Trajectory Estimation of an Optically Radiating Source

by

Craig Andera

Submitted to the
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Bachelor of Science in Electrical Science and Engineering

and

Master of Engineering in Electrical Science and Engineering

at the Massachusetts Institute of Technology

May 1995

Author _____
Department of Electrical Engineering and Computer Science
Monday, May 22, 1995

Certified by _____
Kurt Nordhaus
Engineering Fellow, Loral Infrared and Imaging Systems

Certified by _____
Jacob White
Professor of Electrical Engineering, Massachusetts Institute of Technology

Accepted by _____
F. R. Morgenthaler
Chairman, Department Committee on Graduate Students

Thesis Release Letter

Professor F. R. Morgenthaler

Department of Electrical Engineering and Computer Science

Room 38-476, M.I.T.

Cambridge, MA 01239

Subject: Master of Engineering Thesis of Craig Andera

Dear Professor Morgenthaler:

I have reviewed the attached thesis of Craig Andera on behalf of Loral Infrared and Imaging Systems. The thesis is within the scope of the thesis proposal as previously approved and does not contain any material that is objectionable to Loral. It is also approved for its technical content.

It is understood that the actual thesis document will be the permanent property of M.I.T. and will be placed in the M.I.T. Library within one month after the date of submission. Loral agrees that M.I.T. shall have the nonexclusive right to reproduce, publish, and distribute the thesis.

Authorized Official of Loral Infrared and Imaging Systems

Acknowledgments

There are many, many people to thank, and I don't think I could list them in order of importance if I tried. So the following are in no particular order, other than that this is how I felt like writing it.

First of all, I have to thank Kurt Nordhaus and Prof. Jacob White, who were my company and MIT thesis advisors (respectively, as if the "Prof" wasn't clue enough). Both of them were of great help to me during the development of this work, pushing me off in a different direction when I was hitting my head against the wall.

The entire SIP lab gets a hand as well. Although I might have been a little more productive if they weren't around, it would have been because I was so bored during lunch that I would have wound up working. Jeff Andre, Tim Boyd, Mike Bulpett, Jennifer Roy, and Robin Williams, stand up and take your bows. (That's enough, Mike.) And who could forget Pinky and The Brain, Wayne, Garth, Death, and all the rest who did all the actual work on this thesis while I was home for Christmas.

My parents deserve a big Thank You for helping to finance this 5-year hockey game I call a college education, and for generally being cool.

I'd also like to thank my partners at BBE, not for any help that they gave me on the thesis, but merely as a shameless plug for our business. We're available for work now.

Last and far from least (setting earlier comments of precedence aside), I would like to thank Shin-I for way too many reasons to list here. This work is really dedicated to her, but I can't resist a joke.

For Someone

Trajectory Estimation of an Optically Radiating Source

by

Craig Andera

Submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degrees of Bachelor of Science in Electrical Science and Engineering and Master of Engineering in Electrical Science and Engineering

Abstract

Consider a detector array and a radiating source in relative motion through the atmosphere. The objective of this thesis is to determine if sufficient information is contained in the scattered radiation field to provide accurate estimates of source *time to intercept* (TTI) and *angle of arrival* (AOA). The time to intercept is defined as the amount of time remaining until source and array are co-located. The definition of angle of arrival is self-evident. This thesis is a continuation of the research done by Mark R. Luettgen in his thesis by the same title for M.I.T. in 1990.

Models are developed using Loral's Threat Warning Engagement Simulator (TWES), based on atmospheric parameters from the LOWTRAN software package. Assumptions are made about the trajectories of the source and the detector platform. Based on these assumptions, models are developed for the atmosphere and the relative source trajectory, and an inner-product cost function is derived. A global search of the state space using an oct tree derivative provides the estimate of source trajectory.

It is shown that AOA and TTI can be estimated well for high SNR situations. Further, AOA can be estimated within a few degrees even in low SNR conditions. TTI estimation performance falls off rapidly as noise levels increase.

Table of Contents

List of Figures and Tables

1. Introduction

1.1 Threat Warning

The business of threat warning is the business of letting people know when they are in danger. Specifically, we will consider here the problem of alerting combat aircraft when they are under attack by a missile - commonly a Surface to Air Missile (SAM) - but the system developed here is equally applicable to air-to-air engagements.

The motivation for developing a threat warning system is clear: each aircraft costs several million dollars, to say nothing of the pilot inside it. Alerting the pilot to incoming missiles in a timely manner gives him or her several options to increase the aircraft's chances of survival.

1.2 What Are We Looking For and Why?

There are two parameters which are of particular interest concerning an incoming missile. These are Time To Intercept (TTI) and Angle Of Arrival (AOA).

1.2.1 TTI

Time To Intercept is defined as the time remaining before the aircraft and the inbound missile are co-located. An accurate knowledge of TTI is important for several reasons. First, it allows the pilot to prioritize. A missile that is still twenty or thirty seconds away is not yet a major concern and a missile due to impact in three seconds is of immediate importance.

Knowing TTI can also be used to increase the effectiveness of existing countermeasures. Aircraft commonly carry flares for decoying infrared homing (IRH) missiles. The release of a flare, when coupled with a rapid course change by the aircraft, can decoy an IRH missile. If this maneuver is timed properly - using the TTI provided by a threat warning system - the chances of success can be greatly increased.

Even without flares, TTI can be used to augment aircraft survivability. A rapid course change by a missile's target can cause the missile to miss. Obviously, the closer the missile is when the pilot executes this maneuver, the greater the aspect change. Think of a pair of football players, one trying to tackle the other. If the ball carrier zigs and zags while the defensive player is still far down field, little is accomplished. If, however, the carrier makes the same dodges at the last possible second, the defense will have a much harder time making contact.

1.2.2 AOA

For most of the same reasons, AOA is important to know. If the pilot can be told not only how long there is until impact, but from what direction the missile is approaching, he or she has that much more information on which to act, and thus a greater chance of evading the missile.

