival might be to measure the time at which some feature of the signal occurs (a sharp peak or local maxima in the signal, for instance). By comparing the time at which those features were measured at the spacecraft and at the navigation frame origin, the TDOA could be simply measured by

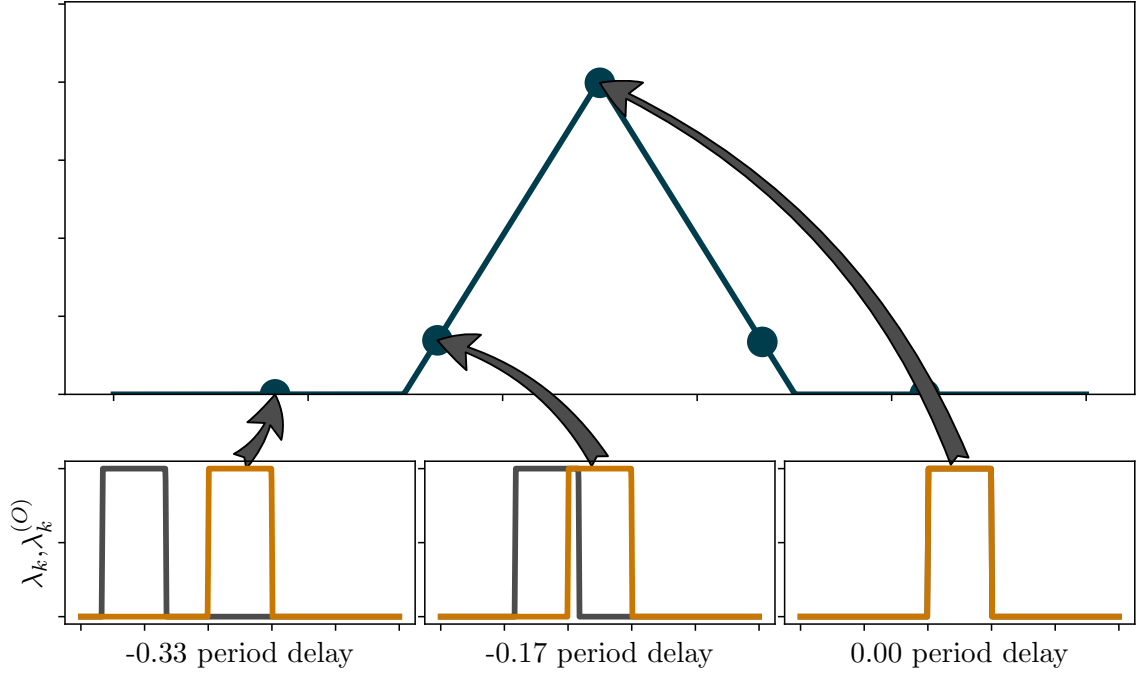


Figure 2.2: Correlation (top panel) is a measure of how similar two signals are as a function of time delay (bottom panels)

taking the difference between the time at which these features were observed. However, this rather simplistic approach is unlikely to yield the best estimate of TDOA, because it generally will not make use of all of the available information in the signal. Additionally, any error or noise in the measurement of the time at which a feature is measured will lead to error in the estimate of TDOA. If only a few features are measured, this error could be quite large depending on the error in the measurements.

A more intelligent method of estimating the time-difference of arrival is to use the cross-correlation method. The correlation function between two time-varying signals, $\lambda_A(t)$ and $\lambda_O(t)$ is defined as follows.

$$c_{\lambda_A, \lambda_O}[p] = \sum_i \lambda_A(t + iT) \lambda_O(t + iT + p) \quad (2.1)$$

If we consider the correlation between a signal, λ_O and a delayed version of that

signal λ_A , the value of the delayed version of the signal may be written as:

$$\lambda_A(t) = \lambda_O(t + \tau) \quad (2.2)$$

In the most basic terms, the correlation is a measure of the “similarity” between two signals. When written as in Equation (3.1), the correlation function measures the similarity between two signals as a function of the time delay of one of the signals. A basic pictorial illustration of the correlation function is shown in Figure 2.2.

It may be shown that the correlation function between a signal, λ_O and a time-delayed version of that signal, λ_A as defined in Equation (2.2), is maximized at the value of the delay between the two signals [32, 33]. This is illustrated in Figure 2.2: The correlation function is maximized at the delay value which causes the two signals to perfectly overlap.

$$\operatorname{argmax} [c_{\lambda_A, \lambda_O} [p]] = \tau \quad (2.3)$$

Therefore, if we could compute the cross-correlation between the photon flux at the spacecraft and at the origin (or, in the case of pulsar signals, we may know the flux rate at the origin ahead of time), then we could easily estimate the TDOA between the two signals simply by finding the location of the maxima of the correlation function.

Unfortunately, we cannot measure the flux rate perfectly. In fact, we cannot even measure the flux rate at all, at least not in a direct sense. What we can do is measure the number of photons that hit a detector over a given period of time. As would be expected, there is a direct relationship between flux rate (power per unit area) and number of photons that hit the unit area. In the most ideal case, we might say that the number of photons we measure during a given instant in time is directly proportional to the flux rate. However, this assumption breaks down under certain conditions.

2.2.2 Poisson Processes: A Practical Example

To understand why we cannot use this direct linear relationship between number of photons and flux rate, an analogy may be useful. Imagine that it’s raining, and we’re

trying to measure how much water is hitting the ground per unit area. Imagine further that all the drops are the same size. Now, suppose that we go outside, and measure an area on the ground. We then try to count the number of drops that hit that area over a given period of time, say one minute. If the area is big enough, and we count the drops accurately, we might be able to get a decent estimate of the rate at which water is falling from the sky (in gallons per square mile, or whatever ratio of volume to area units we wanted). For instance, if we measured out a 1 square foot area, and it was raining lightly, we might count 30 drops hitting that area over the length of one minute. We could do the appropriate calculations and make assumptions about drop size to get a flux rate in volume of water per unit area per second.

But now, say we want to track how fast the rain is falling as a function of time. Moreover, we want our time resolution to be one second. How would we do that? We could count the number of rain drops we see per second that hit our square foot of ground. We were only seeing 30 drops per minute, which means that on average, we saw 0.5 drops per second. But, we cannot actually see half drops. This means that during some one-second periods, we might see one drop, while during other one-second periods we might see none. Some seconds we might see two or three drops, and at other times we might go for several seconds without seeing a drop.

If we were to do the same process that we did before (dividing drops counted by time) for each one second time period, we would get a rain-fall rate that varies drastically with time. At some points we would estimate that our rain-fall rate was torrential, while at other seconds we would estimate that it was not raining at all. Clearly, it would be incorrect to estimate the rain fall rates for a given second based solely on the number of rain drops we count in that particular second. But this raises the question: If we cannot estimate the rain-fall rate over a short period of time simply by counting the number of raindrops, how can we estimate it? We could take the average of all our rainfall estimates over time, but then we are right back to what we did at the beginning. The average will definitely be more accurate, but then we lose all the temporal variation that we were trying to get by measuring more frequently.

The problem of estimating photon flux rate is essentially the same problem. In estimating the flux rate, we are able to count photons over a length of time. We want the length of time over which we were counting to be very small, so that we capture every

variation with high temporal accuracy. (We need high temporal accuracy so that we can accurately estimate the distance between the two spacecraft.) But, as we make our time windows smaller and smaller, we see fewer and fewer photons per time window.

In many cases, the average number of photons we will see per time bin is well below one, meaning that at many times, we will see no photons. And yet, we know that the flux rate at that point in time is not zero. How do we estimate the flux rate from this data with as high of temporal resolution as possible, without the drastic variations associated with small time bins?

2.2.3 Statistical Models

In order to answer this question, we must first have a better understanding of the statistics of the signal we actually measure; that is, the statistics of photon arrival numbers and arrival times.

Photon Count Statistics

Photon arrival times are governed by a Poisson process. More specifically, photon arrivals are a non-homogeneous Poisson process, because the flux varies as a function of time. Poisson distributions model the number of times an event is expected to occur in a given window, given a mean arrival rate. It is the limiting case of a binomial distribution, as the number of trials goes to infinity. The binomial distribution itself represents a sum of Bernoulli trials (i.e. coin flips). Thus, we can think of the Poisson process as modeling each infinitesimally short instant in time as a Bernoulli trial. That is, at each instant in time, we flip a weighted coin. The result of that coin flip determines whether or not we see a photon in that infinitesimal segment of time. The coin is weighted such that, over the sum of Bernoulli trials, the expected value of number of heads is equal to the expected photon arrival rate.

The Poisson distribution is expressed as:

$$\Pr \left[\tilde{\Lambda}_k = N \mid \lambda_k \right] = \begin{cases} e^{-\lambda_k} \frac{(\lambda_k)^N}{N!} & \{N \in \mathbb{Z} \mid N \geq 0\} \\ 0 & N < 0 \end{cases} \quad (2.4)$$

where

- $\tilde{\Lambda}_k = N$ is the number of events observed in a time window
- λ_k is the number of events expected to occur in that time window k

We note that N is restricted to positive integer values, which makes sense based on our understanding of what the Poisson distribution represents (we can only see an integer number of photons, and we can never see fewer than zero). We also note, however, that while λ_k must be positive, there is no restriction that it be an integer. Again, this makes sense based on our understanding of what the distribution represents: we may expect to see a non-integer number of photons over the given time interval, on average.

We will use the shorthand $\tilde{\Lambda}_k \sim \text{Pois}[\lambda_k]$ to indicate that the random variable $\tilde{\Lambda}_k$ follows a Poisson distribution. Below are some relevant properties of the Poisson distribution:

$$\mathbb{E}[\tilde{\Lambda}_k] = \lambda_k \tag{2.5a}$$

$$\text{var}(\tilde{\Lambda}_k) = \lambda_k \tag{2.5b}$$

Another relevant property of the Poisson distribution is the fact that it *does not* depend on what has happened at previous time steps.¹ This means that the probability of seeing a given number of photons at a given time step, given a certain flux, is independent of how many photons were seen at the previous time-step.

As noted above, the arrival of photons from a time-varying signal source is a non-homogeneous Poisson process, that is, it varies with time. We may still use the Poisson distribution to model the photon arrival statistics over a given time interval, but we note that the expected number of photons λ_k is actually the integral of the signal over the time window. Mathematically:

¹It will be noted that the Poisson distribution is not memoryless in the truest statistical sense. However it is “memoryless” in the sense that the number of events observed over a given interval is independent of the number of events observed over the previous interval

$$\lambda_k = \int_{kT}^{(k+1)T} \lambda(t) dt \quad (2.6)$$

where $\lambda(t)$ is the photon flux rate (in units of *photons/second*). Depending upon the length of time-step k , it is frequently sufficient to make the following approximation.

$$\int_{kT}^{(k+1)T} \lambda(t) dt \approx \lambda(t) T \quad (2.7)$$

Based on the statistics of the Poisson process, we may mathematically begin to understand the “rain-drop” problem. As we noted above, if we want to measure the rate of rainfall as a function of time (with high time resolution), the error or noise in our measurements will become very large. As we decrease the size of the time window, T , the noise increases. The properties of the Poisson distribution may be used to express this problem mathematically by examining the signal-to-noise ratio (SNR).

The SNR of a Poisson process may be written as follows.

$$\begin{aligned} \text{SNR} [\tilde{\Lambda}_k] &= \frac{\left(\mathbb{E} [\tilde{\Lambda}_k] \right)^2}{\text{var} (\tilde{\Lambda}_k)} \\ &= \frac{\lambda_k^2}{\lambda_k} \\ &= \lambda_k \\ &= \lambda(t) T \end{aligned} \quad (2.8)$$

We see from the last line in Equation (2.8) that as the time resolution increases (i.e. as T gets smaller), the SNR goes to zero. It is important to note that this loss of signal quality occurs even in what would otherwise be considered an entirely “noise-free” environment. That is, even if there are no sources of additive noise, and even if the detector in question is “perfect”, the SNR still tends to zero as time resolution increases.

Photon Arrival Times

While the *number* of photon arrivals over a given time period is modeled as a Poisson process, the *time*, t , at which the photons arrive is modeled by an exponential process. We show this fact first for the case of a homogeneous Poisson process.

We would like to compute the probability of a photon arriving at or before a given time. We first note that the probability of a photon arriving at or before a given time is equal to one minus the probability of a photon arriving *after* the given time. Mathematically:

$$\mathbf{Pr} [t \leq t|\lambda] = 1 - \mathbf{Pr} [t > t|\lambda] \quad (2.9)$$

We also note that stating that a photon has arrived after a given time is equivalent to stating that zero photons arrived during the given time interval. Recalling Equation (2.4), we can write this probability as:

$$\mathbf{Pr} [\tilde{\Lambda}_k = 0 | \lambda_k] = e^{-\lambda_k} \frac{(\lambda_k)^0}{0!} = e^{-\lambda_k} \quad (2.10)$$

Consequently, we may say:

$$\mathbf{Pr} [t \leq t|\lambda] = 1 - e^{-\lambda_k} \quad (2.11)$$

We then recall that the term λ_k refers to the expected number of photons in a given time period. For the case of a constant flux rate, this term is equal to the photon flux rate times the length of the time period, $\lambda_k = \lambda t$. Therefore, the probability of the photon's arrival time being less than or equal to t is:

$$\mathbf{Pr} [t \leq t|\lambda] = 1 - e^{-\lambda t} \quad (2.12)$$

If the time of arrival, t is a random variable, then Equation (2.12) is by definition the cumulative distribution function of the variable. Consequently, we may obtain the

probability density function for \mathcal{t} by taking the derivative of Equation (2.12).

$$f_{\mathcal{t}}(t|\lambda) = \frac{d}{dt} (1 - e^{-\lambda t}) = \lambda e^{-\lambda t} \quad (2.13)$$

The arrival time cannot be negative, so the full probability density function is written as:

$$f_{\mathcal{t}}(t|\lambda) = \begin{cases} \lambda e^{-\lambda t} & t \geq 0 \\ 0 & t < 0 \end{cases} \quad (2.14)$$

where

- \mathcal{t} is the amount of time before the next event occurs
- λ is the photon arrival per unit time, or the arrival rate

Equation (2.14) is the PDF of the exponential distribution. We will use the short hand $\mathcal{t} \sim \exp[\lambda]$ to indicate that \mathcal{t} is exponentially distributed, and below are a few relevant properties of the exponential distribution.

$$\mathbb{E}[\mathcal{t}] = \frac{1}{\lambda} \quad (2.15a)$$

$$\text{var}(\mathcal{t}) = \frac{1}{\lambda^2} \quad (2.15b)$$

As with the Poisson distribution, the exponential distribution is “memoryless.” That is to say that the probability of an event occurring at some future time is independent of when the last event occurred. Mathematically:

$$\mathbf{Pr}[\mathcal{t} > s + t | \mathcal{t} > s] = \mathbf{Pr}[\mathcal{t} > t] \quad (2.16)$$

The preceding derivations assumed a constant photon arrival rate, i.e. $\lambda(t) = \lambda = C$ where C is some constant. We know, however, that the flux will be time-varying.

Consequently, we need to derive an expression for the PDF of photon arrival times in the time-varying case. As before, we start out by writing the probability that the number of photons we see in a given time interval is zero. Substituting Equation (2.6) into Equation (2.12), we may write:

$$\Pr[t \leq t | \lambda(t)] = 1 - e^{-\int_s^t \lambda(t) dt} \frac{\left(\int_s^t \lambda(t) dt\right)^0}{0!} = 1 - e^{-\int_s^t \lambda(t) dt} \quad (2.17)$$

where

- s is the current “starting” time
- t is future arrival time for which the PDF is being computed

Equation (2.17) is by definition the cumulative distribution function for a non-homogeneous exponential process (NHEP). Consequently we may take the derivative to yield the PDF.

$$f_t(t | \lambda(t)) = \frac{d}{dt} \left(1 - e^{-\int_s^t \lambda(t) dt} \right) \quad (2.18a)$$

$$= \lambda(t) e^{-\int_s^t \lambda(t) dt} \quad (2.18b)$$

We note that in the case where $t = s$, the PDF of the NHEP evaluates to be:

$$f_t(t = s | \lambda(t)) = \lambda(s) \quad (2.19a)$$

An alternative derivation of the PDF for the NHEP, along with a method for generating events governed by an NHEP, is presented in Appendix A

2.2.4 Methods For Estimating λ and Correlation

Given the statistics of the signal we expect to encounter, we are interested in finding a way to estimate the TDOA of that signal. We showed that we cannot simply measure the signal with high time resolution and compute the correlation function from that sampling; since SNR goes to zero as our time resolution increases, the “SNR” of the correlation function we compute would also go to zero.

To address this problem, we turn to techniques from digital signal processing for estimating a signal from a noisy, sampled version of that signal. We begin with the solution to the Wiener-Hopf equations, which describe the ideal finite-impulse response (FIR) filter for estimating the true signal from a noisy version of the signal [34]. We start by defining the following terms:

- $\mathbf{\Lambda}_k = [\lambda_k \ \lambda_{k-1} \ \dots \ \lambda_{k-m+1}]^T$ is the vector of true photon flux at the location of spacecraft for the previous m time-steps
- $\mathbf{\Lambda}_{O,k} = [\lambda_{O,k} \ \lambda_{O,k-1} \ \dots \ \lambda_{O,k-m+1}]^T$ is the true photon flux at the navigation frame origin for the previous m time-steps
- $\tilde{\mathbf{\Lambda}} = [\tilde{\Lambda}_k \ \tilde{\Lambda}_{k-1} \ \dots \ \tilde{\Lambda}_{k-m+1}]^T$ is the time-history of the photon arrivals, measuring $\mathbf{\Lambda}$
- $\mathbf{c}_{\tilde{\mathbf{\Lambda}}, \lambda_O}$ is the correlation vector between $\tilde{\mathbf{\Lambda}}$ and $\mathbf{\Lambda}_O$
- $\mathbf{M}_{\tilde{\mathbf{\Lambda}}, \tilde{\mathbf{\Lambda}}}$ is the auto-correlation matrix of the measured signal, $\tilde{\mathbf{\Lambda}}$
- \mathbf{w} is the ideal finite-impulse response (FIR) filter for estimating $\lambda_{O,k}$ from $\tilde{\mathbf{\Lambda}}_k$

We pause to note the units of the quantities presented above. Photon flux as a variable in continuous time (i.e. $\lambda(t)$) has dimensions of $\frac{\text{photons}}{\text{second} \cdot \text{area}}$. Flux in discrete time (i.e. λ_k) is the result of integration over the time-step, as in Equation (2.6) which yields units of $\frac{\text{photons}}{\text{area}}$. Finally, photon counts (i.e. $\tilde{\Lambda}_k$) is simply a count of photons, and consequently has units of *photons*. To achieve compatibility between the units of $\tilde{\Lambda}_k$ and λ_k , there are two equivalent options. The first is to divide $\tilde{\Lambda}_k$ by the area of the detector in question, resulting in units of $\frac{\text{photons}}{\text{area}}$ for both quantities. The other option is to multiply λ_k by the area of the detector, resulting in units of *photons* for both quantities. We choose to multiply flux by the detector area, resulting in both variables having units of *photons*, but the two options are essentially equivalent and equally acceptable.

With these terms defined, the Wiener-Hopf equations give us an expression for an ideal FIR filter that, when applied to the measured signal will give us the desired signal. This solution is given below:

$$\mathbf{w} = \mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}^{-1} \mathbf{c}_{\tilde{\Lambda}, \lambda_O} \quad (2.20)$$

These equations give us an expression for estimating the flux at the origin, λ_O from the photon arrivals measured at the spacecraft. In the context of estimating the TDOA of a pulsar, we are not actually interested in estimating the true signal at the origin, λ_O , since this is assumed known. But, we note that the Wiener-Hopf equations do contain the correlation vector between the measured signal and the true signal, $\mathbf{c}_{\tilde{\Lambda}, \lambda_O}$. As discussed in Section 2.2, if we can compute the correlation function and find its maximum, we can determine the TDOA between those two signals.

The Wiener-Hopf equations do not give us a direct way to estimate $\mathbf{c}_{\tilde{\Lambda}, \lambda_O}$, but they do help us. We note first of all that there are a number of adaptive filtering algorithms, (for instance LMS, nLMS, RLS, and Kalman Filtering) that give us tools for estimating \mathbf{w} . So, if we estimated \mathbf{w} using one of the many algorithms available, and if we knew $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$, then we could compute $\mathbf{c}_{\tilde{\Lambda}, \lambda_O}$ as follows:

$$\mathbf{c}_{\tilde{\Lambda}, \lambda_O} = \mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}} \mathbf{w} \quad (2.21)$$

In general, $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$ is not known beforehand. However, in the case of photon arrival times, we can actually do an excellent job of approximating it. To see how we can approximate it, we examine the matrix $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$ more closely. $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$ is defined by

$$\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}} = \begin{bmatrix} c_{\tilde{\Lambda}, \tilde{\Lambda}}[0] & c_{\tilde{\Lambda}, \tilde{\Lambda}}[1] & \dots & c_{\tilde{\Lambda}, \tilde{\Lambda}}[m-1] \\ c_{\tilde{\Lambda}, \tilde{\Lambda}}[1] & c_{\tilde{\Lambda}, \tilde{\Lambda}}[0] & \dots & c_{\tilde{\Lambda}, \tilde{\Lambda}}[m-2] \\ \vdots & \vdots & \ddots & \vdots \\ c_{\tilde{\Lambda}, \tilde{\Lambda}}[m-1] & c_{\tilde{\Lambda}, \tilde{\Lambda}}[m-2] & \dots & c_{\tilde{\Lambda}, \tilde{\Lambda}}[0] \end{bmatrix} \quad (2.22)$$

where

$$c_{\tilde{\Lambda}, \tilde{\Lambda}}[p] = \mathbb{E} \left[\tilde{\Lambda}_k \tilde{\Lambda}_{k-p} \right] \quad (2.23)$$

Now, we recall Section 2.2, where we stated that $\tilde{\Lambda}_k \sim \text{Pois}[\lambda]$. The statistics of Poisson processes are well characterized, and we can apply that knowledge to the expectation in Equation 2.23. There are, generally speaking, two cases of interest; when $p = 0$ and when $p \neq 0$. When $p = 0$, we can simply write

$$\begin{aligned} c_{\tilde{\Lambda}, \tilde{\Lambda}}[0] &= \mathbb{E} \left[\tilde{\Lambda}_k \tilde{\Lambda}_k \right] \\ &= \text{var} \left(\tilde{\Lambda}_k \right) \\ &= \lambda_k \end{aligned} \quad (2.24)$$

On the other hand, when $p \neq 0$, we evaluate Equation 2.23 as follows:

$$\begin{aligned} c_{\tilde{\Lambda}, \tilde{\Lambda}}[p] &= \mathbb{E} \left[\tilde{\Lambda}_k \tilde{\Lambda}_{k-p} \right] \\ &= \text{Cov} \left[\tilde{\Lambda}_k, \tilde{\Lambda}_{k-p} \right] + \mathbb{E} \left[\tilde{\Lambda}_k \right] \mathbb{E} \left[\tilde{\Lambda}_{k-p} \right] \\ &= \lambda_k \lambda_{k-p} \end{aligned} \quad (2.25)$$

The last line in Equation 2.25 makes use of the fact that two measurements are independent realizations of the same Poisson process, and so the covariance between them is zero. Substituting 2.24 and 2.25 into 2.22 yields the following result.

$$\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}} = \begin{bmatrix} \lambda_k & \lambda_k \lambda_{k+1} & \dots & \lambda_k \lambda_{k+m-1} \\ \lambda_k \lambda_{k-1} & \lambda_k & \dots & \lambda_k \lambda_{k+m-2} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_k \lambda_{k-m+1} & \lambda_k \lambda_{k-m+2} & \dots & \lambda_k \end{bmatrix} \quad (2.26)$$

Next we note that since $\lambda_k \ll 1$, this implies $\lambda_k \lambda_{k-p} \approx 0$. With this approximation, $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$ simply becomes a scaled version of the identity matrix, $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}} = \lambda \mathbf{I}_{m \times m}$. Under

this assumption then, we can say that $\mathbf{w} \approx \lambda \mathbf{c}_{\tilde{\Lambda}, \lambda_O}$. We note that we are only interested in finding the peak of the correlation vector, so the fact that the correlation vector is scaled by flux does not impact our ability to determine TDOA.

As we stated above, we have many tools at our disposal to find \mathbf{w} , so according to this analysis, we simply need to select an appropriate adaptive filtering algorithm, run it on the incoming signal $\tilde{\Lambda}$, and the resultant filter should be the correlation vector between $\tilde{\Lambda}$ and the true signal, λ_O . Furthermore, not only do we have tools to estimate \mathbf{w} , but these estimators will also estimate the *covariance* of \mathbf{w} . These two points are the key takeaways from this section: *We can estimate the correlation between the true signal and the observed signal, and its variance, simply by estimating the optimal FIR filter for the observed signal, given that the observed signal is weak enough.*

It is interesting to note that the assumption of $\lambda \ll 1$ is a key component of this analysis, (as it lets us make the identity approximation), but at the same time, the fact that $\lambda \ll 1$ was what made the problem challenging. We might reasonably be skeptical of the claim that $\lambda \ll 1$ actually makes our estimation problem easier rather than harder. In fact, it does some of both.

On the one hand, $\lambda \ll 1$ does make our analysis easier, not only because it makes our approximation $\mathbf{c}_{\tilde{\Lambda}, \lambda_O} \approx \mathbf{w}$ reasonable, but also because it makes the matrix $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$ well-conditioned. Even though we will never actually invert $\mathbf{M}_{\tilde{\Lambda}, \tilde{\Lambda}}$, the fact that it is well-conditioned actually helps the algorithms that we will end up using maintain stability.

On the other hand, as noted multiple times previously, $\lambda \ll 1$ is essentially the source of the difficulty to begin with. If λ were not much less than 1, then we would not need to estimate $\mathbf{c}_{\tilde{\Lambda}, \lambda_O}$ to begin with; we could just compute it directly. At some large enough value of λ , the effect of the individual photon arrivals would be smoothed out by the characteristics of the signal, and so the particular challenge addressed here would disappear. So, as we would suspect intuitively, $\lambda \ll 1$ is less desirable than $\lambda \gg 1$. But, there is a “silver lining” in that $\lambda \ll 1$ allows the algorithms described here to function.

It turns out that, in cases where the expected photon flux is large, this method can still be used with some modification, as will be shown in Section 3.1.3. So, in the cases where flux is large enough for the $\lambda \ll 1$ approximation to be incorrect, but where flux

is still not large enough to estimate correlation directly, this technique can still be used.

Nomenclature

Correlation Vector Estimation

$c_{A(t),B(t)}[p]$	The correlation between signals $A(t)$ and $B(t+p)$
n	Dimension (number of taps) in the correlation vector
p	Index to denote a specific correlation vector or FIR filter tap
$\mathbf{M}_{A(t),B(t)}$	The correlation matrix (auto or cross) between a signal and itself ($\mathbf{M}_{A(t)}$), or a signal and another signal ($\mathbf{M}_{A(t),B(t)}$)
\mathbf{w}_λ	Ideal FIR filter for estimating signal λ from noisy measurement of signal

Data Association

p	Indices to designate signal source
-----	------------------------------------

Probability and Statistics

$\Pr[A]$	Probability of event A occurring
$\text{var}(a)$	The variance of random variable a
$\mathbb{E}[a]$	The expected value of random variable a
$\exp[\lambda]$	An exponential distribution with flux λ
$f_X(x)$	The probability density function of a random variable X evaluated at x

Signals and Measurements

$\hat{\Lambda}_{A,k}$	The number of photons measured at location A over time-step k (<i>Photons</i>)
$\lambda_k^{(p)}, \lambda^{(p)}(t)$	Expected number of photons at time index k , or expected photon flux at time t , from signal source p (<i>photons, photons/second</i>)
t_i	Time-of-arrival measurement of photon i (<i>seconds</i>)
S^i	i^{th} signal source

State Estimation

O	Denotes the navigation frame origin
j, k	Time-step indices
t	Current time (<i>seconds</i>)
$\tau_t^{A(S)}$	Time-difference of arrival of signal from signal source S measured between vehicle A and navigation frame origin (<i>seconds</i>)
m	Dimension of state vector
\mathbf{x}	State vector
n	Length of time between filter time-steps (<i>seconds</i>)
$[\cdot]^-$	Indicates <i>a priori</i> value
$[\cdot]^+$	Indicates <i>a posteriori</i> value
$\hat{\cdot}$	Indicates estimated value
\mathbf{a}	Indicates vector quantity
a	Indicates scalar quantity

Misc

$\mathbf{I}_{n \times n}$	The $n \times n$ identity matrix
i	Iterator to designate summation index
A, B	Indices to designate spacecraft(s)

Chapter 3

Recursive Range Estimation

In this chapter, an Extended Kalman formulation of the TDOA estimation problem for pulsar signals is presented [35]. Unlike the formulations in [15, 14, 13, 17, 14], the approach presented is recursive, in that the TDOA estimate and kinematic state of the spacecraft may be updated after each photon arrival. The updated kinematic state may then be used in future photon measurements to compute the optimal TDOA solution. The recursive nature of the algorithm presented would be particularly beneficial in cases where the original estimate of the kinematic state is of low accuracy, because it would allow the vehicle to update its kinematic state after each photon arrival, rather than waiting until the end of a predetermined observation period.

Another advantage of the method proposed here is that it does not rely on the particular periodic nature of the pulsar signal, unlike the methods described in [15, 14, 13, 17, 14]. In principle, the method proposed here could be used on a signal which is aperiodic in nature, such as signals produced by quasars or γ -ray bursts, although if this algorithm were to be used with an aperiodic signal, the time-history of that signal at the navigation frame origin would still be needed in order for the algorithm to function. This general functionality of the algorithm proposed here comes at the cost of increased complexity, which arises from the fact that this algorithm makes no assumptions about the source signal other than that it has stationary statistics.

A high level architecture of the approach proposed is shown in Figure 3.1, and is reminiscent of Vector Delay Lock Loops (VDLL) used to track weak GNSS signals [36]. A key feature of this algorithm is that it does not process the photons in a batch manner.