For example, consider the case in which the pilot is alerted to an approaching missile at ten seconds TTI. The pilot might want to take two completely different courses of action depending on whether the threat is directly ahead or immediately behind. If the threat warning system can provide the pilot with this information, aircraft survivability could be increased.

In addition to evasion enhancement, AOA information could theoretically be provided to some sort of directional countermeasures system. Unlike flares, which are simply released as a decoy, a directional countermeasures system would attempt to confuse or disable the missile by affecting it directly. One such system might consist of a powerful laser mounted on the aircraft. If accurate AOA could be provided, the laser could be pointed at the incoming missile, blinding or destroying the missile's detector array.

1.3 A Specific Sensor System

Clearly, in order to provide estimates of a missile's trajectory, the aircraft must possess some way of gathering information from the environment. For purposes of this analysis, we have chosen to model an existing system, already installed on many aircraft.

The system we will consider consists of six passive non-spatial sensors, with a combined field-of-view (FOV) of 4π steradians. The individual fields-of-view overlap each other. The sensors are arranged orthogonally, four along the intercardinal axes in the horizontal plane, and the remaining two along the vertical axes. See Figure 1.1 and Figure 1.2. Cross hatched areas show regions where fields-of-view overlap. The **x** axis points in the direction of aircraft travel, and the **z** axis points skyward.

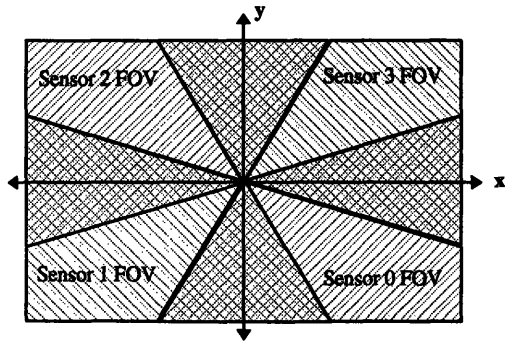

Figure 1.1 - Horizontal Cross Section Showing FOV

Passive systems such as this one have the advantage that they do not radiate detectable energy. While TTI and AOA are readily available with a high degree of accuracy to an active system such as radar, when employing radar the aircraft necessarily emits radiation that could make an opponent's job of detection much easier. If passive trajectory estimation can be accomplished, then the best of both worlds is achieved - knowledge and stealth.

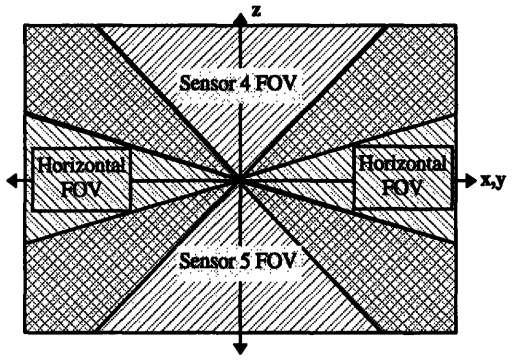

Figure 1.2 - Vertical Cross Section Showing FOV

Imaging sensors are available that give spatial information about the location of objects within each sensor's FOV. However, cost and size considerations make it worthwhile to investigate the use of non-imaging sensors to accomplish AOA and TTI estimation. A non-spatial sensor provides what is essentially a surface integral over scene brightness. This is similar to the problem of guessing from which direction a sound is emanating. Our ears, like our sensors but unlike our eyes, do not provide imaging information. We can estimate direction only by comparing the relative characteristics of the sound as detected by each of our ears. Surprisingly, even though our sensor system is what amounts to a six-pixel system, the algorithm is able to estimate AOA to a high degree of accuracy.

Several characteristics of the atmosphere in the sensor's operational waveband give rise to the phenomenon of scattering. Energy emitted by the source (in this case the missile) is redirected to arrive at the sensors from all directions. It is by modeling this effect and analyzing the proportion of energy received by each sensor that we are able to estimate AOA and TTI more accurately than would be possible using only the line-of-sight[1] (LOS) sensor.

1.4 Desired Solution Qualities

As with any engineering problem, developing a threat warning system that provides TTI and AOA involves balancing conflicting design parameters to achieve what is in some sense the "optimum" system. For our system, these parameters include the following:

- **Accuracy :** Clearly the system must be able to estimate TTI and AOA to some degree of accuracy. Our stated goal is to provide information to the pilot of an aircraft that will be useful in increasing the chances of surviving a missile engagement. Accuracy is therefore of primary concern.
- **Robustness :** The system must be able to handle a wide variety of threats, from slow to fast, near to far, bright to dim. An enemy will seldom provide his target with information about when, where, what and how he is going to shoot, so the system must be able to accurately predict TTI and AOA over a range of engagement parameters.
- **Real-time :** A system which predicts a five second TTI with ten second latency is of limited usefulness. The system must not only be able to provide accurate estimates of

[1] The LOS sensor is defined as the sensor for which the source is within the FOV. All other sensors, which receive only scattered energy are defined as non-line-of-sight (NLOS) sensors.

TTI and AOA over a wide range of engagement parameters, it must do so in a timely manner.

These parameters drove the development of the system. At times they directly contradicted each other. When this was the case, the tradeoffs and the particular choice will be discussed.

1.5 Past Work

Mark Luettgen's M.S. thesis, also entitled *Trajectory Estimation of an Optically Radiating Source,* presented a similar problem. Throughout this document, results from that study will be presented for comparison. Differences in the exact situation addressed make a direct, point-by-point comparison impossible. For example, Luettgen considers estimating TTI and AOA in a two-dimensional situation, where this thesis develops a system for three-dimensional trajectory estimation.

Luettgen's work was also done at Loral Infrared & Imaging Systems.

1.6 Summary/Overview

Section 2 of this thesis presents the development of the estimation algorithm. The atmospheric modeling is discussed, including the tools used to synthesize the model and the parameters over which the model was iterated. Kinematic modeling is similarly presented: how do the aircraft and the missile move through the atmosphere? Simplifying assumptions are discussed and explained.

With appropriate models determined, the Measurement Space Direction Matching method (MSDM) is introduced. MSDM is a method for collapsing the iterative search space by one dimension. In essence, the uncertainty in source intensity is eliminated as a search parameter.

Local minima in the state space abound, eliminating gradient descent methods as a viable option for searching the state space. Global behavior is still favorable, however, suggesting that an oct tree search could provide convergence. Such a search is presented and examined in section 2.3.

Section 3 of this thesis presents the main results and analysis. Test parameters are reviewed, and performance in terms TTI and AOA given. It is shown that AOA can be predicted to a fairly high degree of accuracy over a wide range of engagements. TTI is somewhat less well predicted, especially for low signal-to-noise ratios. The effects of test parameters - such as initial range and velocity - on algorithm performance are discussed in section 3.2.

Section 4 summarizes these results and suggests methods for further research by way of examining how more realistic modeling could be done.

2. The Algorithm

This section describes the development of the trajectory estimation algorithm. It begins with a description of the atmospheric and kinematic modeling. It then moves on to describe a cost function for comparing the choice of a particular state space vector to a measurement vector. Finally, it presents a method for searching through state space to find the "best guess", in this case the estimate of the state of the system.

2.1 Models

Any estimation system must model the phenomena it is trying to predict. Certain simplifying assumptions must be made, as available computational power is generally finite. Consequently, the models are not perfect representations of the system they attempt to emulate. It is the job of the engineer to create a model that is *sufficiently accurate*, i.e. not so accurate as to require an inappropriate amount of time to compute, but sufficiently accurate that results may be obtained within the specified error boundaries.

For our particular case, the modeling that must take place can be broken into two categories: atmospheric and kinematic.

2.1.1 Atmospheric

Our atmospheric model attempts to accurately represent the way energy is transmitted through the atmosphere, to eventually arrive at our sensor array. Changes in source intensity, atmospheric conditions, and sensor configuration all affect the signal that the sensor array receives.

2.1.1.1 Phenomenology

The phenomenology of the atmosphere is a two-edged sword. On the one hand, it makes possible high-accuracy AOA prediction. The same factors that give rise to this beneficial effect make estimating TTI much more difficult, however.

2.1.1.1.1 Vacuum

Consider for a moment how our system would function in a vacuum. Radiation emitted by the source would travel undisturbed in a straight line in all directions. Only the LOS sensor would have any flux across its surface. Ignoring for a moment the noise present in the system, we would have

$$\vec{z} = \begin{bmatrix} z_0 \\ z_1 \\ z_2 \\ z_3 \\ z_4 \\ z_5 \end{bmatrix} = \begin{bmatrix} \frac{k_1 I_0 \cos\phi}{r^2} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Equation 2.1

where **z** is the sensor output vector, $z_{0..5}$ are the outputs of the six sensors in counts, k_1 is the conversion efficiency, I_0 is the source intensity in Watts per steradian, r is the distance between the source and the detector, and ϕ is the angle between the detector normal and the vector which points from the sensor to the source[2].

Figure 2.1 - Solution is a cone in 3D

In this case, if the source were of constant intensity, non-maneuvering, and moving at a constant velocity,[3] estimating TTI would simply consist of finding when the output of sensor zero goes to infinity. Three samples of the output (ignoring the effects of noise for a moment) would be sufficient to make this prediction.

The angle of arrival, on the other hand, would be impossible to determine. Examining Equation 2.1, we see that unless intensity is known, ϕ cannot be uniquely obtained. Furthermore, even if the source intensity were known exactly, ϕ specifies only the angle between the source's approach vector and the detector normal. In three dimensional space, this is the surface of a cone (see Figure 2.1).

2.1.1.1.2 Absorption and Scattering

No aircraft can fly in a vacuum. This is both good news and bad news for our threat warning system. The introduction of atmosphere into our scenario makes possible AOA prediction, but we no longer have such an easy way of obtaining TTI. The reason for this is twofold.

---

2 We are assuming here that only sensor 0 has the source in its FOV. Further, the *cos* ϕ term applies only to a lenseless system. But regardless of the actual form of this term, there will still be some dependency on the angle between the sensor surface normal and the source position vector, and it will therefore not in general be possible to uniquely distinguish AOA. Consequently, we can use *cos* ϕ for the remainder of this section without loss of generality.

3 We will make exactly these assumptions in formulating the kinematic model later. As it turns out, they are not quite as unrealistic as it might seem.

As a photon moves through the atmosphere, it will occasionally "strike" something. The photon might then be re-directed or re-emitted in a pseudo-random[4] direction. This phenomenon is known as scattering. It might also be absorbed and not re-emitted. This phenomenon is known as extinction or absorption.

These twin (admittedly related) phenomena are what make AOA prediction possible and TTI prediction difficult. Occasionally, a photon will be redirected in such a way that it is detected by one of the NLOS sensors. Furthermore, the number of photons that reach each of the sensors in the array is determined by the geometry of the source with respect to the sensor array. For example, if a source at fixed range rotates counterclockwise from a position directly on the axis of detector zero (see Figure 1.1), the output of sensors one and two will increase, while those of sensors zero and three will decrease. The output of sensors four and five will remain the same, as the relative motion in this case is perpendicular to their normal, and the flux at their surface is thus unchanged.

As it turns out, the relative signal strengths of each sensor is a strong function of AOA, but a relatively weak function of TTI. That is, two engagements with relatively different AOA will produce distinct sensor measurement vectors, whereas two engagements with relatively different TTI can produce similar measurement vectors. Since our system must work backwards from measurement to source trajectory, we would expect good performance for AOA prediction, and somewhat weaker performance in terms of TTI prediction.

Given that relative sensor output is a weak function of TTI, it is therefore important to do a good job of modeling sensor output: small inaccuracies in the model can translate to relatively large errors in TTI prediction.