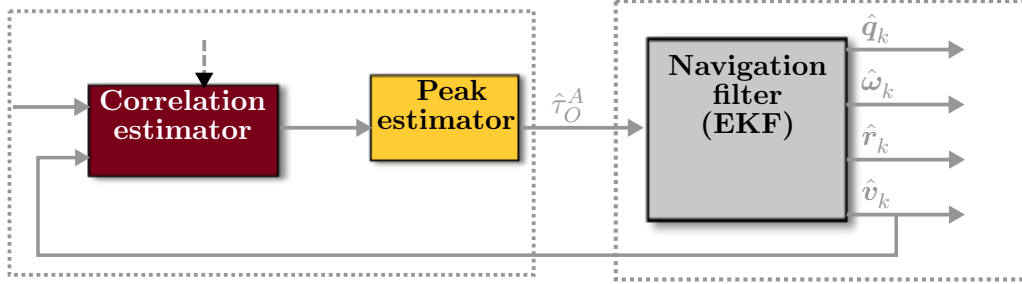


Figure 3.1: X-ray navigation filter block diagram. The filter outputs are the complete PNT state vector consisting of the attitude quaternion \mathbf{q} , the angular velocity rate $\boldsymbol{\omega}$, the position vector \mathbf{r} , and the velocity vector \mathbf{v} .

Rather, each photon is processed individually to update the TDOA estimate, and at each update, an estimate of the vehicle's current kinematic state is used to aide in estimating the current TDOA.

As mentioned previously, the key quantity that needs to be determined in order to effectively navigate via signals of opportunity is the arrival time, or more precisely the time difference of arrival (TDOA) between the user and another location where the signal is measured or known *a priori*. This is challenging with signals from pulsars because, with the detector sizes available for small satellites such as CubeSats, typical photon flux rates are less than one photon per second. Even though the problem becomes easier with largest detectors, it is still very challenging owing to the extremely low signal-to-noise ratios of pulsar signals. While in theory, a large enough detector could help to address some of the problems caused by the low SNR, in practice, the cost and size of a detector required to reduce the noise associated with the Poisson statistics of the signal makes such a detector unfeasible for use on a spacecraft. It is not immediately obvious that it would even be possible to determine the TDOA of such a sparse signal. However, we can take advantage of the fact that the statistics of the underlying signal remain unchanged throughout the duration of the observation.

The most obvious means of determining the TDOA between two signals is to compute the cross-correlation vector between those signals. The discrete-time cross-correlation function between a signal $\lambda_{A,k}$ and a delayed version of that signal $\lambda_{O,k}$ is given by:

$$c_{\lambda_A, \lambda_O}[p] = \sum_i \lambda_A(t + iT) \lambda_O(t + (i + p)T) \quad (3.1)$$

where

- T is the discrete time-step length
- k is the time index
- p is the cross-correlation vector index
- $D = \frac{\Delta T}{T}$ is the delay between the two signals, in discrete time-steps
- i is the index of summation

The cross correlation bears many similarities to the discrete convolution, without the time-reversal. It is well-known that the cross-correlation between a signal and a delayed version of the signal is maximized at the value of the offset between the two signals [32, 33], i.e. $p = D = \frac{\Delta T}{T}$. Note that the signal seen at the detector is a version of the signal at the origin, delayed by $\frac{\Delta T}{T}$ time-steps. Referring to the NAVSOP concept shown in Figure 1.2, this relationship between the two signals may be written mathematically as follows:

$$\lambda_A(t) = \lambda_O\left(t + \frac{\mathbf{r}_t^A \cdot \boldsymbol{\ell}^{(S)}}{c}\right) = \lambda_O(t + \tau_O^A) \quad (3.2)$$

Therefore, if we could directly compute the cross-correlation between λ_O and λ_A , we could find the value that maximizes that cross-correlation and from there directly find the TDOA. However, as discussed in Section 2.2, the signals received from pulsars are usually quite weak, and therefore a direct cross-correlation is dominated by the noise of individual photon arrivals.

We make use of the fact that for very weak signals, and in particular signals which are a realization of a Poisson process with a very small rate parameter, the correlation

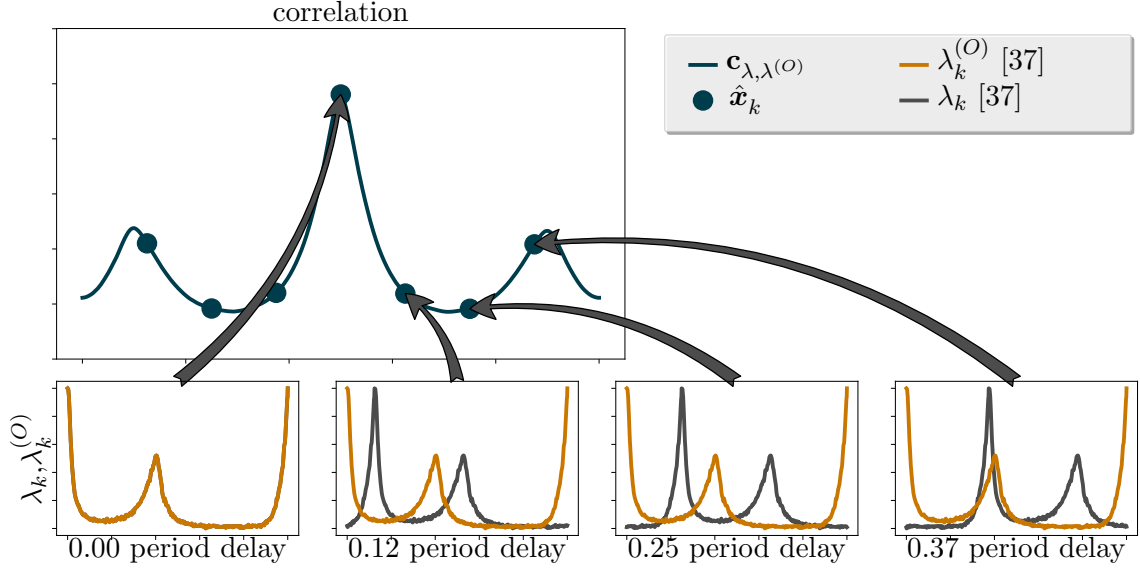


Figure 3.2: Correlation function as state vector

vector between the measurement and the true signal is equivalent to the ideal finite-impulse response (FIR) filter for reconstructing the measured signal, as was shown in Section 2.2.4. Here, we make use of that fact, and formulate an Extended Kalman Filter with a state vector that consists of the ideal FIR filter taps to estimate the correlation vector.

3.1 Correlation Vector Kalman Filter

The first task to be addressed is the implementation of a filter which will estimate the ideal FIR filter taps, \mathbf{w}_{λ_O} . As noted above, we achieve this by designing a Kalman filter that has as its states values of the correlation function between the signal at the spacecraft's location and the signal at the origin at various delays. The state vector $\hat{\mathbf{x}}_k \in \mathbb{R}^{m \times 1}$ is illustrated in Figure 3.2. The optimal estimate of the state vector is defined as:

$$\begin{aligned}\mathbb{E}[\hat{x}_{k,p}] &= c_{\lambda,\lambda_O} [p + d] \\ &= \sum_i \lambda_{A,k+i} \lambda_{O,k+i+p+d}\end{aligned}\tag{3.3}$$

where d is an offset term which will be described below. Recall from Equation (2.6) and Equation (2.7) that the expected number of photons at a given time ($\lambda_{A,k+i}$) may be approximated by taking the continuous-time photon flux at that time and multiplying it by the desired time-step.

Modern adaptive signal processing provides us with a variety of adaptive filtering techniques for this estimation task, including LMS (Least Mean Squares), RLS (Recursive Least Squares), and Kalman filtering. For this particular estimation problem, the Kalman filter is the most well-suited. As shown in [38], the RLS filter is essentially a Kalman filter without a time update model. Because we know with certainty that the system *is* dynamic, it makes sense to employ an algorithm that allows us to account for that dynamic change. We will follow the standard Kalman filter formulation which uses a time update and measurement update, as described in a number of standard textbooks [38, 39].

The dynamic nature of the problem and the formulation of a Kalman filter time-update requires the introduction of an additional term in the time-index of the origin signal, d . This term is introduced because the state vector indices, p only take integer values from 0 to m , the dimension of the state vector. However, it is entirely possible that the delay $\tau = \frac{\Delta T}{T}$ might fall outside of this range. We are interested in computing the correlation vector near the value of the TDOA, so we will almost certainly have to “shift” the state vector some amount forward or backward in time to stay near the value of the TDOA. The term d allows us to perform this shifting by adjusting the value of d as needed. This procedure will be discussed in more detail in Section 3.1.1 below.

3.1.1 Correlation Vector Time Update

The first step we want to accomplish is to formulate a way to estimate how the state $\hat{\mathbf{x}}_k$ will change from time-step j to time-step k between photon measurements. Even though the state vector does not directly represent the spacecraft’s kinematic state (the




Figure 3.3: Correlation peak shifted by δ

state vector is a correlation between two signals), the state vector will evolve predictably in time such that a dynamic model may be derived.

Suppose that at time-step j , the spacecraft is moving along at a velocity \mathbf{v}_O^A with respect to the solar-system barycenter (SSB), and that there is an estimate of that velocity, $\hat{\mathbf{v}}_O^A$. Suppose further that the direction to the signal source is known in the form of the unit vector, $\ell^{(S)}$. Finally, assume that there is a prior estimate of the state vector, $\hat{\mathbf{x}}_j^+$ (that is, the estimate of the state vector after the measurement update at time j).

Since the spacecraft is moving at some velocity relative to the signal source, and since the distance between the signal source and the spacecraft is changing, the correlation peak ought to move some small amount between time-step j and time-step k . The amount by which the peak is shifted may be calculated, given the estimate of velocity. The evolution of the state vector is illustrated graphically in Figure 3.3. Mathematically, the shift from time-index j to k , $\hat{\delta}_{k \rightarrow j}$ may be determined by solving the following first-order differential equation:

$$\frac{d}{dt} \hat{\delta}_{j \rightarrow k} = \frac{\hat{\mathbf{v}}_O^A(t) \cdot \boldsymbol{\ell}^{(S)}}{cT} \quad (3.4)$$

where T is the time-step of the filter. If the time-step is assumed to be small enough such that velocity may be approximated as constant, then $\hat{\delta}_{k \rightarrow j}$ may be estimated as follows:

$$\hat{\delta}_{k \rightarrow j} = \frac{\hat{\mathbf{v}}_{O,k}^A \cdot \boldsymbol{\ell}^{(S)}}{cT} (k - j)T = \frac{\hat{\mathbf{v}}_{O,k}^A \cdot \boldsymbol{\ell}^{(S)}}{c} (k - j) \quad (3.5)$$

where $\hat{\mathbf{v}}_{O,k}^A$ is the velocity vector at time index k .

The expression for $\hat{\delta}_{k \rightarrow j}$ may now be used to construct the state transition matrix, $F_{j \rightarrow k}$. The state transition matrix must, when multiplied by $\hat{\mathbf{x}}_j^+$, shift the peaks by the amount $\hat{\delta}_{k \rightarrow j}$. There are, in general, three possibilities that should be considered. The first, and simplest situation is when $\hat{\delta}_{k \rightarrow j} = 0$. This would indicate that the spacecraft has neither moved closer to nor farther from the signal source. This situation would arise either when the spacecraft is not moving, or when the spacecraft is moving in a direction perpendicular to the line of sight vector $\boldsymbol{\ell}^{(S)}$ between the source and spacecraft, resulting in the dot product in Equation (3.5) equaling zero. In this scenario, the state transition matrix would simply be the identity matrix. This would give the desired result, namely that the peaks of the state vector remain in exactly the same location.

Another scenario is when $\hat{\delta}_{k \rightarrow j} = z$, where z is an integer. This scenario is very unlikely, for the obvious reason that the probability of $\hat{\mathbf{v}}_{O,k}^A \cdot \boldsymbol{\ell}^{(S)}(k - j)$ being an exact multiple of c is very low. This would have the effect of moving all the features of the state vector exactly z steps, which is the desired result if $\hat{\delta}_{k \rightarrow j}$ is truly an integer.

The last, and by far the most probable scenario is when $-1 < \hat{\delta}_{k \rightarrow j} < 1$. That is, the desired delay is some fractional value with a magnitude less than one. For this scenario, we need a state transition which will result in a fractional delay. Put another way, the state transition matrix $F_{j \rightarrow k}$ must “interpolate” between the values of $\hat{\mathbf{x}}_j^+$. The resultant state vector should consist of the values that lie “between” the current values of $\hat{\mathbf{x}}_j^+$. This effect is equivalent to the effect produced by a fractional delay filter.

Much work has been done on the subject of fractional delay FIR filters [40, 41, 42, 43].

For the purposes of this dissertation, a simple sampled $\text{sinc}\left(n + \hat{\delta}_{k \rightarrow j}\right)$ function will provide sufficiently accurate results for all three possible cases noted above, since the *sinc* function is the ideal interpolator. Therefore, in general the state transition matrix is described by

$$F_{j \rightarrow k} = \begin{bmatrix} \text{sinc}\left(\hat{\delta}_{k \rightarrow j}\right) & \text{sinc}\left(\hat{\delta}_{k \rightarrow j} + 1\right) & \dots & \text{sinc}\left(\hat{\delta}_{k \rightarrow j} + m - 1\right) \\ \text{sinc}\left(\hat{\delta}_{k \rightarrow j} - 1\right) & \text{sinc}\left(\hat{\delta}_{k \rightarrow j}\right) & \dots & \text{sinc}\left(\hat{\delta}_{k \rightarrow j} + m - 2\right) \\ \vdots & \vdots & \ddots & \vdots \\ \text{sinc}\left(\hat{\delta}_{k \rightarrow j} - m + 1\right) & \text{sinc}\left(\hat{\delta}_{k \rightarrow j} - m + 2\right) & \dots & \text{sinc}\left(\hat{\delta}_{k \rightarrow j}\right) \end{bmatrix} \quad (3.6)$$

This matrix may be used to write the time update of covariance from time-step to time-step, via the discrete Lyapunov equation as follows:

$$\mathbf{P}_k^- = F_{j \rightarrow k} \mathbf{P}_j^+ F_{j \rightarrow k}^T + L_j \mathbf{Q}_{j \rightarrow k, \delta} L_j^T \quad (3.7)$$

The process noise covariance matrix $\mathbf{Q}_{j \rightarrow k, \delta}$ and the process noise mapping matrix L_j need to be determined. To do this, note that the main uncertainty in propagation comes from the uncertainty in the velocity of the space vehicle, which in turn makes the quantity $\hat{\delta}_{k \rightarrow j}$ uncertain. Consider the following model for the estimate of velocity:

$$\hat{\mathbf{v}}_j = \mathbf{v}_j + \mathbf{w}_{j, \mathbf{v}} \quad (3.8a)$$

$$\mathbf{w}_{j, \mathbf{v}} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{j, \mathbf{v}}) \quad (3.8b)$$

Then, from Equation (3.5), the estimate of shift will be given by

$$\begin{aligned}
\hat{\delta}_{j \rightarrow k} &= \frac{\hat{\mathbf{v}}_j \cdot \boldsymbol{\ell}^{(S)}(k-j)}{c} \\
&= \frac{(\mathbf{v}_j + \mathbf{w}_{j,\mathbf{v}})^T \cdot \boldsymbol{\ell}^{(S)}(k-j)}{c} \\
&= \frac{\mathbf{v}_j^T \cdot \boldsymbol{\ell}^{(S)}(k-j)}{c} + \frac{\mathbf{w}_{j,\mathbf{v}}^T \cdot \boldsymbol{\ell}^{(S)}(k-j)}{c} \\
&= \hat{\delta}_{k \rightarrow j} + \frac{\mathbf{w}_{j,\mathbf{v}}^T \cdot \boldsymbol{\ell}^{(S)}(k-j)}{c} \\
&= \hat{\delta}_{k \rightarrow j} + w_{j,\delta}
\end{aligned} \tag{3.9}$$

The variance of the error in the estimate of $\hat{\delta}_{j \rightarrow k}$ may then be written as:

$$w_{j,\delta} \sim \mathcal{N}[0, Q_{j,\delta}] \tag{3.10a}$$

$$Q_{j,\delta} = \left(\frac{(k-j)}{c} \right)^2 \boldsymbol{\ell}^{(S)} \mathbf{w}_{j,\mathbf{v}} \boldsymbol{\ell}^{(S)T} \tag{3.10b}$$

Although the error in $\hat{\delta}_{j \rightarrow k}$ is additive, it is clear upon examination of Equation 3.6 that the final effect on the uncertainty of the state estimate is not linear. In EKFs, this type of non-linearity is handled by linearization. In order to linearize, we first rewrite the time-update as the following non-linear function of the state, the estimated change in peak location, and noise.

$$\begin{aligned}
\hat{\mathbf{x}}_k^- &= f(\hat{\mathbf{x}}_j^+, \hat{\delta}_{j \rightarrow k}, w_{j,\delta}) = F_{j \rightarrow k}(\hat{\delta}_{j \rightarrow k}, w_{j,\delta}) \hat{\mathbf{x}}_j^- \\
&= \begin{bmatrix} \text{sinc}(\delta + w_\delta) & \text{sinc}(\delta + w_\delta + 1) & \dots & \text{sinc}(\delta + w_\delta + m - 1) \\ \text{sinc}(\delta + w_\delta - 1) & \text{sinc}(\delta + w_\delta) & \dots & \text{sinc}(\delta + w_\delta + m - 2) \\ \vdots & \vdots & \ddots & \vdots \\ \text{sinc}(\delta + w_\delta - m + 1) & \text{sinc}(\delta + w_\delta - m + 2) & \dots & \text{sinc}(\delta + w_\delta) \end{bmatrix} \hat{\mathbf{x}}_j^+
\end{aligned} \tag{3.11}$$

The process noise mapping matrix may then be derived by linearizing the time update equation, $f(\hat{\mathbf{x}}_j^+, \hat{\delta}_{j \rightarrow k}, w_{j,\delta})$ with respect to the process noise.

$$\begin{aligned}
L_j &= \frac{\partial}{\partial w_\delta} f(\hat{\mathbf{x}}_j^+, \delta, w_\delta) \Big|_{w_\delta=0} \\
&= \frac{\partial}{\partial w_\delta} \left[\begin{array}{cccc} \text{sinc}(\delta + w_\delta) & \text{sinc}(\delta + w_\delta + 1) & \dots & \text{sinc}(\delta + w_\delta + m - 1) \\ \text{sinc}(\delta + w_\delta - 1) & \text{sinc}(\delta + w_\delta) & \dots & \text{sinc}(\delta + w_\delta + m - 2) \\ \vdots & \vdots & \ddots & \vdots \\ \text{sinc}(\delta + w_\delta - m + 1) & \text{sinc}(\delta + w_\delta - m + 2) & \dots & \text{sinc}(\delta + w_\delta) \end{array} \right] \hat{\mathbf{x}}_j \Big|_{w_\delta=0} \\
&= \lim_{\delta \rightarrow \hat{\delta}} \left[\begin{array}{ccc} \left(\frac{\cos(\delta)}{\delta} - \frac{\sin(\delta)}{(\delta)^2} \right) & \left(\frac{\cos(\delta+1)}{\delta+1} - \frac{\sin(\delta+1)}{(\delta+1)^2} \right) & \dots \\ \left(\frac{\cos(\delta-1)}{\delta-1} - \frac{\sin(\delta-1)}{(\delta-1)^2} \right) & \left(\frac{\cos(\delta)}{\delta} - \frac{\sin(\delta)}{(\delta)^2} \right) & \dots \\ \vdots & \vdots & \ddots \end{array} \right] \hat{\mathbf{x}}_j
\end{aligned} \tag{3.12}$$

Note that the final expression contains a singularity at $\hat{\delta} = 0$. However, a simple application of L'Hopital's rule reveals the limit to be zero.

At this point, the main source of error in the time-update has been accounted for. However, we also include an additive noise term, \mathbf{Q}_{lin} , to account for errors introduced by linearization, and sampling of the $\text{sinc}(\cdot)$ function. The magnitude of the uncertainty introduced by the $\text{sinc}(\cdot)$ function and the linearization could be quantified analytically, which would inform the value used for \mathbf{Q}_{lin} . However, in this work, the parameter was experimentally tuned using real data and simulation data.

In summary, the EKF time update equations for propagating the state from one time-step to the next are given below.

$$\hat{\mathbf{x}}_k^- = \mathbf{F}_{j \rightarrow k} \hat{\mathbf{x}}_j^+ \quad (3.13a)$$

$$\mathbf{P}_k^- = \mathbf{F}_{j \rightarrow k} \mathbf{P}_j^+ \mathbf{F}_{j \rightarrow k}^T + \mathbf{L}_j \mathbf{Q}_{j,\delta} \mathbf{L}_j^T + \mathbf{Q}_{\text{lin}} \quad (3.13b)$$

$$\mathbf{F}_{j \rightarrow k} = \begin{bmatrix} \text{sinc}(\hat{\delta}) & \dots & \text{sinc}(\hat{\delta} + m - 1) \\ \vdots & \ddots & \vdots \\ \text{sinc}(\hat{\delta} - m + 1) & \dots & \text{sinc}(\hat{\delta}) \end{bmatrix} \quad (3.13c)$$

$$\mathbf{L}_j = \lim_{\delta \rightarrow \hat{\delta}} \begin{bmatrix} \frac{\cos(\delta)}{\delta} - \frac{\sin(\delta)}{(\delta)^2} & \dots \\ \vdots & \ddots \end{bmatrix} \hat{\mathbf{x}}_j \quad (3.13d)$$

$$\hat{\delta}_{j \rightarrow k} = \frac{\mathbf{v}_{O,k}^A \cdot \boldsymbol{\ell}^{(S)}(k - j)}{c} \quad (3.13e)$$

3.1.2 Alternative implementation of time update

In Section 3.1.1, the time update equations for the correlation vector were derived. In this section, an alternative means of performing the time update is presented. This alternative method leverages the derivations presented in Section 3.1.1, but reduces computational complexity and reduces some of the error associated with the use of the $\text{sinc}(\cdot)$ function for performing the time update.

We begin by noting that the end-goal is to estimate TDOA by estimating the correlation vector. We also note that the correlation vector consists of values which are inherently linked to a given TDOA. Recall from Equation (3.3) that the expected value of a single component of the state vector is an estimate of the correlation vector between the signal at the spacecraft and the signal at the origin shifted by some amount $p + d$.

In Section 3.1.1, we derived time-update equations which would have the effect of shifting the state vector forward or backward in time by some amount δ . However, we could achieve the same effect by keeping the state vector values the same, and instead shifting the value of the delays to which each state vector component corresponds; essentially shifting the “x-axis” of the correlation function rather than shifting the correlation

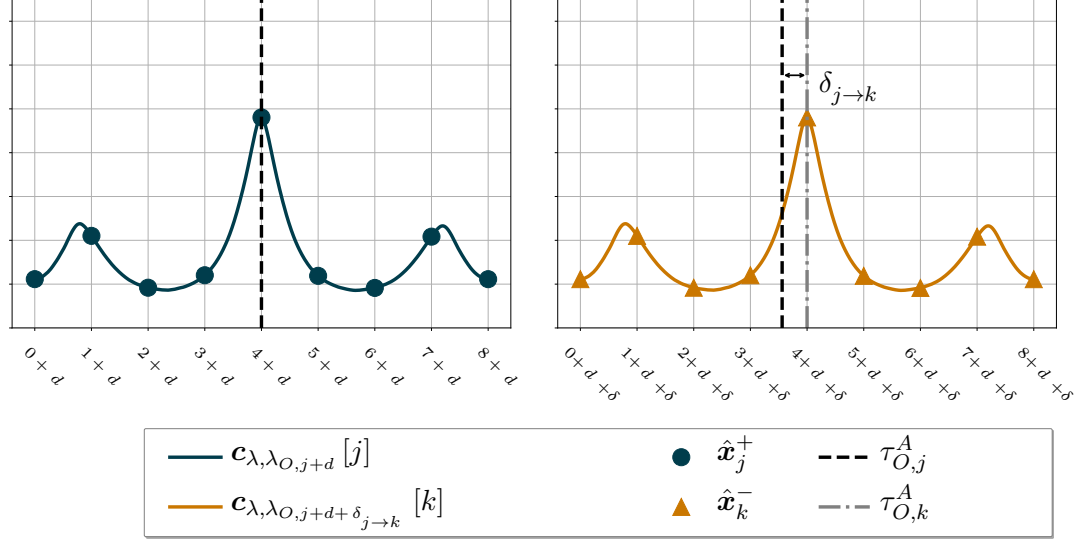


Figure 3.4: TDOA axis shifted by δ

vector itself. Mathematically we may express this as follows:

$$\mathbb{E} \left[\hat{x}_{j,p}^+ \right] = c_{\lambda, \lambda_O} [p + d] \quad (3.14a)$$

$$\mathbb{E} \left[\hat{x}_{k,p}^- \right] = \mathbb{E} \left[\hat{x}_{j,p}^+ \right] = c_{\lambda, \lambda_O} [p + d + \hat{\delta}] \quad (3.14b)$$

Equation (3.14) highlights the fact that in this alternative time update, the value of the state vector remains unchanged; instead the *meaning* of that value changes. This alternative update to the state vector is illustrated in Figure 3.4.

While the *values* of the state vector remain unchanged, the covariance is expected to increase with the time update, due to uncertainty in the velocity. We recall the derivation for the updated covariance from Section 3.1.1.

$$\mathbf{P}_k^- = \mathbf{F}_{j \rightarrow k} \mathbf{P}_j^+ \mathbf{F}_{j \rightarrow k}^T + \mathbf{L}_j \mathbf{Q}_{j,\delta} \mathbf{L}_j^T + \mathbf{Q}_{\text{lin}} \quad (3.15)$$

We note that there are three terms for the time-updated covariance. The term $\mathbf{F}_{j \rightarrow k} \mathbf{P}_j^+ \mathbf{F}_{j \rightarrow k}^T$ has the effect of interpolating the covariance, just as the term $\mathbf{F}_{j \rightarrow k} \hat{\mathbf{x}}_j^{+T}$ interpolates the state vector. Since in this time-update formulation, we do not shift the

state vector, the covariance matrix does not need to be shifted either.

The second term in the time-update covariance, $\mathbf{L}_j \mathbf{Q}_{j,\delta} \mathbf{L}_j^T$ models the uncertainty that is introduced due to the uncertainty in the velocity. This uncertainty is still present regardless of the formulation used, as long as there is uncertainty in the velocity, so this term remains in the time-updated covariance.

Finally, the term \mathbf{Q}_{lin} is representative of additional, unmodeled noise sources, including noise introduced by linearization. Since the same linearization is still used in this alternative formulation, we also still add this final additive noise source. However we note that the magnitude of the additive noise source needed to achieve good results is typically lower than what is needed using the time-update formulations derived in the previous section.

In summary, the alternative time-update equations are given as follows.

$$\hat{\mathbf{x}}_k^- = \hat{\mathbf{x}}_k^+ \quad (3.16a)$$

$$\mathbf{P}_k^- = \mathbf{L}_j \mathbf{Q}_{j,\delta} \mathbf{L}_j^T + \mathbf{Q}_{\text{lin}} \quad (3.16b)$$

$$\mathbf{L}_j = \lim_{\delta \rightarrow \hat{\delta}} \begin{bmatrix} \frac{\cos(\delta)}{\delta} - \frac{\sin(\delta)}{(\delta)^2} & \cdots \\ \vdots & \ddots \end{bmatrix} \hat{\mathbf{x}}_j \quad (3.16c)$$

$$\hat{\delta}_{j \rightarrow k} = \frac{\mathbf{v}_{O,k}^A \cdot \ell^{(S)}(k-j)}{c} \quad (3.16d)$$

It should be noted that while in this formulation of the time update, the value of the state vector remains constant, the value of the state vector will change during the measurement update. The equations which update the value of the state vector based on measured values will be derived below in Section 3.1.3 and Section 3.1.4.

3.1.3 Correlation Vector Measurement Update

We now turn our attention to the measurement update equations. Recall that our measurement consists of a series of photon arrival times. The measurement vector, $\tilde{\mathbf{A}}_k$ is formulated as the vector consisting of the number of photons that arrived in each of the last m time indices. We seek a measurement model that will relate this measurement vector to the state vector we are trying to estimate. To do this, we make use of the

fact that the correlation vector is equivalent to a scaled version of the ideal FIR filter, as shown in Section 2.2.4. We note the result from signal processing theory [38], that the ideal FIR filter \mathbf{w} of length m which enables the reconstruction of $\lambda_{O,k}$ from noise corrupted measurement $\tilde{\Lambda}_k$ is given by:

$$\lambda_{O,k} = \tilde{\Lambda}_k \mathbf{w}^T + v \quad (3.17)$$

where the error, v , is minimized in the least-squared sense by the ideal FIR filter.

To cast Equation 3.17 into a form useful for estimation, we use y_k to denote our measurement and recall that ideally our state estimate \mathbf{x} would be equal to \mathbf{w} . Thus, the measurement update may be written as follows.

$$y_k = \lambda_{O,k} = h(\hat{\mathbf{x}}_k^-, \mathbf{\Lambda}_k, \mathbf{v}_{\text{pois}}) \quad (3.18a)$$

$$\hat{y}_k = \tilde{\Lambda}_k \hat{\mathbf{x}}_k^T = (\mathbf{\Lambda}_k + \mathbf{v}_{\text{pois}})^T \hat{\mathbf{x}}_k \quad (3.18b)$$

This measurement update equation does not follow the “signal + noise” model of a linear problem. Rather, the noise is introduced through $\tilde{\Lambda}(k)$ and multiplied by the state vector. We recall that the state vector is defined as the optimal FIR filter, which is defined such that the expected value of the dot product of the state vector and the signal vector is equal to the signal at the origin (as shown in Equation 3.17). Therefore the expected value of the measurement residual may be written as follows:

$$\begin{aligned} \delta \mathbf{y} &= y_k - \hat{y}_k \\ &= \lambda_{O,k} - \tilde{\Lambda}_k \hat{\mathbf{x}}_k^{-T} \\ &= \lambda_{O,k} - (\mathbf{\Lambda}_{A,k} + \mathbf{v}_{\text{pois}}) \hat{\mathbf{x}}_k^{-T} \end{aligned} \quad (3.19)$$

The variance of this measurement residual may be analyzed as follows. Note that the following analysis makes use of the following properties:

- $\text{var}(\lambda_O) = 0$
- $\text{var}(\lambda_A) = 0$

- $\mathbb{E} \left[\mathbf{v}_{\text{pois}} \right] = 0$
- $\text{var} \left(\mathbf{v}_{\text{pois}} \right) = \lambda_A$

The analysis also assumes that the individual terms in the summation resulting from the vector dot-product are uncorrelated. Finally, second order terms are neglected.

$$\begin{aligned}
S &= \text{var} (\delta y) \\
&= \text{var} \left(\lambda_{O,k} - \tilde{\mathbf{\Lambda}}(k) \hat{\mathbf{x}}_k^{-T} \right) \\
&= \text{var} \left(\tilde{\mathbf{\Lambda}}(k) \hat{\mathbf{x}}_k^{-T} \right) \\
&= \tilde{\mathbf{\Lambda}}(k) \text{var} (\hat{\mathbf{x}}_k^-) \tilde{\mathbf{\Lambda}}(k)^T + \hat{\mathbf{x}}_k^- \text{var} \left(\tilde{\mathbf{\Lambda}}(k)^T \right) \hat{\mathbf{x}}_k^{-T} + \text{var} (\hat{\mathbf{x}}_k^-) \text{var} \left(\tilde{\mathbf{\Lambda}}(k)^T \right) \\
&= \tilde{\mathbf{\Lambda}}(k) \mathbf{P}_k^- \tilde{\mathbf{\Lambda}}(k)^T + \hat{\mathbf{x}}_k^- \mathbf{R}_{k,\tilde{\mathbf{\Lambda}}} \hat{\mathbf{x}}_k^{-T}
\end{aligned} \tag{3.20}$$

where the matrix $\mathbf{R}_{k,\tilde{\mathbf{\Lambda}}}$ is defined below. Note that this same result could be obtained by linearizing the measurement equation $y_k = h(\hat{\mathbf{x}}_k^-, \mathbf{\Lambda}_k, \mathbf{v}_{\text{pois}})$ with respect to the state vector and measurement noise.

The remaining task is to estimate the covariance of the noise associated with the photon arrival vector, $\mathbf{R}_{k,\tilde{\mathbf{\Lambda}}}$. We first recall the statistics of a Poisson process and note that the variance of such a process is equal to the flux of that process at a given time step. We next note that the number of photons which arrive over each time step is statistically independent of the photons which arrived at any other time step. Consequently, the measurement noise may be written as:

$$\mathbf{R}_{k,\tilde{\mathbf{\Lambda}}} = \begin{bmatrix} \lambda_{A,k} & 0 & \dots & 0 \\ 0 & \lambda_{A,k+1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_{A,k+m-1} \end{bmatrix} = \mathbf{I}_{m \times m} \mathbf{\Lambda}_{A,k} \tag{3.21}$$

Because the exact values of $\mathbf{\Lambda}_{A,k}$ are a function of the spacecraft's position, they are not known precisely. However, the time-averaged flux, $\bar{\lambda}_k$ is known precisely from the signal statistics, or can be approximated by taking the average photon arrival rate over

a long period of time. Using this average, we make the following approximation.

$$\mathbf{R}_{k,\tilde{\Lambda}} = \mathbf{I}_{m \times m} \mathbf{\Lambda}_{A,k} \approx \mathbf{I}_{m \times m} \bar{\lambda}_k \quad (3.22)$$

Consequently, the final measurement residual variance may be written as follows.

$$S = \tilde{\mathbf{\Lambda}}_k^T \mathbf{P}_k^- \tilde{\mathbf{\Lambda}}_k + \bar{\lambda}_k \hat{\mathbf{x}}_k^- \hat{\mathbf{x}}_k^{-T} \quad (3.23)$$

3.1.4 Alternate Formulation of Measurement Update Equations

The measurement update outlined in Section 3.1.3 will, in general, only update one state within the state vector at a time. This may be seen by examining the measurement update matrix $\mathbf{\Lambda}_k$ and noting that, in general, the photon arrival rate is much less than one photon per period. When the photon arrives, the first measurement matrix will be zeros except for the first quantity, which will be one. At the next time step, the first value in the matrix will be zero and the second will be one. The one, representative of the single photon arrival, will progress to the end of the measurement matrix, at which point the state update will be complete.

$$\tilde{\mathbf{\Lambda}}_k = \begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix} \quad (3.24a)$$

$$\tilde{\mathbf{\Lambda}}_{k+1} = \begin{bmatrix} 0 & 1 & \dots & 0 \end{bmatrix} \quad (3.24b)$$

$$\tilde{\mathbf{\Lambda}}_{k+m-1} = \begin{bmatrix} 0 & 0 & \dots & 1 \end{bmatrix} \quad (3.24c)$$

We further note that the “measured” quantity is actually not measured by the spacecraft, but is the photon flux at the origin, $\lambda_{O,k}$. Consequently, it is possible to know the measurement values for all the measurement updates as soon as a photon has arrived. Consequently, when a photon arrives, all the states may be updated simultaneously

rather than sequentially, as follows.

$$\mathbf{y} = \begin{bmatrix} \lambda_{O,k} \\ \lambda_{O,k+1} \\ \vdots \\ \lambda_{O,k+m-1} \end{bmatrix} = \mathbf{\Lambda}_{O,k}^T \quad (3.25a)$$

$$\hat{\mathbf{y}} = \begin{bmatrix} \tilde{\mathbf{\Lambda}}_k \\ \tilde{\mathbf{\Lambda}}_{k+1} \\ \dots \\ \tilde{\mathbf{\Lambda}}_{k+m-1} \end{bmatrix} \hat{\mathbf{x}}_k^{-T} = \begin{bmatrix} \begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix} \\ \begin{bmatrix} 0 & 1 & \dots & 0 \end{bmatrix} \\ \vdots \\ \begin{bmatrix} 0 & 0 & \dots & 1 \end{bmatrix} \end{bmatrix} \hat{\mathbf{x}}_k^{-T} = \mathbf{I}_{m \times m} \hat{\mathbf{x}}_k^{-T} \quad (3.25b)$$

Consequently, the measurement residual and residual variance are given by:

$$\delta \mathbf{y} = \mathbf{\Lambda}_{O,k}^T - \hat{\mathbf{x}}_k^{-T} \quad (3.26a)$$

$$\mathbf{S} = \bar{\lambda}_k \mathbf{I}_{m \times m} + \mathbf{P}_k^- \quad (3.26b)$$

This vectorized form of the measurement update offers two advantage of the sequential update derived above. First, it saves the computation time associated with performing m state updates. It will be noted that the vectorized formulation requires the inversion of an $m \times m$ residual variance matrix, while in the sequential update the residual variance is a scalar so no matrix inversion is performed. In terms of computation time, performing the matrix inversion is more computationally expensive than performing m divisions. However, as will be discussed in the next section, there are additional steps which must be performed after each measurement update. The added computational time of these additional steps outweighs the computational cost of performing the matrix inversion.

Second, and more fundamentally, the vectorized measurement update offers the advantage of updating all of the states simultaneously. This is advantageous because of the linearization which is used in the measurement update. In the sequential update, the first state in the state vector and its associated covariance, \hat{x}_1 and $P_{(1,1)}$ are updated first, before any of the other states. If the state estimate variance is large, particularly during initialization, the *a posteriori* value of \hat{x}_1 can be significantly off before the filter

converges. When the next state is updated, the erroneous value of \hat{x}_1 is used to compute the new Kalman gain, as seen in Equation (3.23). If the value is off by a significant amount, this can result in the measurement residual being far larger than it should be, resulting in an incorrectly small value for the next state’s Kalman gain matrix. This can lead to very inaccurate estimates of all the states. By updating all the states in the state vector at once, rather than sequentially, this problem may be avoided.

It will also be noted that the use of the vectorized measurement update makes an assumption about the weakness of the signal which may be violated in the case of very strong signals (i.e. the low flux assumption). However, it turns out that even in the case of very strong signals where this assumption does not hold, making the assumption is not problematic. This is because we assumed in the derivation of the filter that the signals are very weak, which leads to the equivalency of the ideal FIR filter and the correlation vector. Enforcing the weak signal assumption here actually leads to the FIR filter/correlation vector equivalency holding true, even for very strong signals. This is because when using the vectorized measurement update, each time the filter taps are updated, the signal they “see” is a very weak signal, even though the signal itself is strong. Consequently the filter taps end up estimating the correlation vector because they are estimating the filter that would be needed to estimate λ_{origin} from a very weak signal, even if the signal is in fact not weak.

3.2 Estimation of TDOA from Correlation Vector

Once we have an accurate estimate of the correlation vector, the next task is to identify the peak $\tau_{O,k}^A$ offset of that vector. In this work, the peak is estimated by fitting a quadratic function to a subset of $\hat{\mathbf{x}}_k^+$ surrounding the maximum value of $\hat{\mathbf{x}}_k^+$. This quadratic function is maximized, and the argument of the maximum of the function is taken as $\hat{\tau}_{O,k}^A$. This method has the advantage of allowing $\hat{\tau}_{O,k}^A$ to take values that are not integer multiples of T . It has the additional advantage of providing a convenient way to estimate the variance of $\hat{\tau}_{O,k}^A$. There are two ways which the quadratic fit method may be used to derive the variance of the TDOA estimate error, both of which are outlined below.

3.2.1 Linearized Transform

The most straight-forward way to determine the variance of the TDOA estimate error is to linearize the expression for determining the TDOA. We note first that the quadratic fitted to the peak and surrounding state vector may be described as:

$$f(\tau) = a \tau^2 + b \tau + c \quad (3.27)$$

Given these parameters, the function is maximized at:

$$\hat{\tau} = \frac{-b}{2a} \quad (3.28)$$

The task then is to estimate the parameters a and b from the state vector. This problem may be formulated as a system of equations:

$$\begin{bmatrix} \hat{x}_{k_1} \\ \hat{x}_{k_2} \\ \vdots \\ \hat{x}_{k_n} \end{bmatrix} = \begin{bmatrix} k_1 & k_1 & 1 \\ k_2^2 & k_2 & 1 \\ \vdots & \vdots & \vdots \\ k_n^2 & k_n & 1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad (3.29)$$

where \hat{x}_{k_i} is the estimated value of the correlation for delay k_i .

In the case of only three points (the maximum and one point on either side), the solution is unique. In this case, we can even derive a direct expression for the location of the peak by using the analytical expression for the inverse of a 3×3 matrix.

We first denote the delay associated with the maximum estimated correlation value as k_0 , and the delay values on either side as k_{-1} and k_{+1} , respectively. Then the expression for the location of the value that maximizes the fitted quadratic is

$$\hat{\tau} = \frac{\hat{x}_{k_{-1}} A + \hat{x}_{k_0} B + \hat{x}_{k_{+1}} C}{2(\hat{x}_{k_{-1}} D + \hat{x}_{k_0} E + \hat{x}_{k_{+1}} F)} \quad (3.30)$$

where

$$A = (k_0^2 - k_{+1}^2) \quad (3.31a)$$

$$B = (k_{+1}^2 - k_{-1}^2) \quad (3.31b)$$

$$C = (k_{-1}^2 - k_0^2) \quad (3.31c)$$

$$D = (k_0 - k_{+1}) \quad (3.31d)$$

$$E = (k_{+1} - k_{-1}) \quad (3.31e)$$

$$F = (k_{-1} - k_0) \quad (3.31f)$$

This is a non-linear expression, and so to determine the variance of $\hat{\tau}$ we must linearize this function by differentiating with respect to the state vector.

$$\mathbf{T} = \begin{bmatrix} \frac{\partial \tau}{\hat{x}_{k_{-1}}} \\ \frac{\partial \tau}{\hat{x}_{k_0}} \\ \frac{\partial \tau}{\hat{x}_{k_{+1}}} \end{bmatrix}^T = \begin{bmatrix} \frac{(AE-BD)\hat{x}_{k_0} + (AF-CD)\hat{x}_{k_{+1}}}{2(D\hat{x}_{k_{-1}} + E\hat{x}_{k_0} + F\hat{x}_{k_{+1}})^2} \\ \frac{(BD-AE)\hat{x}_{k_0} + (BF-CE)\hat{x}_{k_{+1}}}{2(D\hat{x}_{k_{-1}} + E\hat{x}_{k_0} + F\hat{x}_{k_{+1}})^2} \\ \frac{(CD-AF)\hat{x}_{k_0} + (CD-BF)\hat{x}_{k_{+1}}}{2(D\hat{x}_{k_{-1}} + E\hat{x}_{k_0} + F\hat{x}_{k_{+1}})^2} \end{bmatrix}^T \quad (3.32)$$

where A, B, C, D, E , and F are defined as in Equation (3.31).

With this matrix, the variance of $\hat{\tau}$ may be computed as follows:

$$\text{var}(\hat{\tau}) = \mathbf{T} \mathbf{P}_{-1,0,1} \mathbf{T}^T \quad (3.33)$$

where $\mathbf{P}_{-1,0,1}$ is the 3×3 sub-component of the state vector covariance matrix corresponding to the three points used for estimating the peak.

This formulation offers a few advantages. The primary advantage is its computational simplicity. No matrix inversion is required to estimate the location of the peak, and the variance may be directly estimated from the state vector covariance matrix. However, this simplicity comes at the cost of a limited capability of capturing the effects of non-linearity. Also, as formulated, this method allows only three points to be

used for fitting the peak. In many cases this is not an issue, however if the correlation vector to be estimated is very “flat”, (i.e. it does not have a very sharp peak) it may be advantageous to use more than three points to fit the quadratic to better estimate the true location of the peak, which is not achievable with this method as presented.¹

3.2.2 Unscented Transform

The other method of estimating the variance of the estimate $\hat{\tau}_{O,k}^A$ error from $\hat{\mathbf{x}}_k^+$ is to use the unscented transform. A brief overview of the unscented transform is presented below, as described in a number of standard textbooks [44, 39]. The unscented transform is applied as follows. First, a set of $2m$ sigma points are computed based on the covariance matrix, \mathbf{P}_k^+ . The sigma points are defined as follows:

$$\sigma_k^{(i)} = \begin{cases} \hat{\mathbf{x}}_k + (\sqrt{m} \hat{\mathbf{x}}_k) & i = 1, \dots, m \\ \hat{\mathbf{x}}_k - (\sqrt{m} \hat{\mathbf{x}}_k) & i = m + 1, \dots, 2m \end{cases} \quad (3.34)$$

The quadratic fit process described above is applied to each of these sigma points, and from each of these sigma points, an estimate of the peak location is produced. The average of these estimated peak locations is taken to be the final estimate of $\hat{\tau}_{O,k}^A$, while the variance of the estimated peak location is taken to be the variance of the final estimate.

The unscented approach has the advantage that it is capable of capturing the non-linear effects of the quadratic fit. There are a few places where non-linearities may be introduced. One is inherent in the fact that the fitting of a quadratic to a set of points is a non-linear operation. Even in the case of only three points to be fitted (in which case a unique solution exists), the expression for the location of the peak is non-linear in $\hat{\mathbf{x}}_k^+$.

Another source of non-linearities is the fact that the peak is fitted around the maximum correlation value. However, if the variance of the state vector is very large, then the location of the tap with the maximum value may be uncertain. If the variance of

¹It may be possible to overcome this limitation by fitting a quadratic function to more than three points using a standard least squares approach, and using the estimated variations of the quadratic parameters to then estimate the variance of the TDOA solution.

the state vector is large enough, it is possible for the difference in values of the correlation taps to be less than the standard deviation of the error of those estimates. In this case, even the point about which the quadratic is to be fitted is a variable to be estimated. The unscented transform captures these non-linearities, and is therefore particularly useful in cases when the state vector uncertainty is large. On the other hand, the unscented transform is computationally expensive.

3.3 Monte Carlo Simulations

In order to assess the performance of this algorithm when using a variety of pulsars commonly proposed as candidate signal sources for x-ray navigation, a series of Monte Carlo simulations was performed. The purpose of these simulations was to compare the performance of the algorithms proposed here with a theoretically derived performance bound, as well as with the performance of a maximum-likelihood TDOA estimator, discussed in [17].

The simulations were run as follows. In each simulation, a single TDOA was randomly selected. The TDOAs were drawn from a uniform distribution with a minimum value of 0 and a maximum value equal to the pulsar period. Photon arrival data was then generated using a pulse profile for the pulsar, delayed by the TDOA. These photon arrival times were then processed by the EKF algorithm proposed here, as well as a maximum likelihood estimator. Since the TDOA was not changing during the observation time, only the TDOA and not the frequency shift was estimated. This scenario is analogous to a spacecraft which is not moving and is simply estimating its position, rather than attempting to estimate position and velocity. Equivalently, this scenario is also analogous to an observer with perfect knowledge of their velocity. The goal of using this simplified scenario was to illustrate the performance of both algorithms under the same “ideal” conditions, and to compare their accuracies to a theoretically derived lower bound.

t_{obs} (s)	Background Flux ($\frac{erg}{cm^2s}$)	Detector FOV ($^\circ$)	Energy range (keV)	Simulations
1000	1.3E-13	0.1	2 - 10	100
1000	3.2E-12	0.5	2 - 10	100

Table 3.1: Monte Carlo Simulation Parameters

Name	Period (s)	Flux ($\frac{erg}{cm^2 s}$)	Pulsed Fraction (%)	Pulse Width (s)	Duty Cycle ($\frac{Pulse\ Width}{Period}$)
J0437-4715	0.00575	1.71E-12	27	8.79E-05	0.015279
B1937+21	0.00156	4.10E-13	86	3.90E-05	0.025022
B1957+20	0.00160	5.38E-13	60	4.65E-05	0.029057
B1821-24	0.00305	1.25E-12	98	5.18E-05	0.016978

Table 3.2: Pulsar Parameters [9, 45]

This process was repeated for a range of potential detector areas, detector field-of-views (FOVs), and flux densities (or pulsars), given in Table 3.1 and Table 3.2. Observation time was not varied, since for the case where the spacecraft remains motionless, increased observation time is theoretically equivalent to increased detector area. This implies that having a very large detector and a very short observation time is theoretically equivalent to a very small detector and a very long observation time. For each detector area/FOV scenario, four different pulsars, the parameters of which are given in Table 3.1 were evaluated. Each detector area/pulsar combination was run multiple times, and the error between the estimated TDOA and true TDOA was computed for each run. From these errors, a single standard deviation for each detector area/pulsar combination was computed. Details of the computations performed in these simulations are discussed below.

The detector areas in this simulation were chosen to represent a range of likely detector areas that could be flown onboard a spacecraft, ranging from $100cm^2$ (the area of one side of a 1U CubeSat) to $1m^2$. The two detector FOVs were chosen to represent two alternative cases. The wide FOV was chosen to represent an instrument with no spatial resolution, such as what might be achieved with a simple x-ray detector and collimator. This FOV is similar to the FOV on *Suzaku*'s Hard X-ray Detector (HXD) instrument, which is a non-imaging x-ray detector [46]. The narrow FOV case was chosen to represent a detector with spatial resolution, capable of determining the angle-of-arrival of the photons in question. Several modern x-ray telescopes have an angular resolution of much less than 0.1° . For instance, *SWIFT*'s x-ray telescope (XRT) has an angular resolution of 18 arcseconds, or 0.005° [47], so in theory such an instrument could achieve a much lower FOV and background noise. This is because, even though the FOV of the instrument is large, the angular resolution of the detector could be

used to reject all the photons except those that arrive from the area in the sky in the immediate vicinity of the pulsar, resulting in a lower effective FOV, assuming that the spacecraft has acceptably accurate knowledge of its orientation.

3.3.1 Theoretical Accuracy

The theoretical TDOA accuracy is expressed as follows [9]:

$$\sigma_{TOA} = \frac{W}{2\text{SNR}} \quad (3.35a)$$

$$\text{SNR} = \frac{\lambda^{(S)} A_D p_f t_{\text{obs}}}{\sqrt{[\lambda^{(\text{bkg})} + \lambda^{(S)}(1 - p_f)] \left(A_D t_{\text{obs}} \frac{W}{T} \right) + \lambda^{(S)} A_D p_f t_{\text{obs}}}} \quad (3.35b)$$

where

- t_{obs} is the observation time
- p_f is the fraction of the photons that are part of the “pulsed” signal component
- W is the pulse width (full width at half maximum) in seconds
- T is the period of the pulsar
- A_D is the area of the detector
- $\lambda^{(S)}$ is the average photon flux from the pulsar
- $\lambda^{(\text{bkg})}$ is the average photon flux from background noise

The expression above gives the expected accuracy to which a single TDOA could be computed, given an observation interval t_{obs} . The expression does not take into account the changing position of the observer during the observation, or that the velocity of the observer may not be perfectly known. Therefore, in order to perform a meaningful comparison between the estimation techniques discussed here and the theoretical accuracy, the zero-velocity case was chosen.

Equation 3.35 contains two terms for background photon flux. The first term is the product of the photon flux from the pulsar with the “unpulsed fraction,” or $\lambda^{(S)}(1 - p_f)$. The pulsed fraction for the pulsars used in these simulations were taken from [9, 45].

The second background term is $\lambda^{(\text{bkg})}$ which refers to photons coming from sources other than the pulsar, particularly the cosmic x-ray background.

X-ray background flux is a function of energy of the incoming photons. The intensity of this flux is approximated by the following equation [48]:

$$\lambda^{(\text{bkg})}(E) = \begin{cases} 7.877E^{-0.29}e^{-E/41.13} & 3 < E \leq 60\text{keV} \\ 0.0259\left(\frac{E}{60}\right)^{-5.5} + \\ 0.504\left(\frac{E}{60}\right)^{-1/58} + \\ 0.0288\left(\frac{E}{60}\right)^{-1.05} & E > 60\text{keV} \end{cases} \frac{\text{keV}}{(\text{keV})(\text{cm}^2)(\text{s})(\text{sr})} \quad (3.36)$$

where E is the energy of the background x-ray photons.

In order to determine the background flux for a given detector, Equation (3.36) was numerically integrated over the expected energy range and solid angle of the detector to give the energy of the incoming flux per unit area over the detector's energy range.

3.3.2 ML Estimator

The MLE used in these simulations was a variation of the estimator presented in [17] and referred to in that work as an “approximate MLE.” The estimator computes the maximum likelihood estimate of phase offset and frequency shift given an observation of N photons. The estimator is expressed in its full form as follows:

$$(\hat{\tau}_{\text{ml}}, \hat{f}) = \underset{\tau, f}{\operatorname{argmax}} \sum_{k=1}^N \log \lambda_O \left(\tilde{\phi}(T_k) + \tau + f(T_k - t_a) \right) \quad (3.37)$$

where

- T_k is the arrival time of the k^{th} photon
- $\hat{\tau}_{\text{ml}}$ is the resulting ML estimate of time offset

- \hat{f} is the resulting ML estimate of frequency shift
- $\tilde{\phi}(T_k)$ is the predicted phase, based on an initial estimate of the kinematic state of the vehicle from the navigation filter

The estimator is called an approximate MLE because it approximates the true likelihood function based on the assumption of many periods of the signal being observed. The approximated likelihood function results in a function which is more convenient for computational purposes. Since the case treated here is a simplified one, in which the observer's position does not change during the observation time, the estimator may be simplified even further, since there is no need to estimate the frequency shift of the signal. Under these conditions, the estimator simply becomes:

$$\hat{\tau}_{\text{ml}} = \operatorname{argmax}_{\tau} \sum_{k=1}^N \log \lambda_O(T_k + \tau) \quad (3.38)$$

In these simulations, the ML estimate of TDOA was found using a simple grided search algorithm over possible TDOAs from zero to the pulsar period.

3.3.3 TDOA Kalman Filter

For the EKF developed in Section 3.1, the dimension of the state vector was chosen adaptively for each pulsar/detector area combination. Recalling that the state vector is a sampled version of the correlation vector, and assuming that the samples span one period of the correlation function, then the time resolution of the sampled state vector is directly dependent on the number of states in the state vector. Specifically, the time resolution of the state vector will be the period of the pulsar divided by the number of states in the state vector.

$$T = \frac{T}{m} \quad (3.39)$$

Theoretically, the EKF presented here is capable of achieving TDOA estimates of a higher accuracy than the time-step of the filter, since it uses a quadratic best fit approach

to find the location of the maximum value. However, in practice the quadratic fit is only an approximation of the true peak of the correlation function. As the region over which the quadratic approximation is being applied grows, the approximation becomes less accurate. Since a minimum of three taps are required to formulate the quadratic best-fit, the minimum length over which the approximation is applied is $2T$. Therefore, if the time-step is very large, the quadratic approximation must also be applied to a large region, which results in errors in the estimation of TDOA.

The exact magnitude of the error associated with this approximation has not been analyzed in this dissertation. Furthermore, the error associated with this approximation would be specific to each pulsar, since some correlation functions might resemble the approximating quadratic function more closely than others. However, as will be discussed in Section 3.5.1, the effects of the error associated with the quadratic approximation may be dramatically reduced by centering the correlation vector estimate around the peak of the correlation vector.

It is notable that while in these simulations a fixed time-step between taps in the state vector was used, it need not be so. In future work, a coarse-resolution state vector could be used initially to estimate the approximate location of the peak. After the coarse peak location was estimated, the region near the peak could be estimated using a filter with a smaller time-step. This process could result in significant reductions in the number of states in the state vector and thus the number of computations per time-step.

3.3.4 Monte Carlo Results

The results of the Monte Carlo analysis are shown in Figure 3.5 and Figure 3.6. Figure 3.5 shows the results of the MC analysis performed using a narrow FOV detector (resulting in decreased background flux), while Figure 3.6 shows the results using a wide FOV detector. Each sub-plot displays the achieved and theoretical standard deviations as a function of detector area for a given pulsar. The solid lines with circular markers indicate the TDOA error standard deviation achieved by the EKF described in this dissertation, while the square markers indicate the TDOA error standard deviation achieved by the MLE. The theoretical accuracy predicted by Equation (3.35) is shown with dotted lines. Finally, the standard deviation of a uniform distribution, $\mathcal{U}[0, T]$ is

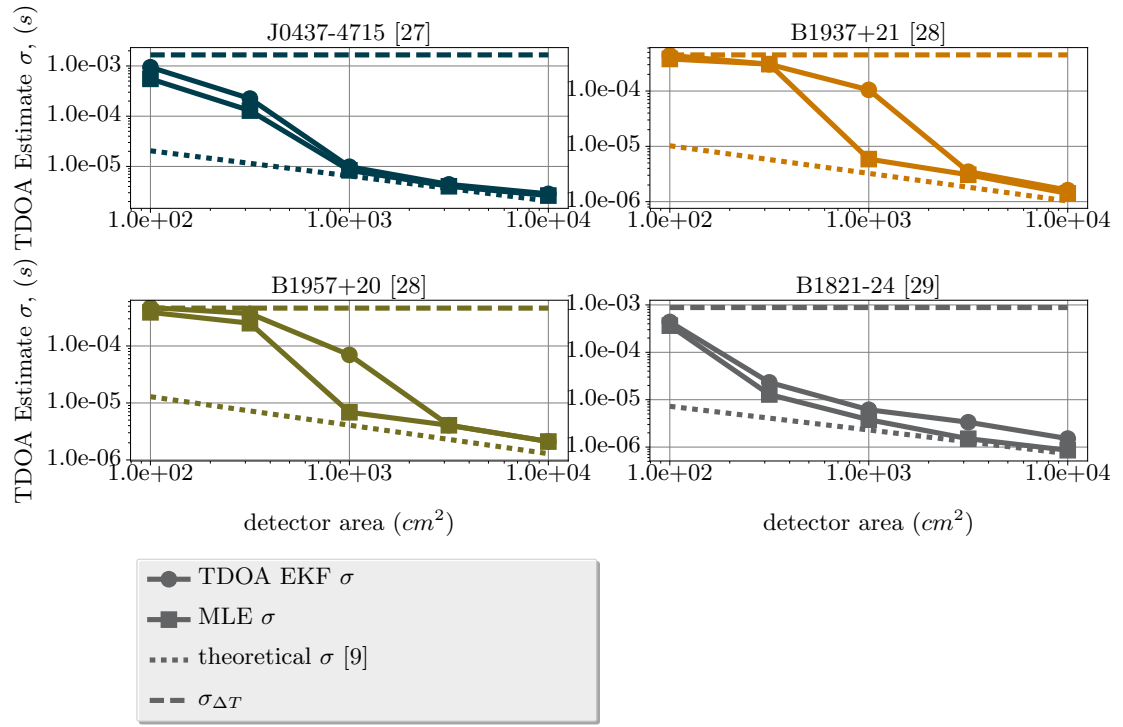


Figure 3.5: Monte Carlo Error Analysis, FOV= 0.1°

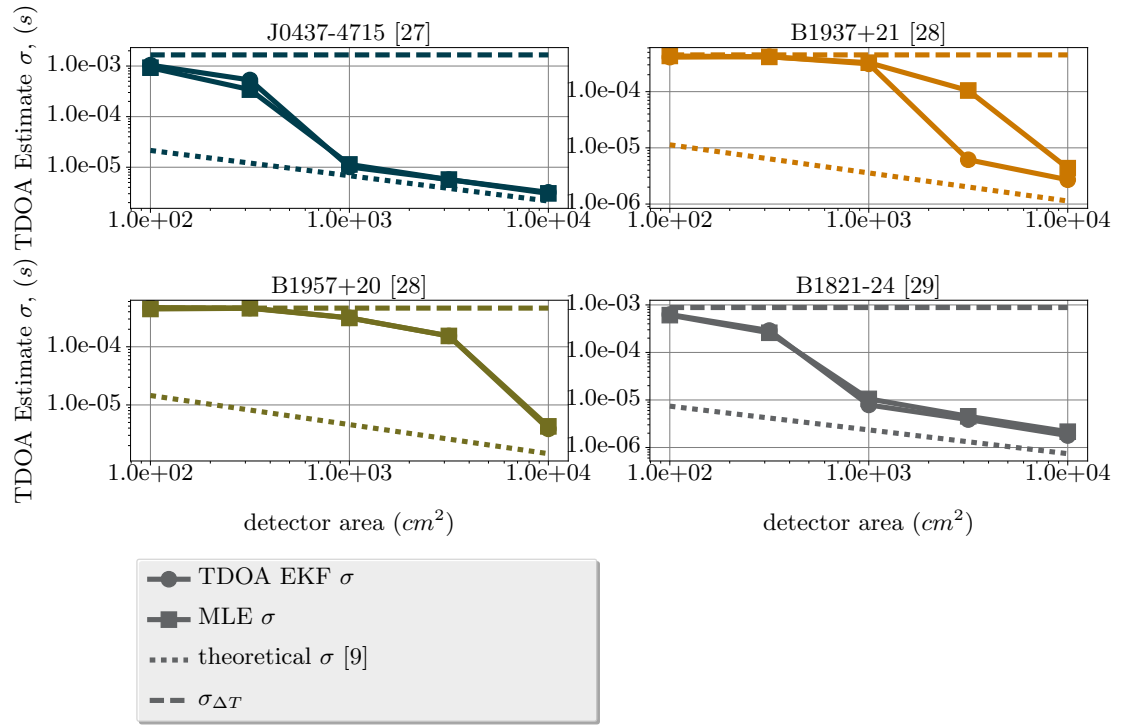


Figure 3.6: Monte Carlo Error Analysis, FOV= 0.5°

shown in the dashed line. This is the distribution from which the TDOAs were originally drawn. This line is shown to indicate the accuracy that could be achieved by guessing “at random.”