A photon en route to the sensor array might be absorbed, scattered once, scattered several times, scattered several times and then absorbed, etc. Furthermore, the degrees of scattering and absorption are sensitive to atmospheric conditions such as haze, visibility, constituents, etc., and these conditions are generally not available as inputs to the system.

Equation 2.1 is now much more complex. Our noiseless sensor outputs can be described as

$$\vec{z} = \begin{bmatrix} z_0 \\ z_1 \\ z_2 \\ z_3 \\ z_4 \\ z_5 \end{bmatrix} = \begin{bmatrix} \dfrac{k_1 I_0 \cos\phi \cdot e^{-k_2 r}}{r^2} + \theta_0(r,\phi,\ldots) \\ \theta_1(r,\phi,\ldots) \\ \theta_2(r,\phi,\ldots) \\ \theta_3(r,\phi,\ldots) \\ \theta_4(r,\phi,\ldots) \\ \theta_5(r,\phi,\ldots) \end{bmatrix} \qquad \textbf{Equation 2.2}$$

4 Pseudo-random because the distribution of re-direction is typically forward-peaked, although this is strongly influenced by several factors, including the size of the molecule or particle with which the photon is interacting.

where all variables are as above, k_2 is an atmospheric extinction coefficient, and the function $\theta(\bullet)$ is the scattering function which takes into account such factors as range, AOA, and atmospheric conditions.

Reducing the myriad possible paths a photon could take to a simple analytic function of source position and atmospherics is not possible. The dimensionality of the problem quickly makes it untractable for any reasonable environment. Consequently, the calculation must be performed off-line, and the results somehow modeled. For any reasonable model, Equation 2.2 is transcendental, and the curve fitting method which works so well in a vacuum is unavailable.

2.1.1.2 The Threat Warning Engagement Simulator

Loral Infrared and Imaging Systems (LIRIS) has developed precisely the tool to accomplish the task of atmospheric modeling. The Threat Warning Engagement Simulator (TWES) has been developed over several years, using an extensive database of field-recorded sensor data, to accurately model scattering and absorption in the atmosphere. Given source characteristics, source trajectory, atmospheric conditions, and sensor configuration, the TWES performs a numeric volume integration to produce noiseless sensor measurements.

These measurements were used to both develop the model and later to generate the inputs to the estimation system. Since the noise characteristic is well-defined (i.e. Poisson), these noiseless measurements could be retained and "noised up" repeatedly to determine the statistical properties of the estimation algorithm.

Since our goal is the *development* of a system for estimating TTI and AOA, and not the *implementation* of one, using the TWES both as basis for and input to the algorithm is a valid approach. We are mostly interested in proving that sufficient information is contained in the scatter field to yield useful estimates of TTI and AOA, and for that purpose, the TWES simulated outputs serve quite well. Further research aimed at producing an implementation of this system would use actual sensor readings taken in the field, and make appropriate adjustments based on the performance of the algorithm given those measurements.

2.1.1.3 Model Parameters

Although it would have been desirable to model the various atmospheric parameters over a wide range of values at high resolution, computational constraints did not allow this. Even picking only a few cases, the TWES took several weeks of computer time to complete.

Given the time constraints involved, a few key cases were chosen as representative of the whole data space.

2.1.1.3.1 Altitude

We assumed for the purposes of this thesis, that the sensor array would be situated at an altitude of 5000 meters. Although it would have been desirable to include data from lower

altitudes, it was decided that the ground model currently used in the TWES was not of sufficient complexity.

The ground, from an optical standpoint, is a complex and variable entity. It can be composed of regions of reflective surfaces interspersed with absorbent ones, as well as regions that are a bit of both. Desert and ocean are very different optically.

The current implementation of the TWES models the ground plane as a perfect absorber. That is, any photon that crosses the plane at altitude is extinguished. The code which would enable the TWES to more accurately simulate the variable effects of the ground plane was under development at the time of this writing.

Given this limitation in the TWES model, the 5000 meter data was chosen as the sole basis for the algorithm model. At 5000 meters, the ground has a much smaller effect than at (for example) 500 meters. Furthermore, altitude is one of the few parameters in the model that will be available to the system at run-time. Since there is no uncertainty in the altitude, the algorithm need not estimate its value, and the model need not provide sensor measurements for altitudes other than the one at which the aircraft is flying.

2.1.1.3.2 Absorption

The influence of extinction in the atmosphere is of primary importance. Relatively small differences in k_2, the coefficient of extinction from Equation 2.2, can have significant impacts on sensor output. Unfortunately, although the spatial distribution of absorbers in the atmosphere does not tend to change over the duration of a missile engagement, this distribution will generally not be known.

For purposes of this model, two different extinction coefficients were chosen as representative of the high and low ends of the spectrum.

2.1.1.3.3 Scattering

Several factors affect the scattering function θ including haze, visibility, and the size of particulate matter in the atmosphere. For purposes of this thesis, we chose four combinations of these factors to represent four varied sets of atmospheric conditions. The LOWTRAN[5] simulation (see reference [2] - the LOWTRAN manual) was then used to generate the exact coefficients required as inputs by the TWES.

2.1.1.3.4 Range

Preliminary results suggested that iterating over range at 75 meter resolution would be sufficient. That is, using linear interpolation between explicitly calculated points every 75 meters, provided a sufficiently smooth curve. Sufficiently smooth in this case was taken to mean that the impact of sensor noise was much greater than the error produced by the interpolation.

---

5 LOWTRAN is a Department of Defense software package that predicts transmittance and radiance parameters for a given set of atmospheric conditions.

2.1.1.3.5 Azimuth

Azimuth refers to the longitudinal angle the source approach vector makes with the aircraft reference axis. Since the atmosphere is presumed to be horizontally homogeneous, it was only necessary to iterate over 45 degrees. Any other configurations could exploit the symmetry of the problem to obtain the sensor output vector.

Preliminary results revealed 5 degrees to be sufficient resolution for accurate reproduction.