A summary of the Monte Carlo simulation parameters, and properties of the pulsars used in this study are shown in Tables 3.1 and 3.2. The pulse profiles for the pulsars used in this simulation are shown in Figure 2.1.

There are a few points worth noting from these results. The first is that both the MLE estimator and the EKF estimator fail to achieve the level of accuracy predicted by Equation 3.35 at low detector areas, particularly in the large FOV case. In fact, for some detector area/FOV/pulsar combinations, the accuracy achieved both by the EKF and the MLE is no better than the standard deviation of a uniform distribution, indicating that for these combinations, the estimation methods shown here do not produce a meaningful estimate of TDOA. The discrepancy between the theoretical error and actual error may indicate that this theoretical accuracy is not a good lower bound in scenarios where the expected number of photons measured by the detector is very small. However, at larger detector areas, both the MLE and EKF estimators tend towards the theoretically achievable accuracies.

Secondly, while there is some difference in performance between the EKF estimator and the MLE estimator, in general both methods achieve similar performance in most detector area/pulsar combinations. Since this study was performed assuming highly idealized conditions, it does not necessarily indicate that the performance of the two methods would be similarly matched in real conditions, for instance when the spacecraft is moving at some unknown velocity. Further Monte Carlo analysis and analysis of real data to compare the two methods under a variety of operating conditions would serve to reveal the performance differences between the two algorithms under more realistic conditions. However, these results indicate that the algorithm described in this dissertation could provide the same level of accuracy as the MLE method, while offering the benefit of incorporating new information from each photon arrival into the navigation solution recursively.




Figure 3.7: *Suzaku*'s position relative to the origin

3.4 Validation Experiment

In order to demonstrate that these algorithms function with more realistic real data, the algorithm was exercised in post-process using photon arrival data from the Japanese Space Agency (JAXA) *Suzaku* probe's observations of the Crab Nebula Pulsar (PSR B0531+21). The photon arrival data was used along with the known characteristics of the pulsar to estimate the TDOA of the signal, and from there the spacecraft's range from the pulsar along the line-of-sight vector was computed. In order to assess the accuracy of the computed range, the actual position of the spacecraft was also computed based on the reported position of the spacecraft relative to Earth, and the position of Earth relative to the SSB. This reported position was also projected onto the line of sight vector to the pulsar, to give the true range. The path of the spacecraft and the range being estimated are illustrated in Figure 3.7.

The photon arrival data was obtained using the Hard X-Ray Detector (HXD) instrument aboard *Suzaku*. This detector provides an interesting validation case because it was similar in area, timing resolution, and energy range to detectors being developed as navigational sensors for CubeSat applications at the University of Minnesota, Twin Cities campus [49, 50]. The specifications of the HXD instrument are available in [51]. The photon arrivals were screened using basic criteria described in [52]. These criteria were used to screen out photons based on event grade and time interval of arrival (e.g. photons that arrived during maneuvers, or while the Earth was blocking the field of view were screened), but did not filter out background photon arrivals.

An estimate of the spacecraft’s velocity is required for the time-update of the state vector. In practice, a high-fidelity orbit model would be used to estimate the vehicle’s velocity. However, for this work, a simple one-dimensional navigation filter was used, in which no prior knowledge of the spacecraft’s trajectory was used. Based on the initial estimates of TDOA, an initial velocity was estimated, and then this velocity was used to provide the velocity feedback in the time-update of the correlation vector. This is somewhat similar to potential “cold-start” operating scenarios, in which the initial position and velocity of the spacecraft are unknown.

The results of several different orbits are shown in Figure 3.8. Each panel in the plot shows the error between the estimated change in range and the true change in range. Additionally, the $\pm\sigma$ estimated standard deviations are plotted, which show the errors to be reasonably well-bounded by the estimated standard deviations.

It is important to note that the results shown here may not be typical, or even achievable for some pulsars at some detector sizes. The Crab Nebula pulsar has significantly higher x-ray photon flux than other pulsars considered for x-ray pulsar navigation, and so the results achieved here are of higher accuracy than what could be expected using other signal sources. Nevertheless, the results shown here illustrate the ability of the algorithm proposed here to estimate velocity in real time without any prior orbit or trajectory knowledge, and to use that estimate of velocity to improve the final position estimate.

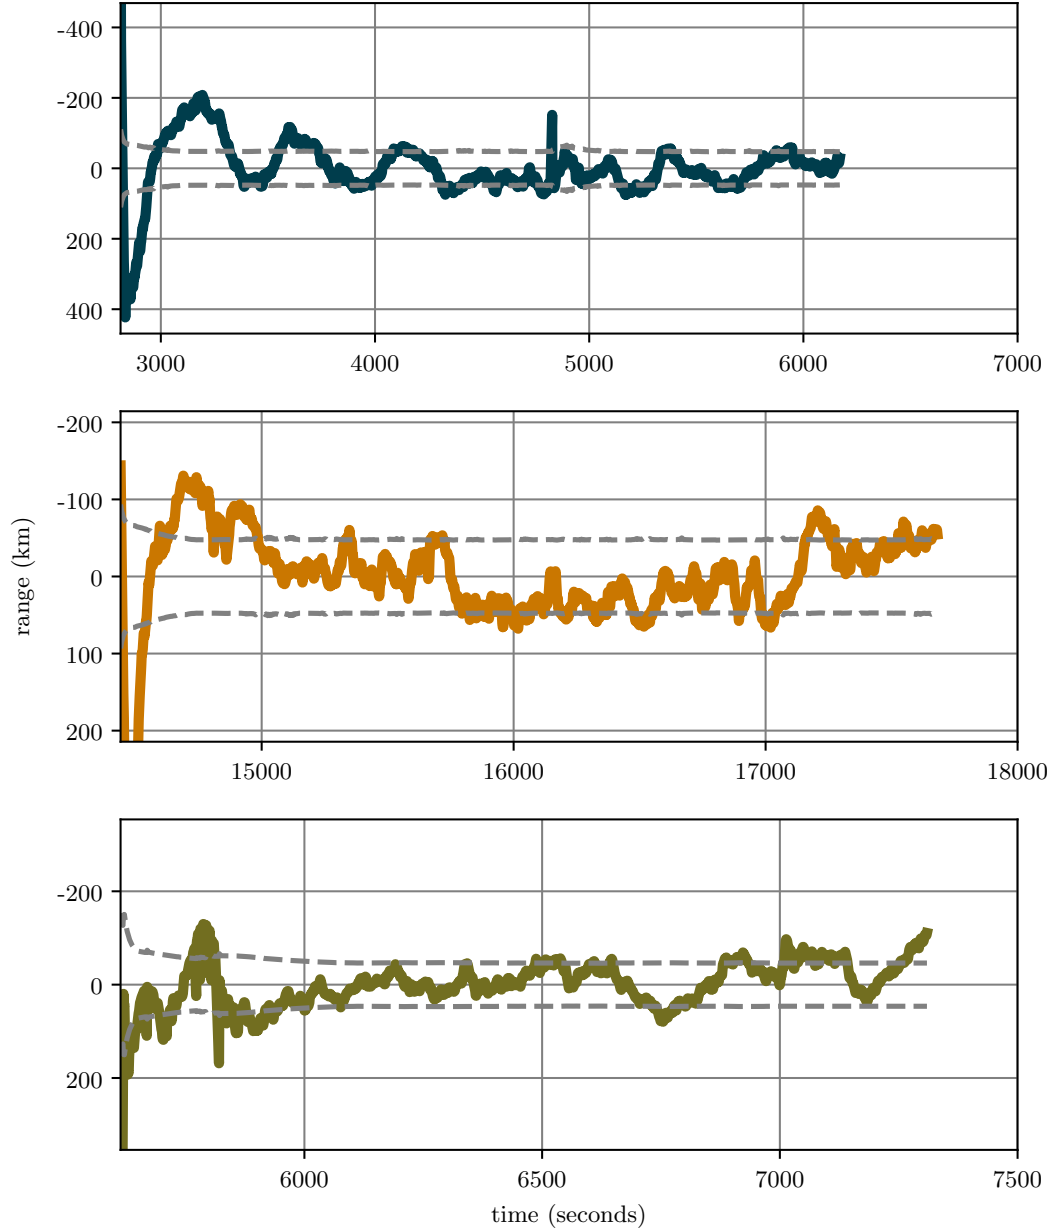


Figure 3.8: Range errors for four different time segments of the *Suzaku* mission. Observation IDs: ae100023020 ($\sigma = 80km$), ae100023020 ($\sigma = 64km$), ae101003010 ($\sigma = 54km$)

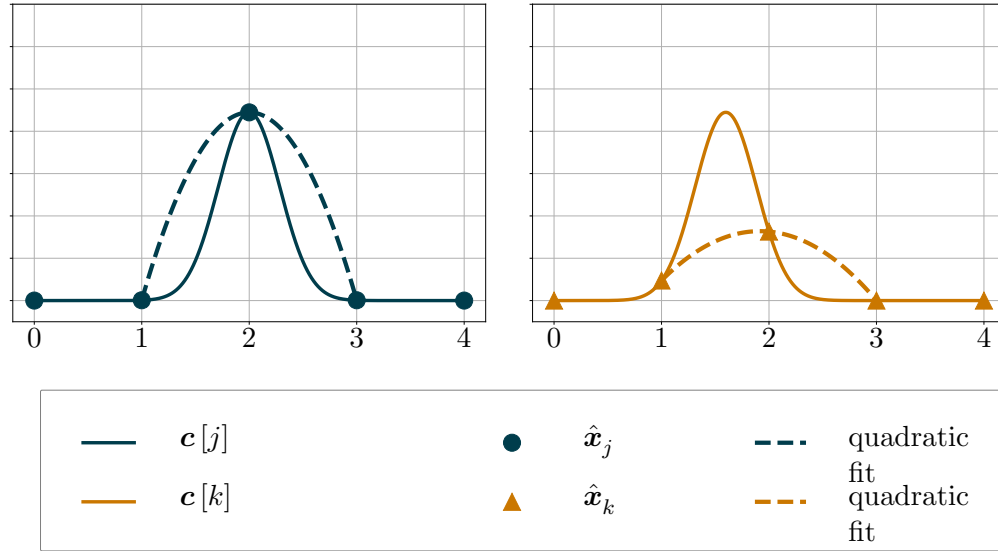


Figure 3.9: Good and bad quadratic fits

3.5 Implementation Details

In the previous sections, the theoretical framework of this estimator was derived. While these equations are sufficient for reproducing the results shown in Section 3.3 and Section 3.4, there are specific details of implementation and “lessons learned” which may aide future research efforts intending to make use of these results. These findings are summarized here.

3.5.1 Peak Centering

The algorithms presented here rely on the use of a quadratic fit the the maximum correlation value and a few surrounding points to determine the true location of the maximum correlation value. The quadratic fit is only an approximation to the true shape of the peak of the correlation function. As such, the accuracy of this approximation may vary based on a number of factors. One factor which has been determined to have a large impact on the accuracy of the approximation is how asymmetric the quadratic fit is about the max point (Figure 3.9).

It may be seen by inspection that the quality of the fit depends on how close the location of the peak correlation value is to the location of the true correlation function

peak. If the peak correlation value lies exactly on the peak, then the maximum of the quadratic will be exactly equal to the true maximum of the correlation vector. If on the other hand the peak of the correlation function is far from the peak correlation value, the quadratic may do a poor job estimating the peak. Consequently, it is desirable to keep the peak “centered” about a correlation tap.

Of course, it is not possible to know the exact location of the true peak. However, it is possible to center the peak based upon the current estimate of TDOA. In fact, the time-update matrix is designed exactly to shift the correlation vector a specified amount. In Section 3.1.1, the time-update matrix was used to shift the correlation vector by an amount determined by the estimated velocity. However, the same matrix can also be used to shift the correlation vector forwards or backwards to align the peak value of the correlation vector with the estimated location of TDOA. This procedure is outlined below.

First, the amount by which the vector needs to be shifted is computed by taking the difference in the delay corresponding to the maximum correlation vector value and the previous estimate of TDOA. Mathematically, if the maximum value of the correlation vector is at index p , then the shift is given by:

$$\delta = \hat{\tau} - (pT + \hat{d}) \quad (3.40)$$

This is the amount by which the correlation vector estimate and covariance matrix need to shift in order to align with the peak. However, it is important to note that we are not intending to move the correlation vector with respect to the x axis. Consequently, we must also update the offset value, \hat{d} to reflect the new time-difference values to which the state vector values now correspond.

This centering step may be performed as a separate, additional step after the measurement update. Alternatively, if the alternate time update formulation is used (Section 3.1.2), then this step may be performed during the time update. Recall that in the alternate time update formulation, the time update matrix is equal to the identity matrix. To incorporate the re-centering, we may simply substitute the identity matrix with the peak-centering matrix, while still using the process noise matrix associated




Figure 3.10: External one-dimensional navigation filter

with velocity.

3.5.2 Integrated Navigation Filter

To this point, nothing has been said about the estimation of velocity, other than that it will be estimated by some external navigation filter with knowledge of the system dynamics and multiple inputs. That velocity is then fed back into the correlation vector estimate during the time update.

However, it is also possible to estimate the velocity of the spacecraft along the line of sight to the pulsar without any other external inputs or knowledge of the spacecraft's acceleration. This is particularly useful in the case where initial position is not known, or is known only very poorly. Two methods for simultaneously estimating the spacecraft's dynamic state are outlined below.

External Navigation Filter

The simplest approach for estimating the spacecraft's dynamic state is to implement a basic one dimensional position and navigation filter that takes as its measurement input the output of the TDOA estimator. This method is similar to the architecture proposed in Figure 3.1, except instead of estimating a full dynamic state of the spacecraft (three dimensions of position, velocity, attitude, etc.), only one dimension of position and velocity are estimated (Figure 3.10).

This implementation has the advantage of relative simplicity. Each component may be treated individual and there is a degree of separation between the estimators which makes the practical implementation somewhat simpler. However, this method suffers

because it fails to account for the correlation between error in TDOA estimate and position and velocity estimate. This can lead to “feed-back” errors in which the estimated variance of TDOA and velocity are far less than the true covariance. In the worst case scenario this feed-back can cause the filter to diverge completely.

A simple “fix” to the issue of un-estimated error correlation is to simply inflate the covariance of the TDOA and velocity estimates. This is a somewhat *ad-hoc* approach which requires significant time to “tune” the inflation values, but once tuned generally produces reasonable results. The external navigation filter method was used in generating the results of Section 3.4.

Joint Navigation Filter

The other alternative method of estimating velocity (and higher-order kinematic states, if desired) is to include those kinematic states in the state vector itself. The new state vector then becomes:

$$\mathbf{x} = \begin{bmatrix} \hat{\mathbf{c}}_{\lambda, \lambda_O} [d] \\ \hat{v} \\ \hat{a} \end{bmatrix} = \begin{bmatrix} c_{\lambda, \lambda_O} [d + T] \\ c_{\lambda, \lambda_O} [d + 2T] \\ \vdots \\ c_{\lambda, \lambda_O} [d + (m - 1)T] \\ \hat{v} \\ \hat{a} \end{bmatrix} \quad (3.41)$$

where $\hat{\mathbf{c}}_{\lambda, \lambda_O} [d]$ is the estimated correlation vector that has been used throughout this chapter, $\hat{v} = \mathbf{v} \cdot \boldsymbol{\ell}$ is the velocity along the line of sight to the pulsar, and $\hat{a} = \mathbf{a} \cdot \boldsymbol{\ell}$ is the acceleration along the line of sight to the pulsar. Note that the number of kinematic states to be estimated is flexible and up to the user. In space, where acceleration is a smoothly varying function of position, it makes sense to estimate acceleration in addition to velocity, however this may not make sense in all applications.

This formulation has the distinct advantage of jointly estimating correlation with the other kinematic states, which allows the correlation between the error estimates to also be estimated. This eliminates the error feedback problem associated with the

external navigation filter. However, this comes at the expense of a larger covariance matrix which is more computationally expensive.

In order to estimate this joint vector, a few modifications must be made to the filtering equations, in particular the time update equations. Recall from Section 3.1.1 that the time-updated state vector is a function of the *a posteriori* state vector and the estimated “peak shift.”

$$\hat{\mathbf{x}}_k^- = f(\hat{\mathbf{x}}_j^+, \hat{\delta}_{j \rightarrow k}, w_{j,\delta}) = F_{j \rightarrow k}(\hat{\delta}_{j \rightarrow k}, w_{j,\delta}) \hat{\mathbf{x}}_j^- \quad (3.42)$$

However, we also note that the “peak shift” itself is a function of the kinematic states of the spacecraft. Under the new formulation, the kinematics are included in the state vector, so the time-update must be reformulated to reflect this relationship.

We begin by noting that the time-updated value of the state vector is given by:

$$\hat{\mathbf{x}}_k^- = \begin{bmatrix} \sum_{i=0}^m \text{sinc}(\delta) \hat{x}_{j,i}^+ \\ \sum_{i=0}^m \text{sinc}(\delta - 1) \hat{x}_{j,i}^+ \\ \vdots \\ \sum_{i=0}^m \text{sinc}(\delta - m + 1) \hat{x}_{j,i}^+ \\ \hat{v}_j + \hat{a}_j (k - j)T \\ \hat{a}_j \end{bmatrix} \quad (3.43)$$

However, as stated above, the “peak shift” term δ is now itself a function of the state vector. Recalling Equation (3.5), the peak shift term is written as

$$\hat{\delta}_{k \rightarrow j} = \frac{\hat{\mathbf{v}}_{O,k}^A \cdot \boldsymbol{\ell}^{(S)}(k - j)}{c} \quad (3.44)$$

Now however, we are estimating acceleration as well as velocity, so this term must be included. We also note that we are estimating a single dimension of velocity along the line-of-sight to the pulsar, so the dot product with the line-of-sight vector may be dropped. Finally, for simplicity, if we estimate velocity in units of fractional speed of light, then the division by the speed of light may be dropped. Therefore, the new peak

shift may be written as:

$$\hat{\delta}_{k \rightarrow j} = \hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} \quad (3.45)$$

Substituting this expression into the expression for the time-updated state vector yields the following:

$$\hat{\mathbf{x}}_k^- = f(\hat{\mathbf{x}}_j^+) = \quad (3.46)$$

$$\begin{bmatrix} \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} \right) \hat{x}_{j,i}^+ \\ \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} - 1 \right) \hat{x}_{j,i}^+ \\ \vdots \\ \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} - m + 1 \right) \hat{x}_{j,i}^+ \\ \hat{v}_j + \hat{a}_j (k-j)T \\ \hat{a}_j \end{bmatrix} \quad (3.47)$$

To derive the time-update matrices, we must take the Jacobian of the time-update function. This yields the following time-update matrix:

$$F_{j \rightarrow k}(\hat{\mathbf{x}}_j^+) = \quad (3.48)$$

$$\begin{bmatrix} \text{sinc}(\delta) & \text{sinc}(\delta+1) & \dots & \text{sinc}(\delta+m-1) & \epsilon_0(k-j) & \epsilon_0 \frac{(k-j)^2}{2} \\ \text{sinc}(\delta-1) & \text{sinc}(\delta) & \dots & \text{sinc}(\delta+m-2) & \epsilon_1(k-j) & \epsilon_1 \frac{(k-j)^2}{2} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \text{sinc}(\delta-m+1) & \text{sinc}(\delta-m+2) & \dots & \text{sinc}(\delta) & \epsilon_{m-1}(k-j) & \epsilon_{m-1} \frac{(k-j)^2}{2} \\ 0 & 0 & \dots & 0 & 1 & (k-j)T \\ 0 & 0 & \dots & 0 & 0 & 1 \end{bmatrix} \quad (3.49)$$

where

$$\epsilon_i = \sum_{i=0}^m \frac{(\delta + i - i) \cos[\delta + i - i] + \sin[\delta + i - i]}{(\delta + i - i)^2} \hat{x}_{j,i}^+ \quad (3.50)$$

This time-update equation bears similarities to the previously derived time update equations. The block corresponding to the correlation vector is in fact identical. The column relating the correlation vector to the velocity is in fact the same as the process noise mapping matrix in Equation (3.7). The column relating acceleration to correlation includes a similar term for the derivative of the *sinc* function, but with the double integral of acceleration rather than the single integral for velocity.

We now turn to the derivation of the new process noise mapping matrix. Now, instead of process noise being noise in velocity, the new process noise will be noise in the derivative of the highest order kinematic state, in this case acceleration.

$$\hat{\mathbf{x}}_k^- = f(\hat{\mathbf{x}}_j^+, w_{\text{acceleration}}) = \quad (3.51)$$

$$\begin{bmatrix} \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} + w \frac{(k-j)^3}{6} \right) \hat{x}_{j,i}^+ \\ \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} + w \frac{(k-j)^3}{6} - 1 \right) \hat{x}_{j,i}^+ \\ \vdots \\ \sum_{i=0}^m \text{sinc} \left(\hat{v}_{O,j}^A (k-j) + \hat{a}_{O,j}^A \frac{(k-j)^2}{2} + w \frac{(k-j)^3}{6} - m + 1 \right) \hat{x}_{j,i}^+ \\ \hat{v}_j + \hat{a}_j (k-j)T + w \frac{(k-j)^2 T^2}{2} \\ \hat{a}_j + w (k-j)T \end{bmatrix} \quad (3.52)$$

Differentiating this expression with respect to the process noise yields the following process noise mapping matrix:

$$\mathbf{L}_j = \begin{bmatrix} \epsilon_0 \frac{(k-j)^3}{6} \\ \epsilon_1 \frac{(k-j)^3}{6} \\ \vdots \\ \epsilon_{m-1} \frac{(k-j)^3}{6} \\ \frac{(k-j)^2 T^2}{2} \\ (k-j)^2 T \end{bmatrix} \quad (3.53)$$

The process noise variance, in this case, should represent how quickly in time the spacecraft’s acceleration is expected to change. Of course, this is a function of position and velocity, so determining an absolute value may not be possible especially in a “lost-in-space” scenario. However, it is possible to come up with some fairly conservative upper bounds on the rate of change of acceleration.

If we assume that the spacecraft is somewhere within our solar system, then we can easily compute the steepest possible gravity gradient based on the masses and radii of bodies in the solar system. We may also make assumptions about the maximum speed of the spacecraft: the spacecraft cannot exceed the speed of light, but most likely will not approach even one percent of that speed (the current record-holder for fastest man-made object is the Parker Solar Probe, which will travel at 0.064% of the speed of light). From these very reasonable assumptions, it is relatively straight-forward to compute a maximum bound on the time-rate of change of acceleration, from which a variance may be approximated. If more is known about the spacecraft’s trajectory and the range of possible accelerations and speeds, the value used for the variance of the time-rate of change of acceleration may be further refined, resulting in higher accuracy estimates.

The time-update matrices presented here do not use the alternative time-update technique presented in Section 3.1.2. However, this may be achieved by following a similar technique to that presented in Section 3.1.2.

The measurement update for the deeply integrated kinematics state vector is the same as that presented in previous sections; the only difference is that this vector will have additional zeros corresponding to the fact that the kinematic states are not directly measured.

Nomenclature

Correlation Vector Estimation

d	Amount by which the correlation vector time indices are offset from zero
$c_{A(t),B(t)}[p]$	The correlation between signals $A(t)$ and $B(t+p)$
n	Dimension (number of taps) in the correlation vector
p	Index to denote a specific correlation vector or FIR filter tap
\mathbf{w}_λ	Ideal FIR filter for estimating signal λ from noisy measurement of signal

Data Association

bkg	Indicates background photon flux
p	Indices to designate signal source

Detector

A_D	Detector area (cm^2)
-------	--------------------------

Physical Constants

c	Speed of light (km/s)
-----	---------------------------

Probability and Statistics

$\text{var}(a)$	The variance of random variable a
$\mathbb{E}[a]$	The expected value of random variable a
$\mathcal{N}[\mu, P]$	A normal distribution with mean μ and variance P
$\mathcal{U}[l, u]$	A uniform distribution with lower bound l and upper bound u

Signals and Measurements

t_{obs}	Total observation time (<i>seconds</i>)
p_f	Fraction of flux from pulsar that is pulsed
W	Full width at half maximum of primary pulsar pulse
$\tilde{\Lambda}_{A,k}$	The number of photons measured at location A over time-step k (<i>Photons</i>)
$\tilde{\Lambda}_{dT=k,f=A}$	The number of photons measured at location A over time-step k (<i>Photons</i>)
$\lambda_k^{(p)}, \lambda^{(p)}(t)$	Expected number of photons at time index k , or expected photon flux at time t , from signal source p (<i>photons, photons/second</i>)
T_p	The period of pulsar p (<i>seconds</i>)
D	Integer delay
S^i	i^{th} signal source
$\text{sinc}(j)$	The sinc function

State Estimation

$\rho_t^{A(S)}$	Range to object A, at time t, along line-of-sight vector to source S (<i>km</i>)
\mathbf{r}_t^A	Position of A, at time t (<i>km</i>)
\mathbf{a}_t^A	Acceleration of object A, at time t (<i>km/s²</i>)
\mathbf{q}_t^A	Attitude of object A, at time t (<i>radians</i>)
$\boldsymbol{\omega}_t^A$	Angular velocity of object A at time t (<i>radians/second</i>)
O	Denotes the navigation frame origin
j, k	Time-step indices
t	Current time (<i>seconds</i>)
$\tau_t^{A(S)}$	Time-difference of arrival of signal from signal source S measured between vehicle A and navigation frame origin (<i>seconds</i>)
m	Dimension of state vector
\mathbf{x}	State vector
\mathbf{P}	Covariance of state vector
$\mathbf{y}_{t,a}$	Measurement of scalar quantity a
$v_{t,a}$	Scalar measurement error of measured quantity a at time t
$\delta_{k,j \rightarrow k}$	Peak shift from time j to k
$\mathbf{R}_{t,a}$	Variance of measurement noise in measured quantity a
$\boldsymbol{\delta y}_{t,a}$	Measurement residual for measured scalar quantity a at time t
$\mathbf{S}_{t,a}$	Variance of measurement residual for measured quantity a at time t
\mathbf{w}_t	Error in time update (process noise) at time t
\mathbf{Q}_t	Variance of process noise at time t
n	Length of time between filter time-steps (<i>seconds</i>)
$[\cdot]^-$	Indicates <i>a priori</i> value
$[\cdot]^+$	Indicates <i>a posteriori</i> value
$\hat{\cdot}$	Indicates estimated value
\mathbf{a}	Indicates vector quantity
a	Indicates scalar quantity

Misc $\mathbf{I}_{n \times n}$ The $n \times n$ identity matrix i

Iterator to designate summation index

 A, B

Indices to designate spacecraft(s)

 $\ell^{(S^i)}$ Unit vector pointing towards i^{th} signal source

Chapter 4

Attitude Estimation

The ability of a spacecraft to determine its attitude is critical for the implementation of an XNAV navigation scheme, as will be shown in Chapter 5. Additionally, attitude determination is important in its own right. The ability of a spacecraft to determine its attitude is critical for many aspects of a space mission.

Star trackers present an ubiquitous and affordable technique for attitude determination [53, 54, 55, 56]. Star trackers function by using optical imagery of visible stars to identify the relative bearing angle from the spacecraft to known stars. Using the angles to multiple known stars, the attitude of the spacecraft may be determined. In theory only two known stars are needed, but in practice three stars are needed for the purpose of uniquely identifying the stars by their relative positions to one another.

While star trackers are both accurate and readily available, they also suffer from limitations. The primary limitation of star trackers is that they cannot function in environments where the star tracker cannot identify guide stars. This can occur when the spacecraft is operating in proximity to objects which block or reflect light. For instance, the star trackers on *Cassini* were suspended for up to five hours when Saturn, the rings, or other bright bodies entered the star-tracker field of view [57, 58]. Other environments in which a star tracker might not function include extraterrestrial volcanic plumes, comet tails, and during landing phases in which dust is stirred up. The Apollo missions were unable to use their star trackers during waste-dumps due to the reflection of light off of waste “fooling” the star trackers [59].

While optical star-tracking suffers from the limitation described here, star-tracking

need not be performed in the optical wavelengths. Many stars are bright in soft and hard x-ray wavelengths. The same principles applied to optical star trackers could theoretically be applied to stars that are x-ray bright. A spacecraft could in theory determine its attitude using angle-of-arrival measurements from x-ray stars.

Such an attitude determination system could potentially function in environments where star-trackers fail, since many of the sources of interference discussed above would not emit x-rays. Furthermore, an x-ray detector capable of measuring both time and angle of arrival of x-ray photons could theoretically be used to determine both position and attitude of a spacecraft, offering a complete, six-degree of freedom navigation solution from a single instrument.

This chapter is devoted to the derivation of attitude estimation algorithms which could be used to estimate attitude from x-ray photon angle of arrival measurements. We assume that the spacecraft is equipped with an inertial measurement unit (IMU) whose angular rate measurements may be numerically integrated to yield an attitude solution. Attitude errors due to gyro noise and biases are corrected by periodic photon angle-of-arrival measurements. In this chapter, we assume for the sake of simplicity that we are able to uniquely associate incoming photons with their signal source of origin. Of course in reality this is not the case; the problem of data association will be addressed in detail in Chapter 5. The purpose of this chapter is to present the estimation algorithms that would be used if one could properly associate photons with their source of origin, and then those algorithms will be used in the Chapter 5 to derive the joint data association and estimation algorithm.

4.1 Quaternions and Spatial Rotation

In this work, attitude is parameterized using the attitude quaternion. For completeness, an overview of how a quaternion may be used to represent rotation and its properties is given here.

We begin by noting that, in general, a three-dimensional rotation from one frame to another may be parameterized by a unit vector axis of rotation, $\mathbf{v} = \begin{bmatrix} v_x & v_y & v_z \end{bmatrix}^T$ and an angle of rotation about that vector, θ (Figure 4.1).

A rotation may be completely characterized by this unit vector and the rotation




Figure 4.1: Rotation from navigation frame N to spacecraft body frame B by axis-angle rotation

about the unit vector. While simply storing the vector and angle (i.e. $\begin{bmatrix} \mathbf{v} & \theta \end{bmatrix}^T$) offers relative simplicity, it is advantageous from a computational standpoint to store the vector in a slightly different form. This form is defined as follows:

$$\mathbf{q} = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} = \begin{bmatrix} \cos\left(\frac{\theta}{2}\right) \\ v_x \sin\left(\frac{\theta}{2}\right) \\ v_y \sin\left(\frac{\theta}{2}\right) \\ v_z \sin\left(\frac{\theta}{2}\right) \end{bmatrix} \quad (4.1)$$

This is the quaternion representation for spatial rotation.¹ While it is not immediately obvious why it would be advantageous to parameterize rotation in such a non-intuitive format, the quaternion has a number of beneficial properties over other formulations of attitude such as Euler angles.

4.1.1 Composition of Spatial Rotations

Two sequential rotations may be composed to form a single rotation. If two rotations are represented by quaternions (\mathbf{p} and \mathbf{q}) the rotation resulting from first rotating by \mathbf{p} and then by \mathbf{q} is given by:

$$\mathbf{w} = \mathbf{p} \otimes \mathbf{q} \quad (4.2)$$

where \otimes is the Hamilton product. The Hamilton product is defined as follows:

$$\mathbf{p} \otimes \mathbf{q} = \begin{bmatrix} p_1 q_1 - p_2 q_2 - p_3 q_3 - p_4 q_4 \\ p_1 q_2 + p_2 q_1 + p_3 q_4 - p_4 q_3 \\ p_1 q_3 - p_2 q_4 + p_3 q_1 + p_4 q_2 \\ p_1 q_4 + p_2 q_3 - p_3 q_2 + p_4 q_1 \end{bmatrix} \quad (4.3)$$

¹It should be noted that the order of the terms is not particularly important; specifically, the placement of the cosine term as the first term in the vector is arbitrary. A number of texts place the cosine term as the last term in the quaternion. Either form is equally acceptable; the important point is to ensure that consistency is maintained throughout.

It is important to note that quaternion multiplication is not commutative, and so the final product is dependent upon the order in which the quaternions are multiplied.

4.2 Attitude State Vector

In many attitude determinations filters which fuse IMU measurements with an attitude measurement device such as a star tracker, the attitude quaternion itself is not stored in the state vector. Rather, the state vector for attitude consists of three attitude error states, along with three gyro bias estimates.

We define the attitude sub-component of the state vector as follows:

$$\hat{\mathbf{x}}_{\mathbf{q}} = \begin{bmatrix} \hat{\boldsymbol{\delta\mathbf{v}}} & \hat{\mathbf{b}} \end{bmatrix}^T \quad (4.4a)$$

$$= \begin{bmatrix} \delta\phi & \delta\theta & \delta\psi & \hat{b}^\phi & \hat{b}^\theta & \hat{b}^\psi \end{bmatrix}^T \quad (4.4b)$$

$$= \begin{bmatrix} (\phi - \hat{\phi}) & (\theta - \hat{\theta}) & (\psi - \hat{\psi}) & \hat{b}^\phi & \hat{b}^\theta & \hat{b}^\psi \end{bmatrix}^T \quad (4.4c)$$

Of course, even though the attitude state vector itself does not contain the attitude quaternion, the attitude quaternion must still be stored in memory as well. We use the multiplicative quaternion error representation to write the following relationship between the estimated attitude quaternion, the attitude error states, and the true quaternion [60].

$$\mathbf{q}_{True} = \delta\mathbf{q}(\boldsymbol{\delta\mathbf{v}}) \otimes \hat{\mathbf{q}} \quad (4.5)$$

This formulation implies that the attitude error term, $\delta\mathbf{q}(\boldsymbol{\delta\mathbf{v}})$ is defined in terms of the body frame, while the attitude quaternion and estimated attitude quaternion are both defined in the global frame.

Also, we note that in this formulation we are only estimating the attitude error states $\boldsymbol{\delta\mathbf{v}}$ and the gyro bias vector $\hat{\mathbf{b}}$. However, in general any number of auxiliary attitude-related terms could be estimated in the state vector including gyro-scale factors.

With this definition of the state vector, we may now derive the needed equations

for an attitude estimator. Attitude estimation is not a new topic, and much of the needed algorithms are widely available in literature. For brevity, we omit the derivations which do not differ from what is available in literature, and focus primarily on the new techniques developed in this dissertation which are necessary for the data association filter described in Chapter 5. However, a thorough treatment of attitude estimation is available [60].

4.3 Attitude Measurement Update Equations

We begin with the assumption that when a photon arrives, the detector measures the photon's angles of arrival (AOAs); specifically the arrival azimuth and elevation angle (Figures 4.2 and 4.3). In x-ray observatories such as *Chandra*, the detector pixel at which the photon is detected corresponds directly to the azimuth and elevation angles through a series of conversions which can be derived based on the configuration of the detector during the observation. We assume that any errors associated with the mapping of the of detector pixel to AOA are negligible in comparison to the other sources of AOA measurement error.

Next, we would like to compute a measurement update matrix which relates the photon's AOAs to the attitude state vector. To do so, we assume (for now) that we know the source of the photon in question, and have the source's coordinates in a global reference frame. We present two methods which may be used to produce an attitude measurement matrix.

4.3.1 Unit Vector Matching

For the sake of comparison, we will briefly review a common approach found in literature for updating the attitude state vector based on a angle measurement. In this method, an equivalent unit vector is computed from the angles of arrival [60]. This unit vector is then compared to the unit vector pointing to the source by taking the cross product of the two vectors (either transforming the source unit vector to the estimated body frame, or transforming the photon unit vector to the global frame). The cross product is related to the attitude estimate error, and this vector quantity is used to update the attitude.




Figure 4.2: AOA and unit vector




Figure 4.3: Photon angle of arrival measurement as viewed from the detector focal plane. It will be noted that the actual detector coordinates of the photon arrival will be a “mirror image” of the photon’s sky coordinates.

We note here that, in general, a unit vector \bar{u} may be computed from a set of angle-of-arrival measurements consisting of an azimuth measurement, α and an elevation measurement β , as illustrated in Figure 4.2 and Figure 4.3.

$$\bar{u}_1 = \cos(\alpha) \cos(\beta) \quad (4.6a)$$

$$\bar{u}_2 = \cos(\alpha) \sin(\beta) \quad (4.6b)$$

$$\bar{u}_3 = \sin(\beta) \quad (4.6c)$$

Similarly, the equivalent azimuth and elevation measurements may be computed from a unit vector measurement.

$$\alpha = \arctan 2(\bar{u}_2, \bar{u}_1) \quad (4.7a)$$

$$\beta = \arcsin(\bar{u}_3) \quad (4.7b)$$

As an aside, the reader may find it useful to note the approximate relationship between the unit vector components and the angle of arrival measurements which may be derived via the small-angle approximation. We note that if the angles of arrival are relatively small, then

$$\bar{u}_1 \approx 1 \quad (4.8a)$$

$$\tan(\alpha) = \frac{\bar{u}_2}{\bar{u}_1} \implies \alpha \approx \bar{u}_2 \quad (4.8b)$$

$$\sin(\beta) = \bar{u}_3 \implies \beta \approx \bar{u}_3 \quad (4.8c)$$

While the small-angle approximations are not used in this derivation, the approximate relationships may be useful in relating the two types of arrival vector parameterizations.

In the unit vector matching approach, a unit vector is computed for all attitude measurements, and that unit vector is then used in the state vector measurement update. This approach is convenient because it presents a universal measurement update

equation which may be used for many (if not all) types of attitude measurements. However, it is notable that the vector matching approach may introduce a third, dependent measurement quantity. This can be seen by noting that the unit vector consists of three values, while it is computed from only two measured values. This third dependent variable is not inherently problematic for the measurement update, but as will be shown, it causes problems in data association. Consequently, we seek an alternate approach that does not introduce a dependent variable into the measurement. This approach is outlined below.

4.3.2 Angle Measurements

While the unit-vector matching approach generates a unit vector from the angle-of-arrival measurements and uses that unit vector in the state vector measurement update, it is also possible to directly update the state from the angle-of-arrival measurements without computing a unit vector. In order to do so, we must write the measured quantities $\mathbf{y} = [\alpha \ \beta]^T$ as a function of the state vector. More specifically, we would like to formulate an expression for these measurements in terms of the attitude error; that is, the (intrinsic, body-frame) changes in roll, pitch and yaw that would correct the estimated attitude to the true attitude.

We first of all note that, using the relationship in Equation (4.7), the predicted measurement values may be computed by transforming the point source coordinates into the local body frame.

$$\hat{\mathbf{y}} = [\hat{\alpha} \ \hat{\beta}]^T \quad (4.9a)$$

$$= [\arctan 2(\bar{u}_2^B, \bar{u}_1^B) \ \arcsin(\bar{u}_3^B)]^T \quad (4.9b)$$

We then note that the true measurement may be written as a function of the attitude state error and the predicted measurement. We further note that if there was no error in the roll estimate, then true measurement would be equal to the predicted measurement plus the error in yaw and pitch, and measurement noise, as shown below (Figure 4.4).




Figure 4.4: Estimated and true measurements with pitch and yaw error only

$$\mathbf{y} = \begin{bmatrix} \alpha & \beta \end{bmatrix}^T = \begin{bmatrix} \hat{\alpha} + \delta\psi + v_\alpha \\ \hat{\beta} + \delta\theta + v_\beta \end{bmatrix} \quad (4.10)$$

Of course, the case in which roll is known perfectly but pitch and yaw are unknown is not a meaningful one. In the case where there is roll error in addition to error in pitch and yaw, the true measurement will be the result of the predicted measurement, rotated from the estimated coordinate frame to the true coordinate frame, plus the measurement noise and the error in yaw and pitch (Figure 4.5).

$$\mathbf{y} = \begin{bmatrix} \alpha & \beta \end{bmatrix}^T = \begin{bmatrix} \delta\psi + \hat{\alpha} \cos(\delta\phi) - \hat{\beta} \sin(\delta\phi) + v_\alpha \\ \delta\theta + \hat{\alpha} \sin(\delta\phi) + \hat{\beta} \cos(\delta\phi) + v_\beta \end{bmatrix} \quad (4.11)$$

From here, we see that we may easily derive the measurement matrix by taking the partial derivative of the measurement function with respect to the states to be measured.

$$\mathbf{H}_{\delta\mathbf{v}} = \left. \frac{\partial \mathbf{y}_{\delta\mathbf{v}}}{\partial \delta\mathbf{v}} \right|_{\delta\mathbf{v}=0, \mathbf{v}=0} = \left. \begin{bmatrix} \frac{\partial \alpha}{\partial \delta\phi} & \frac{\partial \alpha}{\partial \delta\theta} & \frac{\partial \alpha}{\partial \delta\psi} \\ \frac{\partial \beta}{\partial \delta\phi} & \frac{\partial \beta}{\partial \delta\theta} & \frac{\partial \beta}{\partial \delta\psi} \end{bmatrix} \right|_{\delta\mathbf{v}=0, \mathbf{v}=0} \quad (4.12a)$$

$$= \left. \begin{bmatrix} -\hat{\alpha} \sin(\delta\hat{\phi}) - \hat{\beta} \cos(\delta\hat{\phi}) & 0 & 1 \\ \hat{\alpha} \cos(\delta\hat{\phi}) - \hat{\beta} \sin(\delta\hat{\phi}) & 1 & 0 \end{bmatrix} \right|_{\delta\mathbf{v}=0, \mathbf{v}=0} = \begin{bmatrix} -\hat{\beta} & 0 & 1 \\ \hat{\alpha} & 1 & 0 \end{bmatrix} \quad (4.12b)$$

The gyro bias states are not directly measured by the photon angle of arrival. Consequently, the entire attitude state update matrix may be written as:

$$\mathbf{H}_{\delta\mathbf{v}, \mathbf{b}} = \begin{bmatrix} -\hat{\beta} & 0 & 1 & 0 & 0 & 0 \\ \hat{\alpha} & 1 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (4.13)$$




Figure 4.5: Estimated and true measurements with roll, pitch and yaw error

It is notable that this measurement equation allows us to compute the measurement update directly from physically measurable quantities, with physically measurable variances, rather than introducing a third dependent measurement variable. It is also notable that this measurement update matrix is relatively intuitive, as the quantities measured correspond rather directly to the state variables of interest.

The estimated *a posteriori* attitude errors may be used to update the attitude quaternion. In order to do so, we compute the quaternion corresponding to the computed attitude error state, take the product of the error quaternion and the *a priori* quaternion to compute the *a posteriori* attitude quaternion estimate, and finally normalize the resulting quaternion.

$$\delta \mathbf{q} \approx \begin{bmatrix} \delta \boldsymbol{\nu}/2 \\ 0 \end{bmatrix} \quad \hat{\mathbf{q}}^* = \delta \mathbf{q} \hat{\mathbf{q}}^- \quad \hat{\mathbf{q}}^+ = \frac{\hat{\mathbf{q}}^*}{\|\hat{\mathbf{q}}^*\|} \quad (4.14)$$

4.4 Overview of Attitude Estimation Algorithm

For completeness, and for reference in Chapter 5, the entire attitude estimation algorithm is presented here. As noted above, substantial portions of this algorithm are taken directly from literature, and thorough derivations of these equations are readily available [60].

The equations for updating the attitude state vector in time from time-step j to time-step k are given below:

$$\hat{\omega}_j^+ = \omega_j^+ - \hat{\mathbf{b}}_j^+ \quad (4.15a)$$

$$\mathbf{q}_k^- = \begin{bmatrix} \cos\left(\frac{1}{2}\|\hat{\omega}_j^+\|\Delta t_{j \rightarrow k}\right) \mathbf{I}_{3 \times 3} - [\Psi]^\times & \Psi \\ \Psi^T & \cos\left(\frac{1}{2}\|\hat{\omega}_j^+\|\Delta t_{j \rightarrow k}\right) \end{bmatrix} \mathbf{q}_j^+ \quad (4.15b)$$

$$\Psi = \frac{\sin\left(\frac{1}{2}\|\hat{\omega}_j^+\|\Delta t_{j \rightarrow k}\right) \hat{\omega}_j^+}{\|\hat{\omega}_j^+\|} \quad (4.15c)$$

$$\mathbf{P}_k^- = \mathbf{F}_{j \rightarrow k} \mathbf{P}_k^- \mathbf{F}_{j \rightarrow k}^T + \mathbf{Q}_{j \rightarrow k} \quad (4.15d)$$

$$\mathbf{F}_{j \rightarrow k} = \begin{bmatrix} \left(\hat{\omega}_j^+\right)^\times & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \end{bmatrix} \quad (4.15e)$$

$$\mathbf{Q}_{j \rightarrow k} = \begin{bmatrix} \left(\sigma_\omega^2 \Delta t_{j \rightarrow k} + \frac{1}{3} \sigma_{\mathbf{b}}^2 \Delta t_{j \rightarrow k}^3\right) \mathbf{I}_{3 \times 3} & -\left(\frac{1}{2} \sigma_{\mathbf{b}}^2 \Delta t_{j \rightarrow k}^2\right) \mathbf{I}_{3 \times 3} \\ -\left(\frac{1}{2} \sigma_{\mathbf{b}}^2 \Delta t_{j \rightarrow k}^2\right) \mathbf{I}_{3 \times 3} & \left(\sigma_\omega^2 \Delta t_{j \rightarrow k}\right) \mathbf{I}_{3 \times 3} \end{bmatrix} \quad (4.15f)$$

where

- $\hat{\omega}_j^+$ is the angular rate measurement measured by the IMU in the body frame
- $\Delta t_{j \rightarrow k}$ is the time elapsed between time-step j and k
- $\hat{\mathbf{b}}_j^+$ is the estimated angular rate measurement bias
- $[\cdot]^\times$ is the skew-symmetric matrix computed from a 3×1 vector
- $\sigma_\omega^2 = \text{var}(\omega_{meas} - \omega_{True})$ is the variance of the angular rate measurement error
- $\sigma_{\mathbf{b}} = \text{var}(\mathbf{b}_k - \mathbf{b}_j)$ is the variance of the angular rate bias random process noise

It is notable in Equation (4.15) that the value of the state vector does not change during the time update. This is because the state vector consists of the attitude error state (which is set to zero after the measurement update) and the IMU angular rate bias state, the dynamics of which are not modeled in this estimator.

The measurement update equations for updating the attitude state vector based on

an angle-of-arrival measurement are given below:

$$\mathbf{y}_k = \begin{bmatrix} \alpha_k & \beta_k \end{bmatrix}^T = \begin{bmatrix} \hat{\alpha}_k + \delta \psi_k + v_\alpha \\ \hat{\beta}_k + \delta \theta_k + v_\beta \end{bmatrix} \quad (4.16a)$$

$$\hat{\mathbf{y}}_k = \begin{bmatrix} \hat{\alpha} & \hat{\beta} \end{bmatrix}^T = \begin{bmatrix} \arctan 2(\bar{u}_2^B, \bar{u}_1^B) & \arcsin(\bar{u}_3^B) \end{bmatrix}^T \quad (4.16b)$$

$$\delta \mathbf{y}_k = \mathbf{y}_k - \hat{\mathbf{y}}_k \quad (4.16c)$$

$$\mathbf{H}_k = \begin{bmatrix} -\hat{\beta}_k & 0 & 1 & 0 & 0 & 0 \\ \hat{\alpha}_k & 1 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (4.16d)$$

$$\mathbf{R}_k = \begin{bmatrix} \sigma_\alpha^2 & 0 \\ 0 & \sigma_\beta^2 \end{bmatrix} \quad (4.16e)$$

$$\mathbf{S}_k = \mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^T + \mathbf{R}_k \quad (4.16f)$$

$$\mathbf{K}_k = \mathbf{P}_k^- \mathbf{H}_k^T \mathbf{S}_k^{-1} \quad (4.16g)$$

$$\begin{bmatrix} \delta \mathbf{v}_k^+ \\ \mathbf{b}_k^+ \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \mathbf{b}_k^- \end{bmatrix} + \mathbf{K}_k \delta \mathbf{y}_k \quad (4.16h)$$

$$\delta \mathbf{q} \approx \begin{bmatrix} \delta \nu / 2 \\ 0 \end{bmatrix} \quad (4.16i)$$

$$\hat{\mathbf{q}}^* = \delta \mathbf{q} \hat{\mathbf{q}}^- \quad (4.16j)$$

$$\hat{\mathbf{q}}^+ = \frac{\hat{\mathbf{q}}^*}{\|\hat{\mathbf{q}}^*\|} \quad (4.16k)$$

$$\delta \mathbf{v}_k = \mathbf{0}_{3 \times 1} \quad (4.16l)$$

where $\sigma_\alpha = \text{var}(\delta \alpha_k)$, $\sigma_\beta = \text{var}(\delta \beta_k)$ are the variances of the angle-of-arrival measurement errors.

The component of the state vector corresponding to the attitude error state, $\delta \mathbf{v}$, is computed solely for the purpose of updating the attitude quaternion. After this is complete, the estimated value of the attitude error state is zero, since the attitude estimate has already been updated to account for the estimated error. Consequently, it is important that this component of the state vector be set to zero after the quaternion is updated, as reflected in Equation (4.16l).

Nomenclature

Probability and Statistics

$\text{var}(a)$ The variance of random variable a

Signals and Measurements

ζ_i Vector of angle-of-arrival measurements at of photon i

State Estimation

\mathbf{q}_t^A	Attitude of object A, at time t (<i>radians</i>)
$\delta \mathbf{v}_t^A$	Error in attitude estimate of object A at time t (<i>radians</i>)
b	Bias of gyro measurement of angular velocity in axis ϕ (<i>radians/second</i>)
\mathbf{b}	Vector of gyro biases in three axes (<i>radians/second</i>)
$\boldsymbol{\omega}_t^A$	Angular velocity of object A at time t (<i>radians/second</i>)
ϕ_t^A	Roll angle of object A at time t (<i>radians</i>)
θ_t^A	Pitch angle of object A at time t (<i>radians</i>)
ψ_t^A	Yaw angle of object A at time t (<i>radians</i>)
α	Azimuth angle (<i>radians</i>)
β	Elevation angle (<i>radians</i>)
N	Denotes the global navigation frame
B	Denotes the spacecraft body frame
$\bar{\mathbf{u}}_A^B$	Unit vector to object A in frame B
j, k	Time-step indices
\mathbf{x}	State vector
\mathbf{P}	Covariance of state vector
$\mathbf{y}_{t,a}$	Measurement of scalar quantity a
\mathbf{H}	Measurement update matrix for measured quantity a
$v_{t,a}$	Scalar measurement error of measured quantity a at time t
$\mathbf{R}_{t,a}$	Variance of measurement noise in measured quantity a
$\delta \mathbf{y}_{t,a}$	Measurement residual for measured scalar quantity a at time t
$\mathbf{S}_{t,a}$	Variance of measurement residual for measured quantity a at time t
\mathbf{w}_t	Error in time update (process noise) at time t
\mathbf{Q}_t	Variance of process noise at time t
$[\cdot]^-$	Indicates <i>a priori</i> value
$[\cdot]^+$	Indicates <i>a posteriori</i> value
$\hat{\cdot}$	Indicates estimated value
\mathbf{a}	Indicates vector quantity
a	Indicates scalar quantity

Misc

$\mathbf{I}_{n \times n}$

The $n \times n$ identity matrix

Chapter 5

Data Association and Joint Range and Attitude Estimation

In Chapter 3 and Chapter 4, estimators for determining range and attitude based on photon arrival measurements from astrophysical signals of opportunity were derived. However, in both of these cases, it was assumed that the estimators “know” which photons originate from which signal sources, or that they are able to correctly associate photons with the correct signal sources.

In reality, the estimator must have a means of correctly associating photons with the correct signal source. In this chapter, we will derive a technique for data association based upon the Joint Probabilistic Data Association Filter (JPDAF) [61]. We will then demonstrate the efficacy of this method using both simulated data and x-ray observatory data.

5.1 Naive Data Association

As a first approximation to the data association problem, we consider a naive approach to associating photons with the correct signal sources. We do so because it helps illustrate the challenge of the data association problem. Later we describe a more sophisticated approach used for PNT estimation derived in this dissertation.

Consider the scenario depicted in Figure 5.1, where a single pulsar is in the field of view (FOV) of the x-ray detector being used for PNT. It is clear by inspection that




Figure 5.1: A naive approach to data association

the photon's angle of arrival, in conjunction with the attitude estimate, can be used to estimate the coordinates from which the photon originated in the navigation frame. Using this information, the angle between the photon's arrival vector and the vector to an arbitrary signal source may be computed. If the angle between these two vectors were below a certain threshold, then the photon would be associated with that signal source (Figure 5.1).

The threshold below which the angle between the two vectors must be depends upon the accuracy to which the spacecraft's attitude is known. A threshold based on a multiple of the attitude error standard deviation estimate would statistically guarantee that a certain percentage of photons originating from the signal source would be associated with that signal source. For example, if the rejection threshold is an angle greater than $2\sigma_{\delta_v}$, then statistically 95% of photons originating from any given signal source would be correctly associated with that signal source, whereas if the threshold was $3\sigma_{\delta_v}$ as shown in Figure 5.1, then 99.7% of the photons would be associated correctly.

If this algorithm is to be used, a trade-off must be made between the number of photons which are correctly associated and the number of background photons which are also incorrectly associated. As the rejection threshold is increased, more background photons will be erroneously associated with that signal source. If the attitude is known to a high enough accuracy, this might not be a problem since $3\sigma_{\delta_v}$ would be a relatively small quantity. However, if the attitude uncertainty is relatively large, and if the background flux is high compared to the source fluxes, the background photons may quickly "drown out" the signal source photons, resulting in a signal to noise ratio that quickly




Figure 5.2: A measurement update in the presence of source uncertainty

approaches zero. Both of these conditions are easily realizable in practice; consequently, a more intelligent data association approach is needed. This is particularly true if more than one x-ray source is in the field of view of the x-ray detector.

5.2 Joint Probabilistic Data Association Filter

To address the problem of data association, we turn to the Joint Probabilistic Data Association Filter (JPDAF) algorithm [61]. In what follows, we will show how the JPDAF algorithm may be used on joint position and attitude estimation using measurements of x-ray photons from pulsars. This work was initially described in [62].

The basic framework of the JPDAF algorithm is depicted graphically in Figure 5.2, and proceeds as follows: Once an x-ray photon is received by the x-ray detector, the

algorithm iterates through all the possible signal sources, and for each signal source computes the probability that the measurement originated from the given signal source. This probability is referred to as the association probability. After the probabilities are computed, the most likely signal sources are kept, while the less likely signal sources are rejected. The remaining signal sources are once again iterated over, and for each signal source, the updated state vector and covariance are computed, assuming that the given association is correct. Finally, the updated state vector is approximated as a Gaussian mixture of all of the possible updated states, weighted according to each association's probability. The time update then proceeds in the same manner as in a standard EKF. Each step in the JPDAF algorithm is given in more detail below.

5.2.1 Association Probabilities

The first step in the JPDAF algorithm is to compute the probability of association for each signal source. To compute the association probability, we use Bayes' rule for conditional probability.

$$\Pr [A_p | t_i, \zeta_i, \hat{\mathbf{x}}_t^-] = \frac{\Pr [t_i, \zeta_i | A_p, \hat{\mathbf{x}}_t^-] \Pr [A_p]}{\Pr [t_i, \zeta_i, \hat{\mathbf{x}}_t^-]} \quad (5.1)$$

where

- A_p is the event that a photon measured by the detector originated from signal source p , where the subscript p denotes the p^{th} signal source
- t_i is the time-of-arrival measurement
- $\zeta_i = \begin{bmatrix} \alpha & \beta \end{bmatrix}$ are the angle-of-arrival measurements
- $\hat{\mathbf{x}}_t^-$ is the *a priori* state vector estimate

The derivation of each term in the expression for probability is briefly presented below. We begin by noting the denominator on the right-hand side of Equation 5.1 is independent of the association event A_p and is therefore a constant that will appear in each probability term. Consequently we do not need to compute it, since we can simply divide it out from all terms. The other two terms in the equation, however, must be computed directly.

Measurement Probability

The first term in the numerator of Equation 5.1 is the probability of measuring a photon at time t_i and angles $\zeta_i = \begin{bmatrix} \alpha & \beta \end{bmatrix}$, given that the photon in question did in fact originate from signal p , expressed mathematically as $\mathbf{Pr} [t_i, \zeta_i | A_p, \hat{\mathbf{x}}_\ell^-]$. We begin by making the approximation that the angle of arrival and time of arrival are statistically independent of each other under the assumption that we know the origin of the photon. We justify this assumption by noting that for a given signal source, t_i depends on the distance between the signal source and the detector, while ζ_i depends on the attitude of the detector. While the position and attitude estimate errors of the spacecraft are coupled, the effects of that coupling are small enough to be neglected, which allows us to assume independence. Therefore, this probability may be split into two separate terms.

$$\mathbf{Pr} [t_i, \zeta_i | A_p, \hat{\mathbf{x}}_\ell^-] = \mathbf{Pr} [\zeta_i | A_p, \hat{\mathbf{x}}_\ell^-] \mathbf{Pr} [t_i | A_p, \hat{\mathbf{x}}_\ell^-] \quad (5.2)$$

The first term on the right-hand side of Equation 5.2 is the probability of measuring a photon from signal source p at angle-of-arrival ζ_i , and incorporates the current attitude estimate. We will assume that the state estimate error and the AOA measurement error are both Gaussian. This implies that, given the photon did originate from source p , the attitude measurement residual may be described as:

$$\delta \mathbf{y}_\zeta \sim \mathcal{N} [\mathbf{0}, \mathbf{S}_\zeta] \quad (5.3)$$

where \mathbf{S}_ζ is the measurement residual variance, as defined in Equations 4.16c and 4.16f, and is a function of the *a priori* state vector, $\hat{\mathbf{x}}_\ell^-$. Consequently, the angle-of-arrival probability may be written from the standard normal distribution probability density function (PDF) as follows:

$$\mathbf{Pr} [\zeta_i | A_p, \hat{\mathbf{x}}_\ell^-] = \frac{e^{-\frac{1}{2} \delta \mathbf{y}_\zeta^T \mathbf{S}_\zeta^{-1} \delta \mathbf{y}_\zeta}}{\sqrt{\det(2\pi \mathbf{S}_\zeta)}} d\zeta \quad (5.4)$$

It is important to remember that the PDF is not a direct expression for probability, and normally must be integrated over some range of values to compute the probability of the random variable falling in that range. We avoid this step by including the final

term on the right-hand side, $d\zeta$. This term is an infinitesimally small range of angles. As long as this term is present in every association probability, we may simply divide it out of every term (leveraging the assumption that the probabilities must sum to one). However, we note here the importance of consistency in the choice of attitude measurement matrices. The $d\zeta$ term *must be the same* for every probability that is computed. If we use the attitude measurement update described above, then $d\zeta = d\alpha d\beta$. However, if the more standard unit-vector of arrival approach is used then $d\zeta = du_1 du_2 du_3$. The main point here is that the choice must be consistent across all the computed probabilities.

The next term to be computed is the time-of-arrival probability, $\mathbf{Pr} [\tau_i | A_p, \hat{\mathbf{x}}_\tau^-]$. Unlike angle of arrival, the time-of-arrival is not normally distributed. Rather, photon arrival times are a Poisson process, or in the case of a time-varying flux, the arrival times are non-homogeneous Poisson process (NHPP). Event times governed by an NHPP are distributed as follows [63].

$$f_\tau(t | \lambda(t)) = \lambda(t) e^{-\int_s^t \lambda(t) dt} \quad (5.5)$$

where $\lambda(t)$ is the photon flux from the signal source at time t and s is some start time prior to the photon arrival time. We may make use of the memory-less property of Poisson/exponential processes to simplify this expression. It may be shown that for an exponentially distributed arrival times:

$$\mathbf{Pr} [\tau > t + \delta t | \tau > t] = \mathbf{Pr} [\tau > \delta t] \quad (5.6)$$

That is to say that the fact that an event has or has not occurred in the recent past has no bearing on the probability of an event occurring at a future time. At a given photon arrival time, we can let δt in Equation 5.6 approach zero. Then, our probability may be simplified as:

$$\mathbf{Pr} [\tau | \lambda(t), \hat{\mathbf{x}}_\tau^-] = \lambda(t) dt \quad (5.7)$$

Once again, as with the computation of angle-of-arrival probability, we must include the differential term dt . This will be divided out later, but must be present in every




Figure 5.3: Uncertainty in the spacecraft’s position results in uncertainty in the estimated flux at t .

association probability that we compute.