2.1.1.3.6 Elevation

Elevation is the latitudinal angle the source approach vector makes with the aircraft reference axis. Unfortunately, the same symmetry that applies in the azimuthal direction does not apply to elevation. Even at 5000 meters altitude, the ground plane has a significant impact on the sensor output. In addition, the atmospheric model used by the TWES incorporates the vertical inhomogenities found in the real atmosphere. Thus, model data was generated over 180 degrees of elevation. As with azimuth, preliminary results showed 5 degrees resolution to be sufficient.

2.1.1.3.7 Source

For our purposes, the radiative source was modeled as an isotropic 3000K 1 m^2 blackbody. For reasons discussed in section 2.2.3.2 below, the actual temperature of the source was unimportant. Assuming the source to be isotropic eliminated source geometry considerations in the estimation problem. In reality, detectable radiation in a missile engagement is not emitted in an isotropic manner. Source emission geometry may or may not significantly affect sensor measurements, however the additional dimensionality this would introduce would carry the estimation problem beyond the scope of this thesis.

2.1.1.4 Luettgen's Approach

Luettgen approached his atmospheric modeling in a different manner. Although his thesis was also developed at LIRIS, he did not have available the TWES. Consequently, his work was based on two atmospheric models he developed himself: one single scatter model, and one phenomenological multiple scattering approximation.

In addition to differences in atmospheric modeling, Luettgen was working in only two dimensions. No energy was emitted or scattered out of the horizontal plane. Consequently, his AOA estimator used only data from the four sensors whose normals lay in the horizontal plane (corresponding to sensors zero through three in this thesis). Furthermore, his TTI estimation considered only trajectories on the boresight of the sensor. That is azimuth and elevation were both known exactly.

In contrast, this thesis will examine the case wherein initial source location is constrained only by min and max boundaries on range. Within this region, the source may be located at any range, azimuth, and elevation (r, θ, ϕ) triplet.

2.1.1.5 Synthesis of Final Model

Having run the TWES on all combinations of the parameters discussed above, the results were collected and examined. It was discovered that the haze and visibility have little effect on sensor output.

Because of the particular cost function chosen, the important figure of interest is the inner product of the normalized measurement vectors. That is, two models could be considered equivalent to the algorithm, if, for the same source position, they produced two sensor vectors whose inner product was simply the product of their magnitudes. In measurement space, this corresponds to the two measurement vectors pointing in the same direction.

The reasons for this criteria are made clear in section 2.2.3.2 below. For now, let us examine the effects of scattering and absorption on the modeled sensor measurement vectors.

Examine Table 2.1 and Table 2.2, which show the effect each atmospheric parameter has on the measurement vector direction. The tables were generated by comparing on a point-by point basis the output vectors of the simulation for each possible pair of atmospheric

Table 2.1 - Maximum Angular Deviations in Atmospheric Models (degrees)

Extinction Coefficient		Low	Low	Low	Low	High	High	High
	Scatter Model	A	B	C	D	A	B	C
Low	B	0.55320	-	-	-	-	-	-
Low	C	0.05234	0.59479	-	-	-	-	-
Low	D	0.71244	0.16551	0.75407	-	-	-	-
High	A	5.27778	5.57366	5.26313	5.69258	-	-	-
High	B	5.25572	5.51825	5.24120	5.63579	0.08622	-	-
High	C	5.28079	5.58062	5.26619	5.69959	0.03956	0.96913	-
High	D	5.25531	5.51455	5.24082	5.63207	0.09065	0.03956	0.10087

parameter triplets. For example, the sensor measurements for a source at 5000 meters, 0 degrees azimuth and 0 degrees elevation in an atmosphere with high extinction, and scatter model A was compared to a source at the same position in an atmosphere with high extinction and scatter model B. The angle between the two measurement vectors was recorded.

Table 2.2 - Average Angular Deviations in Atmospheric Models (degrees)

Extinction Coeff.	Scatter Model	Low A	Low B	Low C	Low D	High A	High B	High C
Low	B	0.10163	-	-	-	-	-	-
Low	C	0.00791	0.10935	-	-	-	-	-
Low	D	0.11630	0.01479	0.12405	-	-	-	-
High	A	2.66861	2.73645	2.66347	2.74737	-	-	-
High	B	2.66370	2.79125	2.65858	2.74212	0.00885	-	-
High	C	2.66934	2.73783	2.66420	2.74815	0.00511	0.00968	-
High	D	2.66397	2.79153	2.65885	2.74240	0.00854	0.00489	0.00945

Table 2.1 shows the maximum angular deviation encountered for each pair of atmospheres, while Table 2.2 displays the average angular deviation. Clearly, the absorbers in the atmosphere have by far the largest effect. The difference between low and high extinction coefficients in the atmosphere on average produces a 2.7 degree deviation in measurement vector direction, with deviations as high as 5.7 degrees possible.

Scattering coefficients were of somewhat lesser effect. Differences tended to deviate measurement vectors by no more than 0.75 degrees.

At what point, then, can we consider these deviations to be insignificant? In this case, the noise sets the mark. The addition of noise to each individual sensor output will clearly have some effect on the direction of the measurement vector.

Using the outputs of the TWES runs, it was determined that under typical SNR conditions, an average deviation in sensor measurement vectors of approximately 1 degree was caused by the noise. Even for high SNR readings, the noise could deviate the measurement vector by as much as 0.25 degrees.

As the deviations caused by the scattering coefficients are - for most source locations - below the sensor measurement noise floor, the four low extinction atmospheres were averaged together to yield a single model for a low extinction atmosphere independent of scattering coefficients. A single high extinction atmosphere was similarly obtained. Table 2.3 shows the angular deviations between the eight original models and the two models obtained from them and used by the algorithm. Clearly, deviations are below that produced by the noise, even for high SNR situations.