It is important to note the fact that the term λ refers to the photon flux from the signal source at the spacecraft’s current *time* (t) and location, or range (ρ). This flux is equal to the flux at the origin (λ_O) delayed by the true TDOA ($\tau^{(S)}$). However, the true TDOA is a function of position $\hat{\tau}^{(S)} = \hat{\mathbf{r}}_{t_i} \cdot \boldsymbol{\ell}^{(p)}$, and its true value is not known exactly. Since the TDOA is not known, the flux from the signal source is not known exactly either. Instead, we can only compute the expected value of the flux.

At first, it may seem that this problem is trivial: We may simply evaluate the flux at the estimated TDOA to obtain its expected value. However, this approach is not correct, because flux is non-linear in t . To illustrate why this is problematic, consider Figure 5.3.

This figure illustrates how an incorrect estimate of position can result an incorrect estimate of flux at t . Furthermore, as shown in the figure, the incorrect estimate of flux can vary dramatically from the true flux, especially if the flux time-profile is very discontinuous. As shown in Equation (5.7), the time of arrival probability term simply the estimated flux at that time. Therefore, if the estimate of flux is dramatically different from the true flux, the computed time of arrival probability will either be erroneously high or low, resulting in an incorrect association probability.

To address this problem we may compute the expected value of flux, given the

estimate and uncertainty of the spacecraft's position. The expected value is given by the following equation:

$$\begin{aligned}\mathbb{E} \left[\lambda_O \left(t_i + \hat{\tau}^{(S)} \right) \right] &= \int_{-\infty}^{\infty} \lambda_O(t) f(t|\tau, \sigma_\tau^2) dt \\ &= \int_{-\infty}^{\infty} \lambda_O(t + \hat{\tau}) \frac{1}{\sqrt{2\pi\sigma_\tau^2}} e^{-\frac{(t+\hat{\tau}-\tau)^2}{2\sigma_\tau^2}} dt\end{aligned}\tag{5.8}$$

This integral cannot be directly computed. It could be numerically computed for each photon, but this could become computationally expensive. Instead, we approximate the PDF of the spacecraft's TDOA estimate as a moment-matched uniform distribution, which yields the following approximation.

$$\begin{aligned}\mathbb{E} \left[\lambda_O \left(t_i + \hat{\tau}^{(S)} \right) \right] &\approx \frac{1}{\sqrt{12}\sigma_\tau} \int_{t_i + \hat{\tau} - \sqrt{3}\sigma_\tau}^{t_i + \hat{\tau} + \sqrt{3}\sigma_\tau} \lambda_O(t_i + t) dt \\ &= \frac{F_{\lambda_O}(t_i + \hat{\tau} + \sqrt{3}\sigma_\tau) - F_{\lambda_O}(t_i + \hat{\tau} - \sqrt{3}\sigma_\tau)}{\sqrt{12}\sigma_\tau} \\ &= \frac{F_{\lambda_O} \Big|_{t_i + \hat{\tau} - \sqrt{3}\sigma_\tau}^{t_i + \hat{\tau} + \sqrt{3}\sigma_\tau}}{\sqrt{12}\sigma_\tau}\end{aligned}\tag{5.9}$$

where $F_{\lambda_O} = \int \lambda_O(t) dt$ is the indefinite integral of the flux profile with respect to time.

This expression is the result of approximating the normally distributed TDOA estimate error as a uniform distribution with variance equal to that of the original normal distribution. The integral in the approximation above is easily obtained because the indefinite integral can be numerically computed once for each signal source. The results of this indefinite integral can then be used to compute the definite integral.

In summary, the measurement probability may be written as

$$\mathbf{Pr} [t_i, \zeta_i | A_p, \hat{\mathbf{x}}_t^-] = \frac{e^{-\frac{1}{2} \delta \mathbf{y}_\zeta^T \mathbf{S}_\zeta^{-1} \delta \mathbf{y}_\zeta}}{\sqrt{\det(2\pi \mathbf{S}_\zeta)}} + \mathbb{E} \left[\lambda_O \left(t_i + \hat{\tau}^{(S)} \right) \right] dt d\zeta \quad (5.10a)$$

$$= \frac{e^{-\frac{1}{2} \delta \mathbf{y}_\zeta^T \mathbf{S}_\zeta^{-1} \delta \mathbf{y}_\zeta}}{\sqrt{\det(2\pi \mathbf{S}_\zeta)}} \frac{F_{\lambda_O} \big|_{t_i + \hat{\tau} - \sqrt{3}\sigma_\tau}^{t_i + \hat{\tau} + \sqrt{3}\sigma_\tau}}{\sqrt{12}\sigma_\tau} dt d\zeta \quad (5.10b)$$

Background Probability

The above expression describes the measurement probability for the case in which the photon originates from a discrete point source. However, photons may also originate from the ambient x-ray background. This probability also must be computed.

The time-of-arrival term remains the same, simply replacing the signal source flux with the x-ray background flux. Also, since background flux is assumed to be constant, the expected value is simply equal to the constant value of background flux, so the definite integral term may be replaced with the constant value of background flux. However, the expression derived for angle-of-arrival probability does not accurately represent the background photon case. This is because photons from the x-ray background are distributed uniformly in two dimensions (or at least they may be approximated as such). For the background case, we may say:

$$\zeta_i | A_{\text{bkg}}, \hat{\mathbf{x}}_t^- \sim \mathcal{U} [-f/2, f/2] \quad (5.11)$$

where f is detector field of view. Consequently, the angle-of-arrival probability for a background photon may be written as:

$$\mathbf{Pr} [\zeta_i | A_{\text{bkg}}, \hat{\mathbf{x}}_t^-] = \frac{1}{\Omega} d\zeta = \frac{1}{2\pi(1 - \cos(f))} d\alpha d\beta \quad (5.12)$$

Equation is derived from a two-dimensional uniform distribution over the solid angle subtended by the detector's field of view [64]. This particular expression assumes a conical field of view, but similar expressions could be derived for other field of view geometries.

It is here that we see the advantage of the attitude measurement update equation

that uses azimuth and elevation directly. It is straight-forward to compute the uniform distribution probability with $d\alpha$ and $d\beta$ as the differential quantities. However, with the introduction of a third dependent variable, as is required by the unit vector measurement update, it is not clear how to model the uniform distribution. Consequently, for the computation of probabilities, it is simpler to use the direct approach shown here.

In summary, the measurement probability for a background photon may be written as follows.

$$\mathbf{Pr}[\mathcal{E}_i, \zeta_i | A_{\text{bkg}}] = \frac{1}{\Omega} \lambda_{\text{bkg}} d\mathcal{E} d\zeta = \frac{1}{2\pi(1 - \cos(f))} \lambda_{\text{bkg}} d\mathcal{E} d\alpha d\beta \quad (5.13)$$

Prior Association Probability

The next term in Equation 5.1 is $\mathbf{Pr}[A_p]$. This term is often referred to as the “prior” probability, because it describes the probability of the event occurring prior to the measurement being taken, and independent of the measurement values.

Because this term is to be computed independently of the measurement or the state vector, it is dependent only on the average flux of each signal source. This probability may be written directly as follows for each signal source.

$$\mathbf{Pr}[A_p] = \frac{\mathbb{E}[\lambda^{(p)}]}{\lambda^{(\text{bkg})} + \sum_{i=1}^N \mathbb{E}[\lambda^{(i)}]} \quad (5.14)$$

where

- $\mathbb{E}[\cdot]$ is the expected value
- i is the index for each x-ray source
- N is the number of sources in view.

Similarly, the probability of seeing a photon from the ambient x-ray background may be computed as follows.

$$\mathbf{Pr}[A_{\text{bkg}}] = \frac{\lambda^{(\text{bkg})}}{\lambda^{(\text{bkg})} + \sum_{i=1}^N \mathbb{E}[\lambda^{(i)}]} \quad (5.15)$$

The final step in the computation of probabilities is to sum up all the computed probabilities, then divide each probability by the computed sum. In this way we remove the differential angle and time terms, as well as the constant denominator term.

5.2.2 Joint Measurement Update

Having computed the association probabilities for all the possible association events, the next task is to compute the updated state vector. Before we do this, it is beneficial to go through the association events, and reject those events which have a probability below some very small threshold. This step is not strictly necessary, but it does save on computational time and is therefore beneficial from an efficiency standpoint.

After this step, we iterate through the remaining association events, and compute the updated state vector for each association event, assuming that the association event is correct. We note that the state vector is a combined state vector which contains both the attitude estimate and the TDOA estimate, (or quantities which may be used to directly compute attitude and TDOA). The joint state vector is written as:

$$\mathbf{x} = \begin{bmatrix} \hat{\mathbf{x}}_{\mathbf{q}} \\ \hat{\mathbf{x}}_{\tau} \end{bmatrix} \quad \mathbf{P} = \begin{bmatrix} \mathbf{P}_{\mathbf{q}} & \mathbf{P}_{\mathbf{q}\tau} \\ \mathbf{P}_{\tau\mathbf{q}} & \mathbf{P}_{\tau} \end{bmatrix} \quad (5.16)$$

Once again, we note that a method for estimating $\hat{\mathbf{x}}_{\tau}$ is derived in Chapter 3. For each measurement association, the *a posteriori* state vector is computed as if that measurement association were the correct one. These updated state vectors are then used to compute a sort of “weighted average” state vector, by treating the true state vector’s distribution as a Gaussian mixture. The final updated state vector may then be written as follows.

$$\hat{\mathbf{x}}_k^+ = \sum_{i=1}^N w_i \hat{\mathbf{x}}_{k|A_i}^+ \quad (5.17)$$

where $w_i = \mathbf{Pr}[A_p | \mathcal{I}_i, \zeta_i, \hat{\mathbf{x}}_{\ell}^-]$. Similarly, we compute the updated covariance based upon the covariance of a Gaussian mixture.

$$\mathbf{P}_k^+ = \sum_{i=1}^N w_i \left(\mathbf{P}_{k|A_i}^+ + (\hat{\mathbf{x}}_k^+ - \hat{\mathbf{x}}_{k|A_i}^+)(\hat{\mathbf{x}}_k^+ - \hat{\mathbf{x}}_{k|A_i}^+)^T \right) \quad (5.18)$$

5.3 Results

The algorithm described above is validated using simulated data as well as data collected from x-ray detectors in space. We describe and discuss these validation experiments below.

5.3.1 Monte Carlo Simulation Results

In order to verify the performance of the estimator presented here, a series of Monte Carlo simulations was performed. The simulations show that the full, seven degree of freedom PNT solution (position, orientation and time) can be estimated. However, the solution is sensitive to initial conditions, in particular the error in the initial attitude estimate. Thus, the simulation results are presented parametrically and show the effects of initial attitude uncertainty on the final position estimate error.

The simulations were performed as follows. For each signal source, photon measurements were generated containing angle and time of arrival measurements. Additionally, background photon measurements were generated based on empirical models for hard x-ray background flux.[48]

The JPDAF estimator results were bench-marked against two other estimators. The first estimator is an unrealizable “ideal” estimator. The “ideal” estimator had perfect knowledge of which photons originated from the signal source, and which photons were background photons. The second estimator is the “naive” estimator discussed in Section 5.1. This estimator filtered photons based on its initial estimate of attitude. All photons falling within the $\pm 3\sigma$ angle of arrival window were included in the range estimate, while photons falling outside that window were rejected. Note that both the naive estimator and the JPDAF estimator started with identical initial attitude estimates. The parameters of the pulsar and of the Monte Carlo simulations performed are given in Tables 5.1 and 5.2.

The resulting TDOA estimate error and attitude estimate error standard deviations as a function of initial attitude error are shown in Figures 5.4 to 5.7. The initial attitude error is defined as the standard deviations of the distribution from which the initial attitude error was drawn, which was the same for all three axes. The results of the Monte Carlo simulations reveal that in the case of the naive estimator, performance

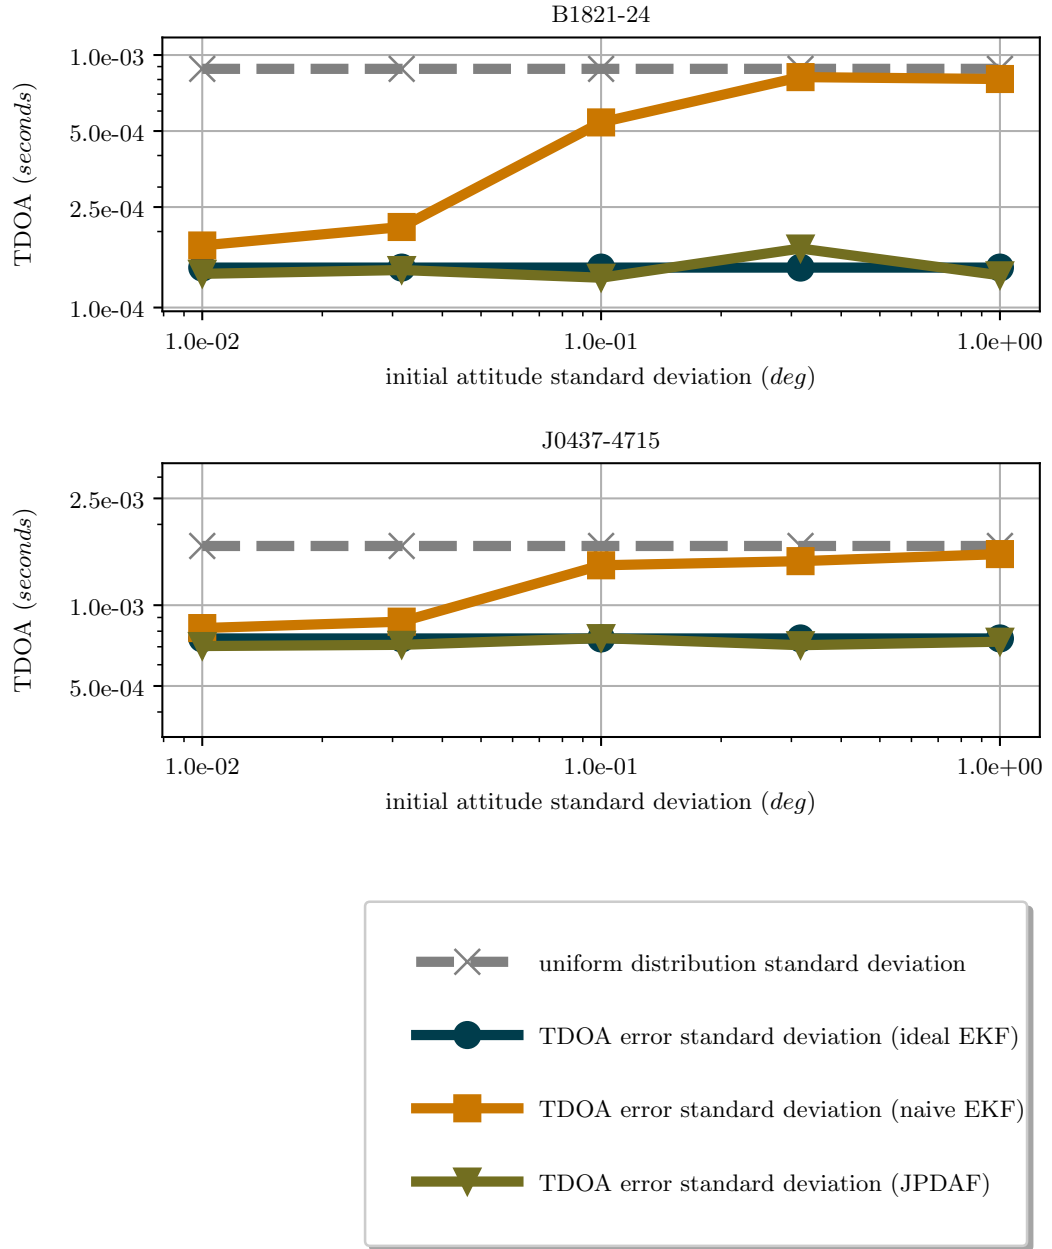


Figure 5.4: TDOA estimate standard deviations, 100 cm² detector area




Figure 5.5: TDOA estimate standard deviations, 1000 cm² detector area

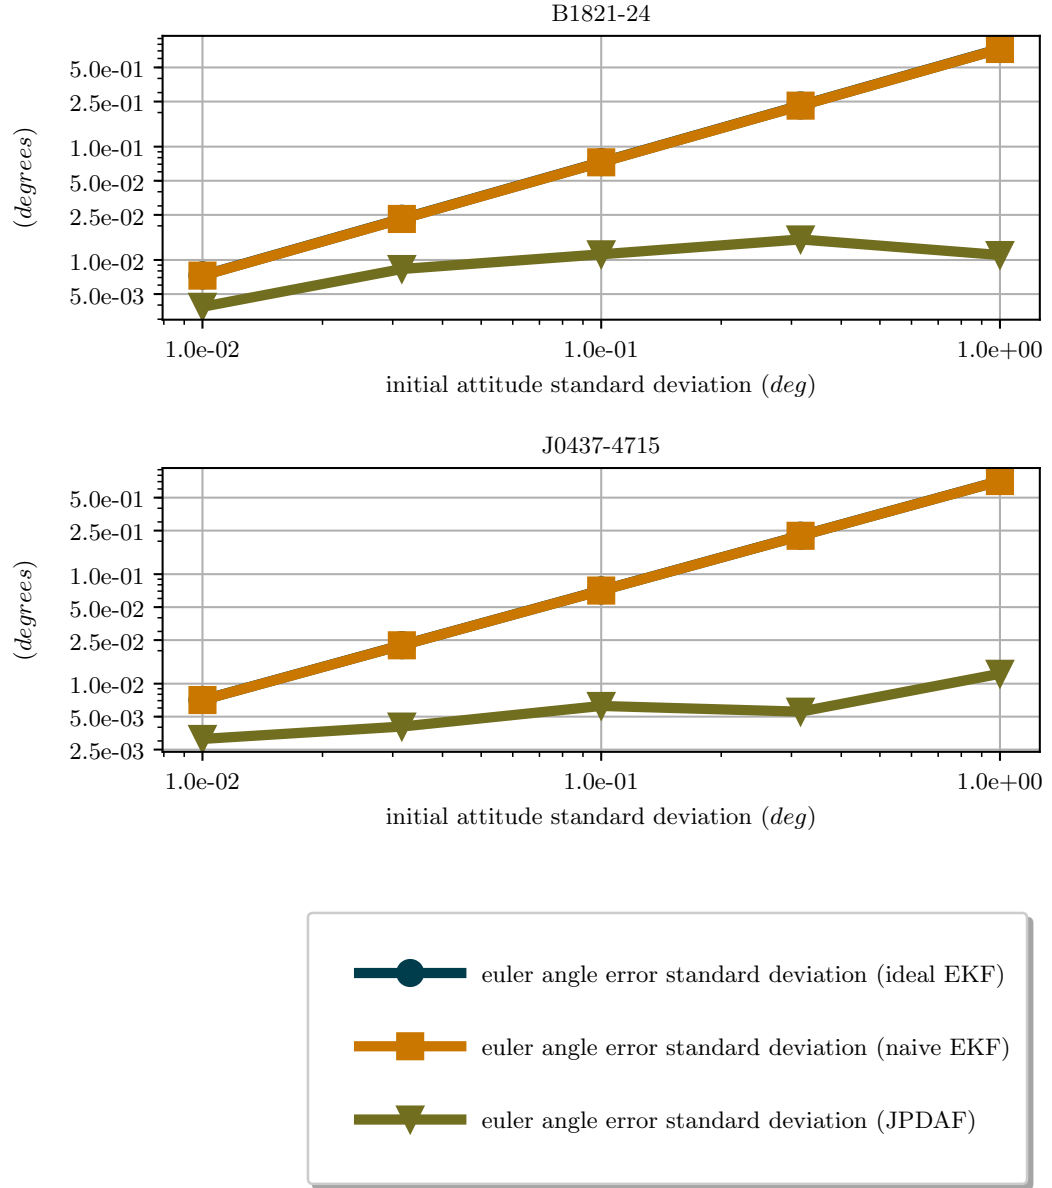


Figure 5.6: Attitude estimate standard deviations, 100 cm² detector area




Figure 5.7: Attitude estimate standard deviations, 1000 cm² detector area

Table 5.1: Pulsar Parameters

pulsar name	area (cm^2)	photon count rate $\left(\frac{photons}{second}\right)$	pulsar period ($seconds$)	pulsed fraction	background count rate $\left(\frac{photons}{second}\right)$
B1821-24	1.00e+02	1.30e-02	3.05e-03	9.80e-01	7.22e-01
B1821-24	1.00e+03	1.30e-01	3.05e-03	9.80e-01	7.22e+00
J0437-4715	1.00e+02	1.78e-02	5.76e-03	2.70e-01	7.22e-01
J0437-4715	1.00e+03	1.78e-01	5.76e-03	2.70e-01	7.22e+00

Table 5.2: Monte Carlo Input Parameters

pulsar name	runtime ($hour$)	area (cm^2)	run count
B1821-24	1	1.00e+02	500
B1821-24	1	1.00e+03	500
J0437-4715	1	1.00e+02	500
J0437-4715	1	1.00e+03	500

dramatically degrades as the initial attitude error increases. This is to be expected, since the background flux increases as a function of the effective field of view. In the naive estimator, as initial attitude error increases, the effective field of view also increases, and consequently the error of the final TDOA solution increases.

On the other hand, the performance of the JPDAF estimator matches very closely the performance of the unachievable “ideal” estimator. This is because the JPDAF estimator is able to update the initial poor estimate of attitude based on the photon AOA measurements (Figures 5.6 and 5.7).

5.3.2 *Chandra* Validation Experiment

In order to further verify the efficacy of the estimation techniques presented here, a series of validation experiments were performed on pulsar observation data from the x-ray observatory *Chandra*. *Chandra* is an x-ray telescope that includes x-ray detectors with high angular and temporal resolution. Its highly elliptical orbit allows for very long duration continuous observations of up to 55 hours in length. In these experiments, the pulsar observation data was used in the algorithms outlined above to estimate the range




Figure 5.8: In both *Chandra* observations, at least one static x-ray star was visible in the detector’s field of view in addition to the pulsar being observed. This allowed for a full three dimensional attitude solution to be estimated.

Table 5.3: HRC-S Parameters (timing mode)

Effective (cm ²)	detector area	Timing resolution (μ s)	Angular Resolution (arc-sec)
224	16		0.4

and attitude of the spacecraft. The estimated range and attitude were then compared to the true range and attitude as reported by the spacecraft.

Of the instruments aboard *Chandra*, the high-resolution camera (HRC) offered the best spatial and temporal resolution for the estimation algorithm to be tested, and some of the relevant parameters of the instrument are outlined in Table 5.3. Due to the HRC timing error, only a small number of *Chandra*’s observations contained valid timing measurements, which limited the number of available observations [65]. Of those observations, only a small subset were observations of pulsars which could be used for navigational purposes.

Of the available observations, observations of pulsars PSR B1509-58 and PSR B0540-6919 were chosen for analysis. These two were chose because they had a number of favorable characteristics, including:

- **Uniform background:** The algorithm presented above assumes a uniform background noise. While the estimator could be modified to account for a spatially varying background, that capability has not yet been incorporated. As a result, pulsars with dense, spatially varying backgrounds (for instance the Crab pulsar

in the Crab nebula) do not perform well with this estimator.

- **Non-binary systems:** Currently the estimator does not account for binary pulsars and the associated variation in frequency. This capability could be added to the estimator at a later time.
- **Relatively strong signals:** The two pulsars in the chosen observations are among the brighter of the available x-ray pulsars. While still about an order of magnitude weaker than the Crab pulsar, they offer a signal that is one to two orders of magnitude stronger than some of their weaker counterparts. While the algorithms presented here can function with any signal strength, the observation times needed to converge for a weaker pulsar are generally longer than the available observation times from *Chandra*.

The author was unable to locate accurate ephemerides, including accurate phase and frequency information for the pulsars in the specified observation at the time at which the observations were taken. This is due to the fact that these pulsars are less stable than the more accurate millisecond pulsars, so the published ephemerides typically have a shorter valid time range. This does not necessarily imply that these pulsars would be unusable for x-ray navigation; only that the spacecraft would have to receive more frequent ephemeris information in order to use these signals. Due to the lack of accurate timing information, the observations themselves were used to generate pulse phase and frequency parameter files. This was done by computing the photons' true time of arrival at the solar-system barycenter, using the spacecraft's true position, and then using those "barycentered" times to compute the phase and frequency of the pulsar signal during the observation.

In addition to the pulsars, there were other x-ray bright stars in the field of view of the HRC in both observations. These signals are possible signal sources, and must also be included in the data association and state update as possible signal sources. Position and flux data from the XMM Newton Slew Survey[66] was used to initialize all sources within the possible field of view.

Using these generated parameters to model the pulsar signal, the observations were processed using the algorithms outlined above. The algorithms were initialized with a poor estimate of attitude and a poor initial estimate of velocity (Table 5.4). No

Table 5.4: Standard deviations on simulation inputs

Target object	Initial velocity error (km/s)	Angular velocity measurement error (deg/s)	Initial attitude (deg)		
			roll	pitch	yaw
B0540-6919 (Obs. ID 1735, 1736)	1.00e+00	1.00e-04	1.00e+01	1.00e-01	1.00e-01
B1509-58 (Obs. ID 5515)	1.00e+00	1.00e-04	1.00e+01	1.00e-01	1.00e-01

information was included regarding the orbit model of the spacecraft.

The results of these simulations are shown in Table 5.5, as well as Figures 5.9 and 5.10.

There are a few points which should be noted about these results. The first is that despite being given no *a priori* knowledge about a given photon’s source of origin, the algorithm successfully calculates the change in range of the spacecraft along the line of sight to the pulsar. As is shown in Figure 5.9, the algorithm is able to compute the range of the spacecraft to an accuracy of approximately 1000 km. In the context of deep-space navigation, errors of this order of magnitude are frequently within the specifications of what the spacecraft needs to successfully complete the mission. As an example, the recent *New Horizons* flyby of Pluto required position accuracy of several thousand kilometers in order to achieve initial target acquisition of Pluto [5].

Second, while the algorithm was initialized with an initial attitude error, the pitch and yaw estimates quickly converge to the correct value. Additionally, in both observations, the roll error converges to a constant non-zero value. This is because in both observations, an additional x-ray star was visible in the field of view, resulting in two non-collinear vectors and enabling a full three degree of freedom attitude solution, as shown in Figure 5.8. However, the cataloged coordinates for these sources varied slightly from the observed positions of the sources, which resulted in a small constant offset in roll estimate. Nevertheless, these results demonstrate that is possible to converge from

Table 5.5: Estimate error standard deviations from *Chandra* validation experiments

Target object	Runtime (s)	Range estimate error standard deviation (km)	Attitude estimate error standard deviation (degrees)		
			roll	pitch	yaw
B0540-6919 (Obs. ID 1735, 1736)	2.04e+04	6.85e+02	3.84e+00	1.56e-02	6.53e-03
B1509-58 (Obs. ID 5515)	4.52e+04	1.15e+03	1.10e+00	3.98e-03	4.75e-03

1

an initial attitude error and estimate a full three degree of freedom attitude solution from x-ray photon observations, in addition to a range estimate.