Table 2.3 - Deviations Between Final Model and Full Model (degrees)

Extinction Coeff.	Scatter Model	Max Deviation	Avg Deviation
Low	A	0.03634207	0.02417174
Low	B	0.06106026	0.02466245
Low	C	0.04717093	0.02440490
Low	D	0.04290679	0.02457286
High	A	0.04522724	0.02423578
High	B	0.10887216	0.02567613
High	C	0.05822540	0.02458014
High	D	0.07306450	0.02481426

Thus, the final model consisted of 71,040 data points (2 atmospheres, 96 ranges, 10 azimuths, and 37 elevations), which would require about 277 kilobytes of RAM to store in an implemented system. While this is not prohibitive, an implementation of this system would probably use piecewise surface-fitting techniques to reduce storage requirements. For purposes of this thesis, however, atmospheric model data was obtained by "memorizing" these datapoints and using linear interpolation for source positions not falling precisely on one of the generated points.

2.1.2 Kinematic

The atmospheric model developed allows us to obtain, for an isotropic source of any intensity and at a specific location, the noiseless sensor measurement vector. What remains, then, is to model the evolution of the source position.

2.1.2.1 Missile Flight

A missile flight proceeds in three phases. In the first, the missile is launched and quickly accelerates to its terminal velocity. In the second phase, the missile closes with its target. In the final phase, the missile attempts to make the final maneuvers which will result in the destruction of its target.

We are primarily interested in the second phase of the flight. During its first phase, a missile will be difficult to track, as it is heating up (and thus has a non-constant intensity) and maneuvering as it comes up to speed and attempts to set up for the second phase of its flight. During the third phase, if our algorithm has not already done its job, it is too late.

2.1.2.2 Model Parameters

We assume trajectories are sufficiently specified by the state vector $[r_0, \theta, \phi, v_0, I_0]^T$ and the time, t, since the initiation of the algorithm.

2.1.2.2.1 Velocity (v_0)

Velocity is here defined to be the rate at which the distance between the aircraft and the missile is decreasing, and not the relative ground speed of the missile. Thus the same

missile might be considered to have a much smaller velocity for a tail shot than for a head-on engagement.

The velocity will be assumed constant for the duration of interest. A missile in the second portion of its flight does not experience large accelerations, so this is a valid assumption.

2.1.2.2.2 Approach Vector (θ,ϕ)

During the second phase of flight, most missiles employ what is known as proportional navigation, as it requires the least energy to intercept a target. A beneficial side-effect of this is that the missile will remain at a relatively stable AOA with respect to the aircraft frame of reference. This fact combines with the ranges involved to make our assumption of a constant azimuth and elevation a fairly good one.

Consequently, the direction of approach should remain fairly constant. In any case, the distances involved will generally rule out the possibility of major changes in missile azimuth and elevation, relative to the aircraft reference axis.

2.1.2.2.3 Initial Range (r_0)

Since we assume the approach velocity of the missile to be constant, we need know only the range at any given time to accurately estimate TTI. For purposes of this algorithm, we will estimate the range at the time which coincides with the initiation of the algorithm.

2.1.2.2.4 Intensity (I_0)

Source intensity is assumed constant over the duration of the engagement. Since we are assuming that the missile has accelerated to its final speed during the first stage of its flight, and since we are assuming all sensor measurements are taken on a missile in the second phase of its flight, we can safely assume that no large temporal variations in temperature (and thus in intensity) will occur.

In addition, significant changes in source intensity can generally be detected, if not measured, and this information can be used to invalidate or re-initiate the estimation algorithm.

2.1.2.2.5 Sensor platform in level flight

Our final assumption about the kinematics of the engagements involve the sensor platform itself. We assume that over the course of the engagement, the aircraft is in level flight at 5000 meters. As on-board inertial systems are capable of providing the estimation algorithm with accurate information concerning the trajectory of the aircraft, the algorithm could be altered to incorporate this information. Thus, little generality is lost by considering only level flight paths.

2.1.2.3 Luettgen's Approach

As explained in section 2.1.1.4, Luettgen examined only constant velocity, non-maneuvering engagements in the horizontal plane. Additionally, TTI was only estimated for sources approaching along the boresight of a sensor.

2.2 Cost Function

The atmospheric model combined with the kinematic model allow us to obtain a series of sensor measurement vectors for any trajectory state vector $[r_0, \theta, \phi, v_0, I_0]^T$. We are interested in reversing this process. That is, given a series of sensor measurement vectors, we would like to know the trajectory state vector.

If the atmospheric model were linear, this would be a simple matter of solving simultaneous equations. In fact, however, the function which relates sensor outputs and source position is highly non-linear. Thus, we must arrive at our approximation to the trajectory state vector by making a guess, seeing how far off we are, improving the guess, and looking again. We repeat this process until we have determined, under some criteria, that we are "close enough". Our cost function - that is, our method of determining whether or not we are close enough - uses an inner product of measurement vectors. This has the advantages of being computationally simple and eliminating the need to iterate the search over source intensities.

2.2.1 The Extended Kalman Filter

In Luettgen's work, he employs an Extended Kalman Filter (EKF) to estimate TTI (AOA is computed directly). The body of work on Kalman filtering is large. See [5], [6], and [7] for a more comprehensive explanation than can be given here.

2.2.2 A Bayesian Approach

The actual output of each sensor is a Poisson random variable whose mean is proportional to the modeled photon flux at the surface. Thus, we have

$$\begin{bmatrix} h_0 \\ h_1 \\ \vdots \\ h_5 \end{bmatrix} = f\left(\begin{bmatrix} r_0 \\ \theta \\ \phi \\ v_0 \\ I_0 \end{bmatrix}, t\right)$$

$$\begin{bmatrix} p_z(z_0|h_0) \\ p_z(z_1|h_1) \\ \vdots \\ p_z(z_5|h_5) \end{bmatrix} = \begin{bmatrix} \frac{h_0^{z_0} e^{-h_0}}{z_0!} \\ \frac{h_1^{z_1} e^{-h_1}}{z_1!} \\ \vdots \\ \frac{h_5^{z_5} e^{-h_5}}{z_5!} \end{bmatrix} \qquad \textbf{Equation 2.3}$$

where $[h_0, h_1, \ldots, h_5]^T$ are the modeled, noiseless sensor outputs, f is our atmospheric modeling function which gives the noiseless sensor measurements given source position and intensity, and $p_z(z_n|h_n)$ is the conditional probability of a sensor reading exactly z_n at a given time, given that the noiseless reading would be h_n.

The expression we are interested in is the conditional probability of a trajectory state given our set of measurements. Expressed mathematically, this is

$$p_h(h_n|z_n) = \frac{p_z(z_n|h_n)p_h(h_n)}{p_z(z_n)} \qquad \textbf{Equation 2.4}$$

where $p_h(h_0|z_n)$ is the state probability function that is our goal, $p_z(z_n|h_n)$ is as above, $p_h(h_n)$ is the unconditional probability of the noiseless sensor output, and $p_z(z_n)$ is the unconditional probability of the noisy sensor reading.

We assume $p_h(h_n)$ to be uniform, since without more *a priori* knowledge of the engagement, one measurement is as intuitively likely as any other. The factor $p_z(z_n)$ is simply a normalization factor to satisfy $\int p_h(h_n|z_n)dh_n = 1$.

Taking into account all six sensor readings at once, and noting that h_n can be written as $f_n(\mathbf{x}, t)$, where $\mathbf{x}$ is our trajectory state vector, we have

$$p_x(\vec{x}|Z_k) = c_1 \prod_{s=0}^{5}\prod_{j=0}^{k} p_z(z_{s,j}|f_s(\vec{x}, jT)) \qquad \textbf{Equation 2.5}$$

Our notation in Equation 2.5 introduces the idea of using more than one integration interval's worth of sensor measurements. In this equation, Z_k is a shorthand for the set of all measurements $[z_0, z_1, \ldots, z_5, z_6, \ldots, z_{6k-1}]^T$ for all six sensors from time 0 through time kT, where T is the integration period of the sensor. That is, z_6 is the measurement from the same sensor that produced z_0, only T seconds later. We also denote this as $z_{s,j}$ where s denotes the sensor number and j denotes the integration interval. The constant c_1 encompasses the two normalization terms from Equation 2.4.

Expanding gives us

$$p_x(\vec{x}|Z_k) = c_1 \prod_{s=0}^{5}\prod_{j=0}^{k} \frac{[f_s(\vec{x}, jT)]^{z_{s,j}} \exp[-f_s(\vec{x}, jT)]}{z_{s,j}!} \qquad \textbf{Equation 2.6}$$

or

$$p_x(\vec{x}|Z_k) = c_2 \exp\left[\sum_{s=0}^{5}\sum_{j=0}^{k}\left(z_{s,j} \ln[f_s(\vec{x}, jT)] - f_s(\vec{x}, jT)\right)\right] \qquad \textbf{Equation 2.7}$$

where c_2 now takes into account the $z!$ terms as well, although if we are interested in the maximum likelihood, we can ignore this normalization constant altogether. Similarly, exponentiation is a monotonically increasing function, so we can concern ourselves only with finding the maxima of the double summation.

At first glance, this would appear to be exactly what we need: an exact expression giving the probability that a particular state vector $\mathbf{x}$ describes the true trajectory given a series of sensor measurements. However, this Bayesian approach has several drawbacks. The first of these is that it is mildly computationally expensive. Calculation of a natural log for

every detector, for each integration period, over every point searched in the state space can quickly add up.

Additionally, the natural log of the modeling function is then multiplied by the actual sensor reading, and the modeling function subtracted from this product. Due to the finite precision of digital computers, this could lead to significant error. Since in floating point notation, the computer only keeps track of the most significant digits, the less significant digits that would be needed to accurately compute the difference are not available.

It is possible that a scheme could be developed to overcome these limitations. However, an alternate method was developed that is not only computationally simpler, but has the added advantage that the state space need only be searched in four of the state variables, rather than the five required by this Bayesian scheme.