Nomenclature

Data Association

i	Iterator to designate association event
A_S	The event that a given photon originated from signal source S
bkg	Indicates background photon flux
p	Indices to designate signal source
N	Number of signals being tracked
w_A	Probability of association event A

Detector

f	Field of view of detector (<i>rad</i>)
Ω	Solid-angle of detector (<i>steradian</i>)
A_D	Detector area (cm^2)

Physical Constants

c Speed of light (km/s)

Probability and Statistics

$\Pr [A]$ Probability of event A occurring

$\mathbb{E} [a]$ The expected value of random variable a

$\mathcal{N} [\mu, P]$ A normal distribution with mean μ and variance P

$\mathcal{U} [l, u]$ A uniform distribution with lower bound l and upper bound u

$f_X (x)$ The probability density function of a random variable X evaluated at x

Signals and Measurements

t_{obs} Total observation time (*seconds*)

p_f Fraction of flux from pulsar that is pulsed

W Full width at half maximum of primary pulsar pulse

$\lambda_k^{(p)}, \lambda^{(p)} (t)$ Expected number of photons at time index k , or expected photon flux at time t , from signal source p (*photons, photons/second*)

T_p The period of pulsar p (*seconds*)

i Photon index

ζ_i Vector of angle-of-arrival measurements at of photon i

t_i Time-of-arrival measurement of photon i (*seconds*)

S^i i^{th} signal source

State Estimation

$\rho_t^{A(S)}$	Range to object A, at time t, along line-of-sight vector to source S (<i>km</i>)
\mathbf{r}_t^A	Position of A, at time t (<i>km</i>)
\mathbf{q}_t^A	Attitude of object A, at time t (<i>radians</i>)
$\delta \mathbf{v}_t^A$	Error in attitude estimate of object A at time t (<i>radians</i>)
α	Azimuth angle (<i>radians</i>)
β	Elevation angle (<i>radians</i>)
O	Denotes the navigation frame origin
j, k	Time-step indices
t	Current time (<i>seconds</i>)
$\tau_t^{A(S)}$	Time-difference of arrival of signal from signal source S measured between vehicle A and navigation frame origin (<i>seconds</i>)
\mathbf{x}	State vector
\mathbf{P}	Covariance of state vector
$\delta \mathbf{y}_{t,a}$	Measurement residual for measured scalar quantity a at time t
$\mathbf{S}_{t,a}$	Variance of measurement residual for measured quantity a at time t
$[\cdot]^-$	Indicates <i>a priori</i> value
$[\cdot]^+$	Indicates <i>a posteriori</i> value
$\hat{\cdot}$	Indicates estimated value
\mathbf{a}	Indicates vector quantity
a	Indicates scalar quantity

Misc

A, B	Indices to designate spacecraft(s)
$\ell^{(S^i)}$	Unit vector pointing towards i^{th} signal source

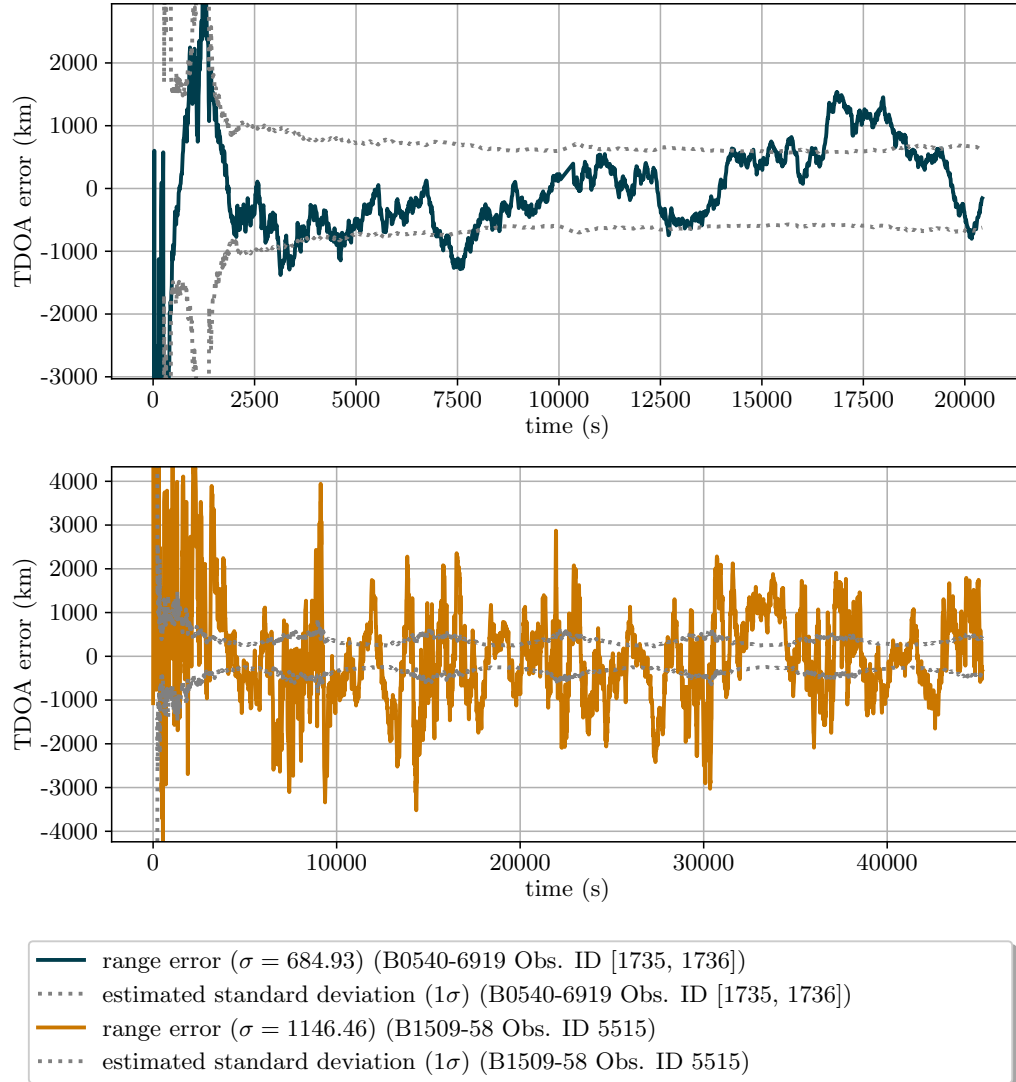


Figure 5.9: Range estimate errors

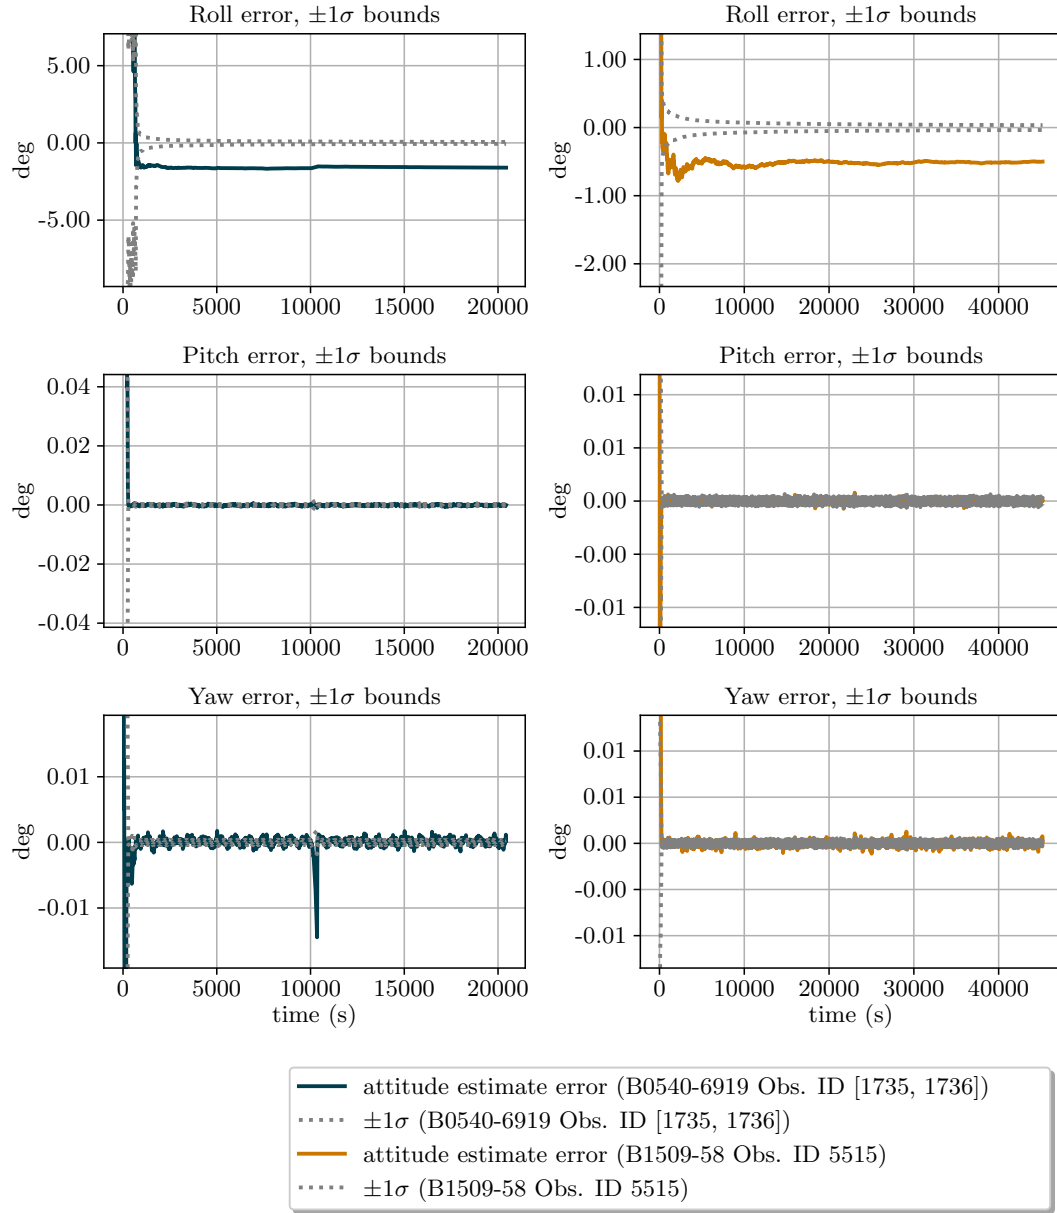


Figure 5.10: Attitude estimate errors

Chapter 6

Conclusions and Future Work

6.1 Conclusions

In this dissertation, we have shown that astrophysical signals of opportunity, in particular those generated by x-ray pulsars can be used to estimate a full six-degree of freedom navigation solution, including three degrees of position and three degrees of attitude. An extensive campaign of Monte Carlo simulations was performed to validate the performance of the algorithms presented here. Additionally, experimentally collected data from two x-ray observatories, *Chandra* and *Suzaku* were used to provide additional validation of the techniques presented here.

There are many areas of inquiry which must be addressed before pulsar-based x-ray navigation becomes a reality. In this closing chapter we describe some of these open questions and present preliminary results. In particular, we discuss the problem of initializing the filters derived in Chapter 3 and Chapter 5.

6.2 Filter Initialization

In this work, as in every proposed x-ray navigation scheme of which the author is aware, the navigation technique relies on an accurate measurement of the signal time-difference of arrival, as outlined in Section 1.2. This method is particularly useful in the case where the starting position and velocity of the spacecraft are known to a high degree of accuracy, allowing the spacecraft to simply track the change in TDOA of the




Figure 6.1: Integer ambiguity. The signal phase ϕ is estimated by algorithms of Chapters 3 and 5. However, N is unobservable in these algorithms.

pulsar signal to estimate distance traveled from the known starting position.

Another, more challenging problem is the so-called “lost-in-space” problem. This problem refers to the scenario in which a spacecraft has no prior knowledge of position or velocity, and is attempting to determine its position in the solar system. There are many challenges associated with this problem. One such challenge is that it is very difficult to estimate the TDOA of a pulsar signal if the spacecraft’s velocity is unknown.

In addition, even if the algorithm is able to estimate the TDOA, there remains the so-called “integer ambiguity” problem. This problem arises from the fact that the pulsar signal is, by its very nature, periodic. Consequently it is impossible to determine how many periods have elapsed between the spacecraft and the navigation frame origin. Put another way, if the time of arrival of a given peak is estimated, there is no way to know *which* peak is being measured; there are in fact an infinite number of peaks, which leads to an infinite number of possible solutions. This problem is illustrated in Figure 6.1.

There are multiple ways in which the integer ambiguity problem could be addressed. An analogy exists between this problem and the carrier-phase GNSS integer ambiguity resolution problem. However, it is not clear if the similarity between these problems is close enough such that the techniques used to address the phase integer ambiguity problem in GNSS could be applied to this problem as well [67].

We propose a different approach to solving the integer ambiguity problem. We note firstly that if the spacecraft had knowledge of its position to an accuracy less than the period of the pulsar times the speed of light, then the spacecraft could correctly determine the phase integer value. Rather than use the signal time-of-arrival directly to estimate position, we propose to instead estimate the velocity and acceleration of the spacecraft based on the measured photon arrival times. This problem is essentially equivalent to estimating the frequency or Doppler shift of the pulsar signal due to the spacecraft velocity, and additionally estimating the rate of change of that Doppler shift.

Because acceleration in space is almost entirely gravitational, a spacecraft's acceleration may be written as a function of its position. Therefore, the acceleration itself may be used as a measurement of the spacecraft's position. Given a sufficiently accurate estimate of acceleration, the spacecraft could estimate its position, velocity, and an approximate orbit model. This information could then be used to perform the more precise TDOA based positioning.

This proposed method of XNAV has the advantage that it is *not limited to extremely accurate millisecond pulsars*. This is significant, because millisecond pulsars are generally much weaker than other non-millisecond pulsars. This method does not inherently rely on the pulsar signal being extremely stable, because it only needs to measure the *change* in the frequency shift of the signal. Consequently, even if the frequency of the pulsar varies significantly from the nominal frequency, as long as it is stable during the course of the observation, the actual of the frequency is not important; all that matters that the derivative of the frequency observed at the spacecraft be estimated.

6.2.1 Algorithm Development

In outer space, the acceleration of a spacecraft is dominated by gravitational terms. Gravitational acceleration is a direct function of the spacecraft's position relative to celestial bodies. The acceleration may be written using Newton's law of universal gravitation as follows:

$$\mathbf{a}^A = f(\mathbf{r}^A, t) = \sum_i \frac{-Gm_i}{|\mathbf{r}^A - \mathbf{r}^i|^3} (\mathbf{r}^A - \mathbf{r}^i) \quad (6.1)$$




Figure 6.2: Acceleration magnitudes near the Earth-moon system. Orbits shown in dotted lines. Notable in this figure is the location between the Earth and the Sun at which the acceleration vector magnitude approaches zero, i.e. where the gravitational force of Earth is equal to and opposite that of the Sun.

This general relationship between acceleration and position may then be used to determine the relationship between acceleration along a given vector and position.

$$a_{\ell}^A = \ell \cdot f(\mathbf{r}^A, t) = \sum_i \ell \cdot \frac{-Gm_i}{|\mathbf{r}^A - \mathbf{r}^i|^3} (\mathbf{r}^A - \mathbf{r}^i) \quad (6.2)$$

Using this expression, we may generate “maps” of acceleration based on position. Such “maps” is shown in Figures 6.2, 6.3, and 6.4. These figures plot the magnitude of the acceleration vector as a function of position. Since a_{ℓ}^A is a vector we can generate




Figure 6.3: Acceleration magnitudes in the vicinity of the inner planets. Orbits shown in dotted lines.




Figure 6.4: Acceleration magnitudes in the vicinity of the outer planets. Orbits shown in dotted lines.

one such map for each component of the vector.

To use this information as a means of estimating position, we must be able to relate acceleration to a position coordinate. A number of estimation techniques could be used at this point. For this analysis, we will use a basic Extended Kalman Filter (EKF) measurement update approach to formulate the measurement update equations. Consequently, the measurement matrix is given by taking the Jacobian of the measurement equation.

$$\mathbf{H} = \begin{bmatrix} \frac{\partial a_{\ell}^A}{\partial r_x^A} & \frac{\partial a_{\ell}^A}{\partial r_y^A} & \frac{\partial a_{\ell}^A}{\partial r_z^A} \end{bmatrix} \quad (6.3a)$$

$$= \begin{bmatrix} \frac{\partial \ell \cdot f(\mathbf{r}^A, t)}{\partial r_x^A} & \frac{\partial \ell \cdot f(\mathbf{r}^A, t)}{\partial r_y^A} & \frac{\partial \ell \cdot f(\mathbf{r}^A, t)}{\partial r_z^A} \end{bmatrix} \quad (6.3b)$$

$$= \ell \cdot \sum_i Gm_i \begin{bmatrix} \frac{1}{|\mathbf{r}^A - \mathbf{r}^i|^3} + \frac{3(r_x^A - r_x^i)(r_y^A - r_y^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{3(r_x^A - r_x^i)(r_y^A - r_y^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{3(r_x^A - r_x^i)(r_z^A - r_z^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} \\ \frac{3(r_x^A - r_x^i)(r_y^A - r_y^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{1}{|\mathbf{r}^A - \mathbf{r}^i|^3} + \frac{3(r_y^A - r_y^i)(r_z^A - r_z^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{3(r_y^A - r_y^i)(r_z^A - r_z^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} \\ \frac{3(r_x^A - r_x^i)(r_z^A - r_z^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{3(r_y^A - r_y^i)(r_z^A - r_z^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} & \frac{1}{|\mathbf{r}^A - \mathbf{r}^i|^3} + \frac{3(r_z^A - r_z^i)(r_x^A - r_x^i)}{|\mathbf{r}^A - \mathbf{r}^i|^5} \end{bmatrix} \quad (6.3c)$$

This measurement matrix is essentially the gravity gradient tensor, projected on the line of sight to the signal source.

Using this measurement matrix, we may perform an analysis to determine what level of accuracy of position may be achieved, given a level of accuracy of acceleration measurement. The Fisher information matrix resulting from a single measurement update of the state with no prior estimate is given by:

$$\mathcal{J}^+ = \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} \quad (6.4)$$

We may solve this equation for \mathbf{R} , if the matrix \mathbf{H} is full rank.

$$\mathbf{R}^{-1} = \mathbf{H}^{-T} \mathcal{J}^+ \mathbf{H}^{-1} \quad (6.5)$$

Using this expression, we could determine a minimum *a posteriori* estimate accuracy (\mathcal{J}^+), and from that requirement determine the needed accuracy of an acceleration measurement (\mathbf{R}^{-1}) to achieve a given position accuracy may be determined. This




Figure 6.5: Map of acceleration required measurement accuracy near the Earth-moon system. Orbits shown in dotted lines.

type of analysis is used to generate Figures 6.5, 6.6 and 6.7. These contour plots show the required acceleration measurement accuracy needed to determine position to an accuracy of 1000 km. These plots were generated assuming that acceleration was measured with equal accuracy in three dimensions.

Examination of the figure indicates that acceleration must be measured to a very high accuracy in order to reach a position estimate of relatively low accuracy, with position error on the order of 1000 km. The required accuracies range from $1 \times 10^{-6} \frac{m}{s^2}$ near Earth to $1 \times 10^{-14} \frac{m}{s^2}$ in the outer reaches of the solar system. This makes sense on an intuitive level. Since the magnitude of the acceleration due to gravity is inversely proportional to the square of the distance between the objects, the relative change




Figure 6.6: Map of acceleration required measurement accuracy in the vicinity of the inner planets. Orbits shown in dotted lines.




Figure 6.7: Map of acceleration required measurement accuracy in the vicinity of the outer planets. Orbits shown in dotted lines.

in gravitational acceleration will diminish, especially far away from massive bodies. Consequently, an extremely accurate acceleration measurement is needed in order to determine any sort of useful position information.

It will be noted that this formulation of the measurement equation assumes an initial estimate of position, since the measurement equation is linearized about that initial estimate. Consequently, this formulation would not be particularly helpful in the “lost-in-space” scenario. However, this formulation allows us to develop initial accuracy requirements on the acceleration measurement, and to develop an idea of what types of position estimate accuracies could be expected given an acceleration measurement.

In a practical implementation, the position of the spacecraft could be resolved to a first order by neglecting other planets and only considering the gravity of the sun. This would allow a range and bearing angle to be determined directly from the magnitude and direction of the acceleration vector. From there, the EKF formulation could be used to refine the position estimate further.

It will also be noted that in some cases, the measured acceleration vector could be matched to more than one unique location. This is particularly likely in the case that the spacecraft is near a planet. This situation could be addressed by additionally estimating the gravitational gradient, or by simply comparing the rate of change of the acceleration vector and the estimated change in the spacecraft’s position.

We then note from Section 3.5.2 that it is possible to estimate the kinematic states of the spacecraft, including velocity and acceleration, jointly with the correlation estimator presented in this work.

6.2.2 Monte Carlo Results

To demonstrate the efficacy of the methods proposed here, a series of Monte Carlo simulations was conducted. The goal of the simulations was to answer whether or not it is feasible to determine gravitational acceleration to the accuracies needed in order to measure position. Furthermore, if it is possible, under what conditions might such a measurement be achieved? Of particular interest in this study was the required observation time for a detector of a size that could reasonably be flown aboard a spacecraft.

To answer these questions, the Monte Carlo simulations were conducted in the following manner. A model of the pulsar signal based on previous observations was used

Table 6.1: Signal Parameters

pulsar name	area (cm^2)	photon count rate $\left(\frac{photons}{second}\right)$	pulsar period ($seconds$)	pulsed fraction	background count rate $\left(\frac{photons}{second}\right)$
B1509-58	2.00e+02	2.45e-01	1.51e-01	8.77e-01	1.44e-01
B1509-58	4.00e+02	4.91e-01	1.51e-01	8.77e-01	2.89e-01

to generate photon arrival times, based on the area of the detector.

Additionally, background photon measurements were generated based on empirical models for hard x-ray background flux [48]. Details of the signals used are given in Table 6.1. These photon arrivals were used in the estimator presented above. The estimator was initialized with a poor initial estimate of velocity and acceleration. For the purposes of this study, acceleration and velocity were assumed to be constant. The parameters of the simulation are given in Table 6.2.

It should be noted that the standard deviation of the initial velocity and acceleration error includes a large fraction of feasible velocities and accelerations that a spacecraft could experience, especially during the cruise phase of a mission, which is the phase of the mission where XNAV techniques would be most likely to be used. Indeed, the range of velocity errors modeled by the standard deviations encompasses some of the highest velocities achieved by man-made spacecraft. Similarly, the acceleration error standard deviation encompasses all but the most extreme accelerations that a spacecraft might experience, for instance in a close solar flyby. Thus, while an initial estimate was used to initialize the filter, the errors on those estimates are so large that, for all practical purposes, the spacecraft is initialized with no initial knowledge of its position or velocity.

This procedure was repeated for a variety of runtimes ranging from hours to days. For each runtime, a number of simulations was performed, each time with a new set of generated photon observations and a new set of poor initial estimates. The final estimates were filtered based on whether the estimator’s standard deviation indicated it had converged to a solution. Note that this procedure was *not* a post-simulation rejection of outliers that could only be performed with knowledge of the true solution.

⁰Full input files and results may be viewed here (Record ID 20190415-181018, 20190415-191554). Scripts available on Python Package Index in the Modular Estimator package (pip install modest)

Table 6.2: Monte Carlo Parameters

pulsar name	area (cm^2)	initial velocity std dev ($\frac{km}{s}$)	initial acceleration std dev ($\frac{km}{s^2}$)	run count
B1509-58	2.00e+02	2.00e+01	1.00e-02	50
B1509-58	4.00e+02	2.00e+01	1.00e-02	50

Rather, it was a rejection of runs in which the estimator did not believe it had converged to a meaningful solution yet, based on the estimated covariance matrix values. This rejection could be performed by the spacecraft in real time; no knowledge of the true solution was required. The final estimates of acceleration were used to compute the standard deviation of the estimate error for each run-time. These standard deviations are shown in Figure 6.8.

The plots show three different sets data points. The solid line plots show the standard deviation of the acceleration estimate errors. This line reflects how accurate the estimator’s average acceleration solution was after a given run-time. The dashed lines show the mean of the estimated standard deviation of acceleration; this reflects how accurate the estimator believed its solution was after the given run-time. Finally, the scatter points are the absolute value of error for each run.

These results suggest a few conclusions. First, the results imply that, in the simplified case considered here, it is possible to estimate acceleration to a high degree of accuracy given a long enough observation time. It will further be noted that, especially after a 10 day run-time, it is possible to achieve a much lower error than the mean standard deviation.

6.2.3 *Chandra* Validation Experiment

In order to further validate the techniques proposed in this section, validation experiments were conducted using x-ray pulsar observation data taken from the *Chandra* x-ray observatory. The *Chandra* observations are well-suited for verifying the performance of this technique for a number of reasons. First, *Chandra*’s High Resolution Camera (HRC), when operating in timing mode, provides very high-resolution measurements of the photons’ time of arrival, which is the primary input to this filter. The




Figure 6.8: Standard deviation of acceleration estimate as a function of runtime

parameters of the HRC were given in Table 5.3.

Second, the orbit of the *Chandra* spacecraft is highly elliptical (eccentricity = 0.743972). This means that for long periods of time, the spacecraft's acceleration is dominated by acceleration due to the sun, with very small contributions from gravitational acceleration due to Earth. This acceleration is similar to the acceleration a spacecraft would experience in the cruise phase of a mission, which is when these algorithms are most likely to be used. Finally, *Chandra* observed multiple pulsars with relatively high x-ray photon fluxes, providing multiple candidate observations for analysis.

From *Chandra*'s observations, three observations were selected for analysis. For each of these observations, the filter described here this section was used on the photon time of arrival measurements to estimate correlation, velocity, and acceleration. As with the Monte Carlo analysis in the previous section, the filter was initialized with a poor initial estimate of velocity and acceleration. Photons were filtered based on their angle of arrival and *Chandra*'s reported orientation. This step was taken to help decrease the number of background photons incorporated into the estimate. The filter was run for the entire duration of each observation. It should be noted that the durations of the observations are less than one day, which is not long enough to reach the acceleration measurement accuracies needed for high accuracy position estimation.

The error between the true and estimated acceleration for each observation is shown in Figure 6.9. It should be noted that the error plotted is the absolute value of the error. This is done so that the error can be displayed on a log scale. Additionally, the estimated standard deviation is plotted in dotted lines.

There are a few results that should be noted from these plots. First, in no observation does the acceleration reach the level required to achieve even a very poor estimate of position. Even in the most favorable of scenarios, such as relatively near a large planet, these acceleration estimates would result in position estimates with accuracies on the order of 1 million kilometers.

While the observations as presented would not yield a high-accuracy position estimate, it is also notable that the accuracy achieved after a (relatively) short observation of less than a day matches the predicted accuracy of the Monte Carlo simulations after a similar run-time. This seems to suggest that, at least upon initial examination, the




Figure 6.9: Acceleration estimate errors for three *Chandra* observations

Monte Carlo simulations are a fairly accurate predictor of the accuracy which could be achieved given a detector size and observation time. If this is indeed the case, this suggests that a longer observation, on the order of 10 days, could in fact result in an estimate with much more reasonable error on the order of 1000 km. It is also notable that the detector size required to achieve this accuracy is only 400 cm².

Finally, it is notable that in each case, the estimated standard deviation roughly matches the magnitude of the error. This indicates that the relationship between photon TOA measurements and acceleration estimate error is well-modeled.

We believe that these initial studies show that the method discussed here has significant potential for addressing the “lost-in-space” problem. The ability to accurately measure acceleration using pulsars in space, combined with the “gravity-maps” presented here could be used by the spacecraft to localize its position to an accuracy sufficient for the spacecraft to solve the integer ambiguity problem, and to track the phase of much fainter millisecond pulsars using its more accurate velocity estimate. However, more work is needed to determine the feasibility of this approach, including Monte Carlo studies that more accurately model the real conditions of a spacecraft on an interplanetary trajectory. Additionally, we have not considered the effects of additional timing errors, including but not limited to the Shapiro delay and delays due to the interstellar medium. Further analysis is needed to determine the magnitude of these effects on the resulting acceleration estimate error.

6.3 Derivation of Lower Bounds on Error

Literature on XNAV frequently cites a lower bound for the accuracy of an XNAV position solution [9]. However in reality, it has been shown through simulation and analysis of experimentally collected data that this lower bound on error is overly optimistic; that is, it over-estimates the accuracy of the XNAV position solution. This work indicates that the failure of the model is due to its inaccurate representation of the non-Gaussian statistics of the signal [35]. Additional sources of inaccuracy include a failure of the bound to incorporate the effects of velocity and attitude uncertainty.

The derivation of a more accurate lower bound which addresses these short-comings and offers a more accurate expression of the expected accuracy of an XNAV solution is

an potential area of future research. Such a bound would include uncertainties in the spacecraft’s initial velocity, orbit model, and attitude solution, as well as detector area. A theoretical bound on the accuracy of an attitude estimate using x-ray signals could also be derived as a function of the detector’s area, angular resolution, and the drift of the spacecraft’s angular velocity measurements. Such bounds would be valuable tools in assessing the performance of XNAV algorithms, and would also help to inform the design of detectors for XNAV use.

6.4 Detector Hardware Development

The research presented here has assumed to a large degree the existence of an idealized detector. There is a large opportunity for research into the design of a cost-effective x-ray detector which is capable of meeting the requirements of the navigation techniques proposed here. In particular, the algorithms developed in this research could be used to determine the exact performance specifications needed from a detector to achieve a given navigational solution accuracy.

6.5 Use of Alternative Signal Sources

As discussed in Section 2.1, pulsars are one of many sources of signals available in space which could be potentially used for navigation. While pulsars offer a promising means of navigation, it may also be beneficial to continue exploring other signals that could be used in addition to pulsar signals. The algorithms needed to use asynchronous signal sources for navigation would most likely involve cooperative navigation techniques, which presents a new set of research challenges. In particular, the cooperative spacecraft would need to a way to share information in an efficient manner while still transmitting the minimum necessary amount of data to successfully navigate. Additionally, the spacecraft would need to address the problem of decentralized estimation, and the problem of unknown correlations between errors.

Nomenclature

Signals and Measurements

S^i i^{th} signal source

State Estimation

r_t^A Position of A, at time t (km)

a_t^A Acceleration of object A, at time t (km/s^2)

t Current time (*seconds*)

\mathbf{x} State vector

$y_{t,a}$ Measurement of scalar quantity a

\mathbf{H} Measurement update matrix for measured quantity a

$v_{t,a}$ Scalar measurement error of measured quantity a at time t

$\mathbf{R}_{t,a}$ Variance of measurement noise in measured quantity a

$[\cdot]^-$ Indicates *a priori* value

$[\cdot]^+$ Indicates *a posteriori* value

$\hat{\cdot}$ Indicates estimated value

\mathbf{a} Indicates vector quantity

a Indicates scalar quantity

Misc

A, B Indices to designate spacecraft(s)

$\ell^{(S^i)}$ Unit vector pointing towards i^{th} signal source

References

- [1] Gregor Reisch. *Margarita Philosophica*. Freiburg im Breisgau, Germany, 1504.
- [2] Peter Apia. *Cosmographia Petri Apiani*. Parisiis, vaneunt apud Viuantium Gaultherot, via Iacobeae: sub intersignio D. Martini, 1553.
- [3] National Aeronautics and Space Administration. NASA Technology Roadmaps TA 5: Communications, Navigation, and Orbital Debris Tracking and Characterization Systems. *NASA Technology Roadmaps*, (July), 2015.
- [4] D S Bagri and W A Majid. Estimating accurate relative spacecraft angular position from deep space network very long baseline interferometry phases using X-band telemetry or differential one-way ranging tones. *The Interplanetary Network Progress Report*, 42(172):1–10, 2008.
- [5] James K Miller, Dale R Stanbridge, and Bobby G Williams. New Horizons Pluto Approach Navigation. In *AAS Guidance and Control Conference*, pages 1–12, Breckenridge, CO, 2004.
- [6] David Ivan Vilaseca and Juan Ignacio Giribet. Indoor navigation using WiFi signals. In *Embedded Systems (SASE/CASE), 2013 Fourth Argentine Symposium and Conference on*, pages 1–6. IEEE, 2013.
- [7] Trevor Layh, Jordan Larson, John Jackson, Yunus Agamawi, Brian Taylor, Demoz Gebre-Egziabher, Joel Krause, and Jacob Gustafson. Cell Phone Navigation. Technical report, University of Minnesota, 2015.
- [8] A Robert (Johns Hopkins University Applied Physics Laboratory) Golshan and Inc.) Sheikh, Suneel I (ASTER Labs. On pulse phase estimation and tracking of

- variable celestial X-ray sources. In *Institute of Navigation 63rd Annual Meeting*, number April, pages 413–422, Cambridge, MA, 2007. Institute of Navigation.
- [9] Suneel I. Sheikh, Darryll J. Pines, Paul S. Ray, Kent S. Wood, Michael N. Lovellette, and Michael T. Wolff. Spacecraft Navigation Using X-Ray Pulsars. *Journal of Guidance, Control, and Dynamics*, 29(1):49–63, jan 2006.
 - [10] Chuck S Hisamoto and Suneel I Sheikh. Spacecraft Navigation Using Celestial Gamma-Ray Sources. *Journal of Guidance, Control, and Dynamics*, 38(9):1765–1774, sep 2015.
 - [11] G.S. Downs. Technical Report 32-1594: Interplanetary Navigation Using Pulsating Radio Sources. Technical report, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 1974.
 - [12] Amir A Emadzadeh, Cassio G Lopes, and Jason L Speyer. Online time delay estimation of pulsar signals for relative navigation using adaptive filters. In *2008 IEEE/ION Position, Location and Navigation Symposium*, pages 714–719. IEEE, 2008.
 - [13] Amir A Emadzadeh and Jason L Speyer. A new relative navigation system based on X-ray pulsar measurements. In *2010 IEEE Aerospace Conference*, pages 1–8. IEEE, IEEE, mar 2010.
 - [14] Amir Abbas Emadzadeh and Jason L Speyer. Relative Navigation Between Two Spacecraft Using X-ray Pulsars. *IEEE Transactions on Control Systems Technology*, 19(5):1021–1035, sep 2011.
 - [15] Amir A. Emadzadeh and Jason L. Speyer. X-Ray Pulsar-Based Relative Navigation using Epoch Folding. *IEEE Transactions on Aerospace and Electronic Systems*, 47(4):2317–2328, 2011.
 - [16] Kevin D Anderson, Darryll J Pines, and Suneel I Sheikh. Validation of pulsar phase tracking for spacecraft navigation. *Journal of Guidance, Control, and Dynamics*, 38(10):1885–1897, 2015.

- [17] Luke M B Winternitz, Munther A Hassouneh, Jason W Mitchell, Jennifer E Valdez, Samuel R Price, Sean R Semper, Wayne H Yu, Paul S Ray, Kent S Wood, Zaven Arzoumanian, Keith C Gendreau, H Yu Wayne, Paul S Ray, Kent S Wood, Zaven Arzoumanian, and Others. X-ray pulsar navigation algorithms and testbed for SEXTANT. In *2015 IEEE Aerospace Conference*, pages 1–14, Big Sky, MT, USA, mar 2015. IEEE, IEEE.
- [18] Jason W Mitchell, Luke B Winternitz, Munther A Hassouneh, Samuel R Price, Sean R Semper, Wayne H Yu, Paul S Ray, Michael T Wolff, Matthew Kerr, Kent S Wood, Zaven Arzoumanian, Keith C Gendreau, and Paul Demorest. SEXTANT X-Ray Pulsar Navigation Demonstration: Initial On-Orbit Results. In *41st Annual American Astronautical Society (AAS) Guidance and Control Conference*, pages 1–12, Breckenridge, CO, 2018.
- [19] Luke M B Winternitz, Jason W Mitchell, Munther A Hassouneh, Jennifer E Valdez, Samuel R Price, Sean R Semper, Wayne H Yu, Keith C Gendreau, Paul S Ray, and Kent S Wood. SEXTANT X-ray Pulsar Navigation Demonstration : Flight System and Test Results. In *2016 IEEE Aerospace Conference*, pages 1–11. IEEE, 2016.
- [20] Science Mission Directorate) (National Aeronautics and Space Administration. Gamma Rays, 2010.
- [21] Chryssa Kouveliotou, Charles A. Meegan, Gerald J. Fishman, Narayana P. Bhat, Michael S. Briggs, Thomas M. Koshut, William S. Paciesas, and Geoffrey N. Pendleton. Identification of Two Classes of Gamma-Ray Bursts. *The Astronomical Journal*, 413(2):L101–L104, 1993.
- [22] G.F. Marani, R.J. Nemiroff, J. Norris, and J.T. Bonnell. On Similarities among GRBs. In *American Astronomical Society Meeting Abstracts #190*, page 839, 1997.
- [23] R. J. Gould and G. R. Burbidge. High Energy Cosmic Photons and Neutrinos. *Annales d’Astrophysique*, 28:171, 1965.
- [24] Didier Barret. The broad band x-ray/hard x-ray spectra of accreting neutron stars. *Advances in Space Research*, 28(2-3):307–321, 2001.

- [25] Walter H. G Lewin and Paul C. Joss. X-ray Bursters and the X-Ray Sources of the Galactic Bulge. *Space Science Reviews*, 28(1):3–87, 1981.
- [26] Ari Laor, Fabrizio Fiore, Martin Elvis, Belinda J. Wilkes, and Jonathan C. McDowell. The Soft XRay Properties of a Complete Sample of Optically Selected Quasars. II. Final Results. *The Astrophysical Journal*, 477(1):93–113, 1997.
- [27] A A Abdo, M Ackermann, M Ajello, A Allafort, E Antolini, W B Atwood, M Axelsson, L Baldini, J Ballet, G Barbiellini, D Bastieri, B M Baughman, K Bechtol, R Bellazzini, F Belli, B Berenji, D Bisello, R D Blandford, E D Bloom, E Bonamente, J Bonnell, A W Borgland, A Bouvier, J Bregeon, A Brez, M Brigida, P Bruel, T H Burnett, G Busetto, S Buson, G A Caliandro, R A Cameron, R Campana, B Canadas, P A Caraveo, S Carrigan, J M Casandjian, E Charles, A Chekhtman, C C Cheung, J Chiang, E Cavazzuti, M Ceccanti, C Cecchi, S Ciprini, R Claus, J Conrad, R Corbet, D S Davis, M Deklotz, A N Cillis, P R Den Hartog, C D Dermer, A De Angelis, A De Luca, F De Palma, S W Digel, M Dormody, P S Drell, R Dubois, D Dumora, D Fabiani, C Farnier, C Favuzzi, S J Fegan, E C Ferrara, W B Focke, P Fortin, M Frailis, Y Fukazawa, S Funk, P Fusco, F Gargano, D Gasparrini, N Gehrels, S Germani, G Giavitto, B Giebels, N Giglietto, P Giommi, F Giordano, M Giroletti, T Glanzman, G Godfrey, I A Grenier, J E Grove, L Guillemot, S Guiriec, M Gustafsson, D Hadasch, Y Hanabata, A K Harding, M Hayashida, A S Johnson, E Hays, S E Healey, A B Hill, D Horan, R E Hughes, G Iafate, W N Johnson, T Kamae, H Katagiri, J Kataoka, N Kawai, M Kerr, R P Johnson, T J Johnson, D Kocevski, M Kuss, J Lande, D Landriu, L Latronico, A M Lionetto, M Llana Garde, F Longo, F Loparco, B Lott, M N Lovellette, P Lubrano, G M Madejski, A Makeev, B Marangelli, M Marelli, E Marsaro, M N Mazziotta, W Mcconville, J E McEnery, P F Michelson, M Minuti, W Mitthumsiri, T Mizuno, A A Moiseev, M Mongelli, C Monte, M E Monzani, E Moretti, A Morselli, I V Moskalenko, S Murgia, H Nakajima, T Nakamori, P L Nolan, J P Norris, E Nuss, M Ohno, T Ohsugi, N Omodei, E Orlando, J F Ormes, M Ozaki, A Paccagnella, D Paneque, J H Panetta, D Parent, V Pelassa, M Pepe, M Pinchera, F Piron, T A Porter, L Poupard, P S Ray, M Razzano, S Razzaque, N Rea, A Reimer, O Reimer, T Reposeur, R Rando, J Ripken, S Ritz,

L S Rochester, A Y Rodriguez, R W Romani, M Roth, D Salvetti, D Sanchez,
 A Sander, P M Saz Parkinson, J D Scargle, T L Schalk, G Scolieri, P D Smith,
 G Spandre, P Spinelli, M S Shaw, E J Siskind, D A Smith, T E Stephens, E Striani,
 M S Strickman, A W Strong, D J Suson, H Tajima, H Takahashi, T Takahashi,
 T Tanaka, J B G Thayer, J B G Thayer, D J Thompson, L Tibaldo, O Tibolla,
 F Tinebra, D F Torres, G Tosti, A Tramacere, Y Uchiyama, T L Usher, A Van
 Etten, V Vasileiou, N Vilchez, V Vitale, A P Waite, E Wallace, P Wang, K Wat-
 ters, B L Winer, K S Wood, Z Yang, T Ylinen, M Ziegler, {Ö} {Ç}elik, E Charles,
 A Chekhtman, C C Cheung, J Chiang, A N Cillis, S Ciprini, R Claus, J Cohen-
 Tanugi, J Conrad, R Corbet, D S Davis, M Deklotz, P R den Hartog, C D Dermer,
 A de Angelis, A de Luca, F de Palma, S W Digel, M Dormody, E do Couto e Silva,
 P S Drell, R Dubois, D Dumora, D Fabiani, C Farnier, C Favuzzi, S J Fegan, E C
 Ferrara, W B Focke, P Fortin, M Frailis, Y Fukazawa, S Funk, P Fusco, F Gargano,
 D Gasparri, N Gehrels, S Germani, G Giavitto, B Giebels, N Giglietto, P Giommi,
 F Giordano, M Giroletti, T Glanzman, G Godfrey, I A Grenier, M.-H. Grondin, J E
 Grove, L Guillemot, S Guiriec, M Gustafsson, D Hadasch, Y Hanabata, A K Hard-
 ing, M Hayashida, E Hays, S E Healey, A B Hill, D Horan, R E Hughes, G Iafrate,
 G J{ó}hannesson, A S Johnson, R P Johnson, T J Johnson, W N Johnson, T Ka-
 mae, H Katagiri, J Kataoka, N Kawai, M Kerr, J Kn{ö}dlseder, D Kocevski,
 M Kuss, J Lande, D Landriu, L Latronico, S.-H. Lee, M Lemoine-Goumard, A M
 Lionetto, M Llena Garde, F Longo, F Loparco, B Lott, M N Lovellette, P Lubrano,
 G M Madejski, A Makeev, B Marangelli, M Marelli, E Massaro, M N Mazziotta,
 W Mcconville, J E McEnery, P F Michelson, M Minuti, W Mitthumsiri, T Mizuno,
 A A Moiseev, M Mongelli, C Monte, M E Monzani, E Moretti, A Morselli, I V
 Moskalenko, S Murgia, H Nakajima, T Nakamori, M Naumann-Godo, P L Nolan,
 J P Norris, E Nuss, M Ohno, T Ohsugi, N Omodei, E Orlando, J F Ormes, M Ozaki,
 A Paccagnella, D Paneque, J H Panetta, D Parent, V Pelassa, M Pepe, M Pesce-
 Rollins, M Pinchera, F Piron, T A Porter, L Poupard, S Rain{ò}, R Rando, P S
 Ray, M Razzano, S Razzaque, N Rea, A Reimer, O Reimer, T Reposeur, J Ripken,
 S Ritz, L S Rochester, A Y Rodriguez, R W Romani, M Roth, H F.-W. Sadrozin-
 ski, D Salvetti, D Sanchez, A Sander, P M Saz Parkinson, J D Scargle, T L Schalk,
 G Scolieri, C Sgr{ò}, M S Shaw, E J Siskind, D A Smith, P D Smith, G Spandre,

- P Spinelli, J.-L. Starck, T E Stephens, E Striani, M S Strickman, A W Strong, D J Suson, H Tajima, H Takahashi, T Takahashi, T Tanaka, J B G Thayer, J B G Thayer, D J Thompson, L Tibaldo, O Tibolla, F Tinebra, D F Torres, G Tosti, A Tramacere, Y Uchiyama, T L Usher, A Van Etten, V Vasileiou, N Vilchez, V Vitale, A P Waite, E Wallace, P Wang, K Watters, B L Winer, K S Wood, Z Yang, T Ylinen, and M Ziegler. Fermi Large Area Telescope First Source Catalog. *The Astrophysical Journal Supplement Series*, 188(2):405–436, 2010.
- [28] L Guillemot, T J Johnson, C Venter, M Kerr, B Pancrazi, M Livingstone, and G H Janssen. Pulsed Gamma Rays from the Original Millisecond and Black Widow Pulsars: A Case for Caustic Radio Emission? *The Astrophysical Journal*, 33(1):33, 2012.
- [29] T J Johnson, L Guillemot, M Kerr, I Cognard, P S Ray, M T Wolff, and G H Janssen. Broadband Pulsations from PSR B1821-24: Implications for Emission Models and the Pulsar Population of M28. *The Astrophysical Journal*, 778(2):106, 2013.
- [30] R.N. Manchester and H.H Taylor. *Pulsars*. W H Freeman and Company, San Francisco, 1977.
- [31] Demetrios Nicholas Matsakis, Joseph Hooton Taylor, and T Marshall Eubanks. A statistic for describing pulsar and clock stabilities. *Astronomy and Astrophysics*, 326:924–928, 1997.
- [32] G Jacovitti and G Scarano. Discrete time techniques for time delay estimation. *IEEE Transactions on Signal Processing*, 41(2):525–533, 1993.
- [33] Charles H. Knapp and G. Clifford Carter. The Generalized Correlation Method for Estimation of Time Delay. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 24(4):320–327, 1976.
- [34] Simon Haykin. Adaptive filter theory Chapter 2. In *Adaptive Filter Theory*, chapter Chapter 2. Pearson Education, Inc., Upper Saddle River, New Jersey, 5th edition, 2013.

- [35] Joel T. Runnels and Demoz Gebre-Egziabher. Recursive Range Estimation Using Astrophysical Signals of Opportunity. *Journal of Guidance, Control, and Dynamics*, 40(9):1–13, 2017.
- [36] Bradford W Parkinson and James J Spilker. Global Positioning System: Theory and Applications. In *Global Positioning System: Theory and Applications*, volume 1, chapter Chapter 7. American Institute of Aeronautics and Astronautics, Inc., 370 L’Efant Promenade, SW, Washington, DC 20024-2518, 5th edition, 1996.
- [37] Jr. Harnden, F. R. and F D Seward. Einstein observations of the Crab nebula pulsar. *The Astrophysical Journal*, 283:279, aug 1984.
- [38] Simon Haykin. Adaptive filter theory Chapter 14. In *Adaptive Filter Theory*, chapter Chapter 14. Pearson Education, Inc., Upper Saddle River, New Jersey, 5th edition, 2013.
- [39] Dan Simon. *Optimal state estimation: Kalman, H and nonlinear approaches*. Wiley-Interscience, 1 edition, 2006.
- [40] Wu Sheng Lu and Tian Bo Deng. An improved weighted least-squares design for variable fractional delay fir filters. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 46(8):1035–1040, 1999.
- [41] T.I. Laakso, V. Valimaki, and M. Karjalainen. Splitting the unit delay. *IEEE Signal Processing Magazine*, 13(1):30–60, 1996.
- [42] Håkan Johansson and Per Löwenborg. On the Design of Adjustable Fractional Delay FIR Filters. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 50(4):164–169, 2003.
- [43] J Vesma and T Saramaki. Optimization and efficient implementation of FIR filters with adjustable fractional delay. In *Proceedings of 1997 IEEE International Symposium on Circuits and Systems. Circuits and Systems in the Information Age ISCAS ’97*, volume 4, pages 2256—2259 vol.4. IEEE, 1997.
- [44] John Junkins, John Crassidis, and Sequential State Estimation. Unscented Filtering. In *Optimal Estimation of Dynamic Systems, Second Edition*, Chapman &

- Hall/CRC Applied Mathematics & Nonlinear Science, pages 192–199. Chapman and Hall/CRC, Boca Raton, FL, 2nd edition, oct 2011.
- [45] Jules P Halpern, Christopher Martin, and Herman L Marshall. Soft X-Ray Properties of the Binary Millisecond Pulsar J0437-4715. *The Astrophysical Journal*, 462:908, may 1996.
 - [46] Tadayuki Takahashi, Keiichi Abe, Manabu Endo, Yasuhiko Endo, Yuuichiro Ezoe, Yasushi Fukazawa, Masahito Hamaya, Shinya Hirakuri, Soojing Hong, Michihiro Horii, Hokuto Inoue, Naoki Isobe, Takeshi Itoh, Naoko Iyomoto, Tuneyoshi Kamae, Daisuke Kasama, Jun Kataoka, Hiroshi Kato, Madoka Kawaharada, Naomi Kawano, Kengo Kawashima, Satoshi Kawasoe, Tetsuichi Kishishita, Takao Kitaguchi, Yoshihito Kobayashi, Motohide Kokubun, Jun’ichi Kotoku, Manabu Kouda, Aya Kubota, Yoshikatsu Kuroda, Greg Madejski, Kazuo Makishima, Kazunori Masukawa, Yukari Matsumoto, Takefumi Mitani, Ryohei Miyawaki, Tsunefumi Mizuno, Kunishiro Mori, Masanori Mori, Mio Murashima, Toshio Murakami, Kazuhiro Nakazawa, Hisako Niko, Masaharu Nomachi, Yuu Okada, Masanori Ohno, Kousuke Oonuki, Naomi Ota, Hideki Ozawa, Goro Sato, Shingo Shinoda, Masahiko Sugiho, Masaya Suzuki, Koji Taguchi, Hiromitsu Takahashi, Isao Takahashi, Shin’ichiro Takeda, Ken-ichi K.-i. Tamura, Takayuki Tamura, Takaaki Tanaka, Chiharu Tanihata, Makoto Tashiro, Yukikatsu Terada, Shin’ya Tominaga, Yasunobu Uchiyama, Shin Watanabe, Kazutaka Yamaoka, Takayuki Yanagida, and Daisuke Yonetoku. Hard X-Ray Detector (HXD) on Board Suzaku. *Publications of the Astronomical Society of Japan*, 59(sp1):S35–S51, jan 2007, 0611232.
 - [47] David N. Burrows. The Swift X-ray telescope. *Space Science Reviews*, 120(3-4):165–195, 2005, 0508071.
 - [48] D. E. Gruber, J. L. Matteson, L. E. Peterson, and G. V. Jung. The Spectrum of Diffuse Cosmic Hard X-Rays Measured with HEAO 1. *The Astrophysical Journal*, 520(1):124–129, 1999, 9903492.
 - [49] Trevor Knuth, Tim Kukowski, Maxwell Yurs, Kendra Bergstedt, Jeffrey Chaffin, Samuel Drechsel, Abigail Valero, Lindsay Glesener, Demoz Gebre-Egziabher, and

- Ryan Vogt. Developing a detector model for the experiment for x-ray characterization and timing (EXACT) CubeSat. (August 2017):49, 2017.
- [50] Josiah DeLange, Seth Frick, Joel Runnels, Demoz Gebre-Egziabher, and Kale Hedstrom. Sensor for small satellite relative PNT in deep-space. In *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 955–963. IEEE, 2016.
 - [51] NASA. Suzaku Mission Overview, 2006.
 - [52] HXD Team. Screening criteria for the HXD data in the version 2.0 processing, 2007.
 - [53] Carl Christian Liebe. Star Trackers for Attitude Determination. *IEEE AES Systems Magazine*, (June):10–16, 1995.
 - [54] Carl Christian Liebe. Accuracy Performance of Star Trackers A Tutorial. *IEEE transactions on Aerospace and Electronic Systems*, 38(2):587–599, 2002.
 - [55] Carl Christian Liebe, Leon Akalai, George Domingo, Bruce Hancock, Don Hunter, Jeff Mellstrom, and Ian Ruiz. Micro APS Based Star Tracker. In *Proceedings of 2002 IEEE Aerospace Conference*, pages 2285–2300, Big Sky, MT, USA, 2002.
 - [56] J.L. (Texas A&M University) Ju, G. Kim, H. Pollock, T. Junkins. DIGISTAR: A Low Cost Micro Star Tracker. In *AIAA Space Technology Conference & Exposition*, number September, Albuquerque, NM, 1999.
 - [57] Allan Y Lee and Thomas A Burk. Cassini Spacecraft Attitude Control System: Flight Performance and Lessons Learned, 1997-2017. In *AIAA Guidance, Navigation, and Control Conference*, number January, pages 1–33, Kissimmee, Florida, 2018.
 - [58] California Institute of Technology) Burk, Thomas A (Jet Propulsion Laboratory. Cassini at Saturn Proximal Orbits Attitude Control Challenges. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, Boston, MA, 2013.
 - [59] David G. Hoag. Apollo Navigation, Guidance, and Control Systems: A Progress Report. Technical report, Massachusetts Institute of Technology, 1969.

- [60] F. Landis Markley and John L. Crassidis. *Fundamentals of spacecraft attitude determination and control*. Springer Science+Business Media, New York, 2014.
- [61] Yaakov Bar-Shalom, Fred Daum, and Jim Huang. The probabilistic data association filter. *IEEE Control Systems Magazine*, 29(6):82–100, 2009.
- [62] Joel Runnels and Demoz Gebre-Egziabher. Inertially Aided Vector Matching for Opportunistic Navigation in Space. In *Proc. of 68th International Astronautical Congress*, Adelaide, Australia, 2017.
- [63] P. A.W. Lewis and G. S. Shedler. Simulation of Nonhomogeneous Poisson Processes By Thinning. *Naval research logistics quarterly*, 26(3):403–413, 1979.
- [64] Eric W. Weisstein. "Solid Angle." From MathWorld—A Wolfram Web Resource.
- [65] Chandra X-ray Center. HRC Timing Issues, 2018.
- [66] R. S. Warwick, R. D. Saxton, and A. M. Read. The XMM-Newton slew survey in the 210keV band . *Astronomy & Astrophysics*, 548(i):A99, 2012, 1210.3992.
- [67] P. J. G. Teunissen. Least-Squares Estimation of the Integer GPS Ambiguities. In *Invited Lecture for Section IV "Theory and Methodology" at the General Meeting of the International Association of Geodesy, Beijing, China, August 1993*, number August, page 16, 1993.

Appendix A

Alternate Derivation of NHEP

Equation (2.18a) gives the analytical expression for the PDF of a nonhomogeneous Poisson process, and was derived based on the expected number of photons generated by an NHEP over a given time period. Here, we present an alternate derivation which gives an intuitive understanding of the NHEP and also provides a straight-forward method of simulating such a process.

We begin by noting that the non-homogeneous Poisson process may be thought of as the combination of a homogeneous Poisson process and a Bernoulli random variable, in which photons are generated at a constant flux, $\lambda^* \leq \max \lambda(t)$, and are deleted with probability $1 - \frac{\lambda^*}{\lambda(t)}$ [63]. Alternatively, we could say that photons are generated at constant flux λ^* , and transmitted with probability $\frac{\lambda^*}{\lambda(t)}$.

Such a process could be physically realized with the setup illustrated in Figure A.1.




Figure A.1: Generation of NHPP by thinning

It is intuitively clear that such a process would result in photons governed by the NHPP with flux $\lambda(t)$. We now show this analytically.

To demonstrate the equivalency between two distributions, we must show that the distributions have the same CDF. Consequently, we want to compute the probability that the first photon generated by the process described above will be less than or equal to a given time. To start, we define the following:

- N_T is the number of photons generated by the “thinned,” non-homogeneous exponential process over a given time-interval, T
- N_T^* is the number of photons generated by the “un-thinned,” homogeneous exponential process over a given time-interval, T
- p_t is the Bernoulli random variable at time t . If it $p_t = 0$, the event at time t (if present) is deleted. If $p_t = 1$, the event at time t (if present) is not deleted

As before, we note that the following is true:

$$\mathbf{Pr}[\mathcal{t} \leq t | \lambda(t)] = 1 - \mathbf{Pr}[\mathcal{t} > t | \lambda(t)] \quad (\text{A.1})$$

So, if we can compute the probability that the first arrival time will be greater than some time t , we can compute the CDF for this process. To compute this probability, we need to find the probability of no photons arriving at each differential time-step between s and t . (For convenience in this derivation we will let $s = 0$.)

We note that, at some infinitely small time-step, Δt , no photons occur if one of the following two events occur:

- No event is generated by the homogeneous Poisson process, i.e. $N_{\Delta t} = 0$, or
- The Bernoulli random variable is zero, i.e. $p_{\Delta t} = 0$

Since either condition is sufficient for no photon to be generated by the combination of the two process, we may write:

$$\begin{aligned}
\mathbf{Pr}[N_{\Delta T} = 0] &= \mathbf{Pr}[N_{\Delta t} = 0 \cup p_{\Delta t} = 0] \\
&= \mathbf{Pr}[N_{\Delta t} = 0] + \mathbf{Pr}[p_{\Delta t} = 0] - \mathbf{Pr}[N_{\Delta t} = 0] \mathbf{Pr}[p_{\Delta t} = 0] \\
&= e^{-\lambda^* \Delta t} + \left(1 - \frac{\lambda(t)}{\lambda^*}\right) - e^{-\lambda^* \Delta t} \left(1 - \frac{\lambda(t)}{\lambda^*}\right) \\
&= 1 - \frac{\lambda^*}{\lambda(t)} + e^{-\lambda^* \Delta t} \frac{\lambda(t)}{\lambda^*} \\
&= 1 - \left(1 - e^{-\lambda^* \Delta t}\right) \frac{\lambda(t)}{\lambda^*}
\end{aligned} \tag{A.2}$$

This expression is only valid for an infinitesimally small time-step, such that the time-varying flux $\lambda(t)$ is constant. To compute the probability that no photon arrives over the interval t , we must compute the probability that no photons arrive at each infinitesimal time-step.

$$\begin{aligned}
\mathbf{Pr}[t > t | \lambda(t)] &= \mathbf{Pr}[N_{\Delta T_1} = 0 \cap N_{\Delta T_2} = 0 \dots \cap N_{\Delta T_N} = 0] \\
&= \mathbf{Pr}[N_{\Delta T_1} = 0] \mathbf{Pr}[N_{\Delta T_2} = 0] \dots \mathbf{Pr}[N_{\Delta T_N} = 0] \\
&= \lim_{\Delta T \rightarrow 0} \prod_{i=0}^{\frac{t}{\Delta T}} \mathbf{Pr}[N_{\Delta T} = 0] \\
&= \lim_{\Delta T \rightarrow 0} \prod_{i=0}^{\frac{t}{\Delta T}} \left(1 - \left(1 - e^{-\lambda^* \Delta t}\right) \frac{\lambda(i \Delta t)}{\lambda^*}\right)
\end{aligned} \tag{A.3}$$

To proceed from here, we recall that the Taylor series for e^x is given by

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots \tag{A.4}$$

Consequently, we can approximate e^x near $x = 0$ as $e^x \approx 1 + x + \mathcal{O}(x^2)$. Since in this case, we are taking the limit as the exponent goes to zero, the error in the approximation also goes to zero. Therefore the probability may be re-written as:

$$\begin{aligned}
\Pr[t > t|\lambda(t)] &= \lim_{\Delta T \rightarrow 0} \prod_{i=0}^{\frac{t}{\Delta T}} \left(1 - (1 - (1 - \lambda^* \Delta t)) \frac{\lambda(i\Delta t)}{\lambda^*} \right) \\
&= \lim_{\Delta T \rightarrow 0} \prod_{i=0}^{\frac{t}{\Delta T}} \left(1 - \lambda^* \frac{\lambda(i\Delta t) \Delta t}{\lambda^*} \right) \\
&= \lim_{\Delta T \rightarrow 0} \prod_{i=0}^{\frac{t}{\Delta T}} (1 - \lambda(i\Delta t) \Delta t)
\end{aligned} \tag{A.5}$$

Recalling rules for logarithms and exponents, this expression may be rewritten as an infinite sum:

$$\begin{aligned}
\Pr[t > t|\lambda(t)] &= \exp \left[\lim_{\Delta T \rightarrow 0} \ln \prod_{i=0}^{\frac{t}{\Delta T}} (1 - \lambda(i\Delta t) \Delta t) \right] \\
&= \exp \left[\lim_{\Delta T \rightarrow 0} \sum_{i=0}^{\frac{t}{\Delta T}} \ln (1 - \lambda(i\Delta t) \Delta t) \right]
\end{aligned} \tag{A.6}$$

Once again, we make use of the Taylor series, but this time for natural logs. Note that:

$$\ln(1 - x) = -x - \frac{x^2}{2} - \frac{x^3}{3} - \dots \approx -x + \mathcal{O}(x^2) \tag{A.7}$$

Again, since we are evaluating the limit as $x \rightarrow 0$, the $\mathcal{O}(x^2)$ error term also goes to zero. Therefore, the expression may be written again as:

$$\Pr[t > t|\lambda(t)] = \exp \left[\lim_{\Delta T \rightarrow 0} \sum_{i=0}^{\frac{t}{\Delta T}} -\lambda(i\Delta t) \Delta t \right] \tag{A.8}$$

The argument of the exponent in this expression is the Riemann sum for the integral,

which by definition is:

$$\lim_{\Delta T \rightarrow 0} \sum_{i=0}^{\frac{t}{\Delta T}} -\lambda(i\Delta t) \Delta t = - \int_0^t \lambda(t) dt \quad (\text{A.9})$$

And so:

$$\Pr[t > t | \lambda(t)] = e^{-\int_0^t \lambda(t) dt} \quad (\text{A.10})$$

And so, we arrive at the same CDF as in Equation (2.17), which implies that the process described here is governed by the same probability distribution as the NHEP derived in Section 2.2.3.

$$\Pr[t \leq t | \lambda(t)] = 1 - e^{-\int_0^t \lambda(t) dt} \quad (\text{A.11})$$

The preceding derivation is useful not only because it presents an intuitive understanding of the NHPP and NHEP, but also because it leads to a simple method for simulating the arrival times generated by such a process. It is straight-forward to generate numbers from an homogeneous exponential distribution and Bernoulli distribution, which may then be used to generate NHEP arrival times, even in the case where the function $\lambda(t)$ is not readily available.

Nomenclature

Data Association

p Indices to designate signal source

Probability and Statistics

$\Pr[A]$ Probability of event A occurring

$\mathbb{E}[a]$ The expected value of random variable a

Signals and Measurements

$\tilde{\Lambda}_{A,k}$ The number of photons measured at location A over time-step k
(*Photons*)

$\lambda_k^{(p)}, \lambda^{(p)}(t)$ Expected number of photons at time index k , or expected photon flux at time t , from signal source p (*photons, photons/second*)

t_i Time-of-arrival measurement of photon i (*seconds*)

State Estimation

j, k Time-step indices

t Current time (*seconds*)

\mathbf{x} State vector

$[\cdot]^-$ Indicates *a priori* value

$[\cdot]^+$ Indicates *a posteriori* value

$\hat{\cdot}$ Indicates estimated value

\mathbf{a} Indicates vector quantity

a Indicates scalar quantity

Misc

i Iterator to designate summation index