Figure 2.2 - Poisson vs. Normal Distribution for Various Means

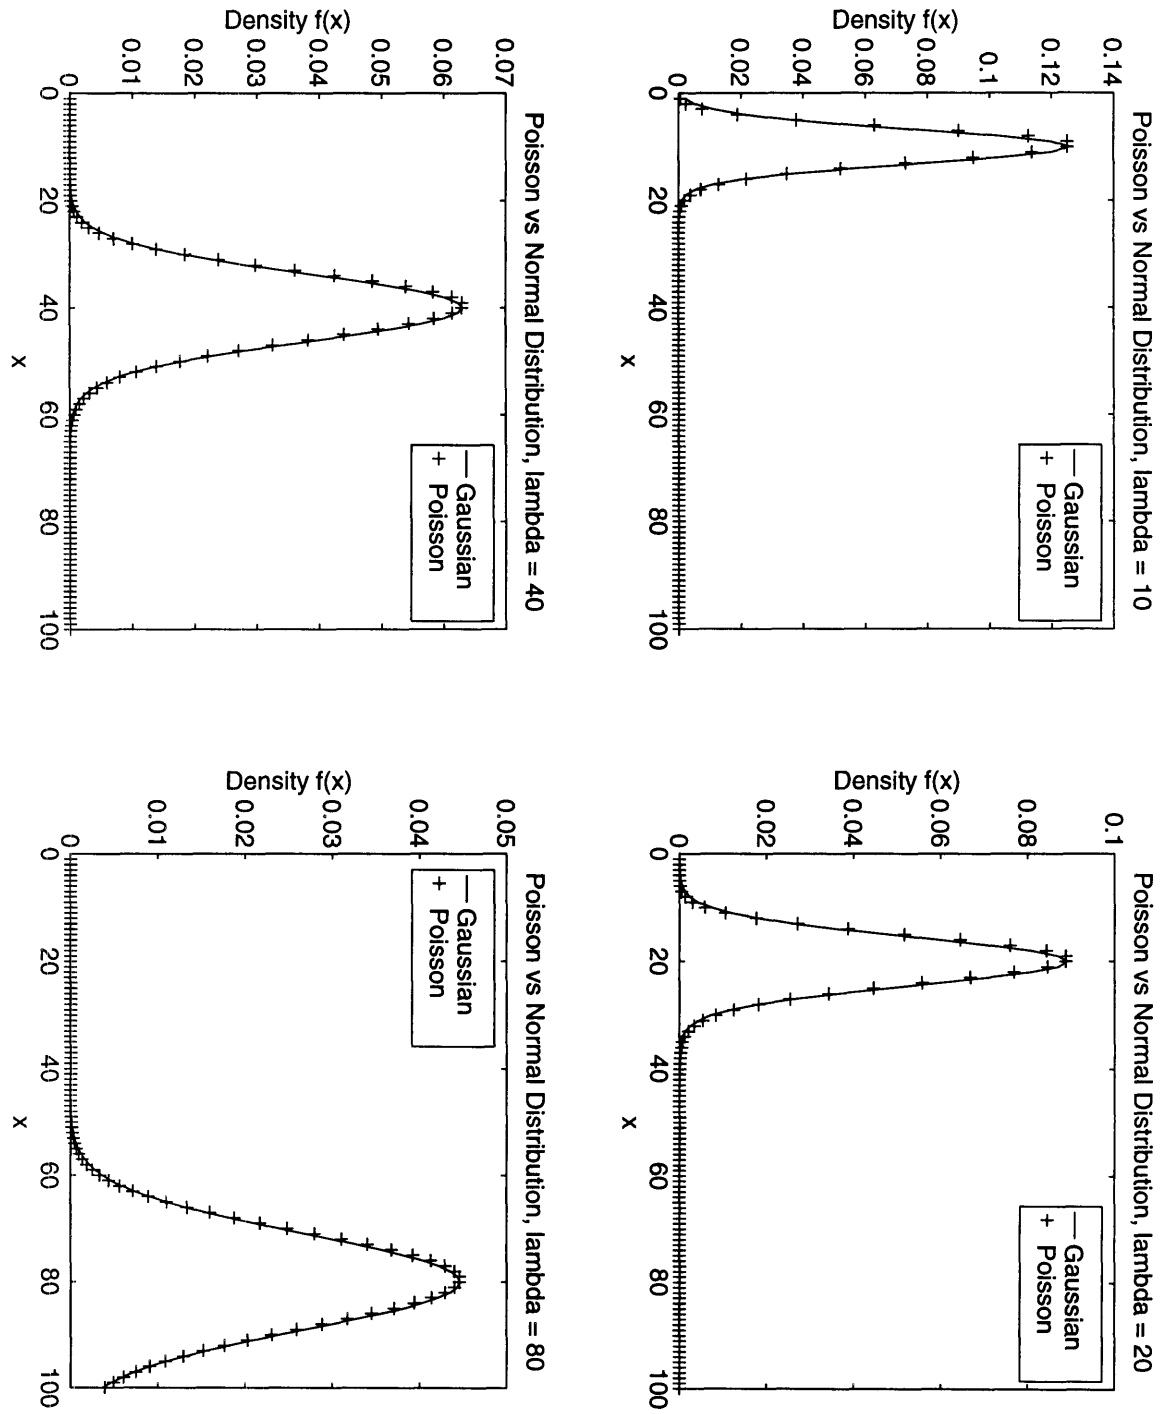

2.2.3 A Better Approach - Measurement Space Direction Matching

2.2.3.1 Basis - Minimum Squared Error

The central limit theorem can be used to show that as the mean of a Poisson random variable increases, it is approximated more and more closely by a normal distribution. In fact, as can be seen from Figure 2.2, even for relatively low means, the two distributions are remarkably similar.

If we consider the noise to be additive and of normal distribution, this suggests a minimum squared-error criteria would be appropriate as a cost function [9]. That is, we would find the state vector **x** for which the Equation 2.8 was minimal.

$$\sum_{s=0}^{5}\sum_{j=0}^{k}\left[\left(f_s(\vec{x}, jT) - z_{s,j}\right)^2\right] \qquad \textbf{Equation 2.8}$$

Geometrically, this is equivalent to minimizing a distance, albeit in at least six dimensions. The "distance" in this case is between our actual, noisy sensor measurements, and the noiseless measurements predicted by our modeling function *f*. For simplicity's sake, consider the case in which we have only two sensors, and we are basing our estimate on only one integration period. Figure 2.3 shows the measurement space, with our noisy, measured data and the vector our model produces for a given trajectory state. The distance between them is shown as a dotted line. It is this distance that our search method would seek to minimize.

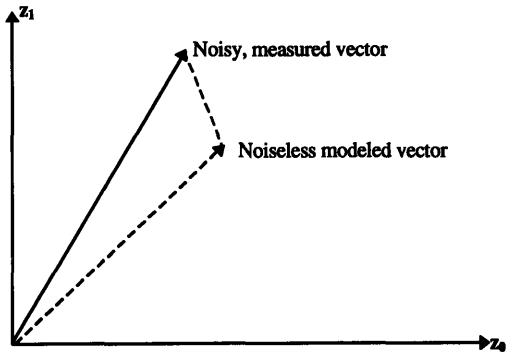

Figure 2.3 - Actual and Modeled Sensor Vectors in Measurement Space

2.2.3.2 Measurement Space Direction Matching

The key to Measurement Space Direction Matching (MSDM) is the assumption of constant source intensity over the course of the engagement. Realizing that the atmospheric modeling function

$$\begin{bmatrix} h_0 \\ h_1 \\ \vdots \\ h_{\frac{6t}{T}-1} \end{bmatrix} = f\left(\begin{bmatrix} r_0 \\ \theta \\ \phi \\ v_0 \\ I_0 \end{bmatrix}, t\right)$$

Equation 2.9

scales linearly in I_0, we see that it could also be written as

$$\begin{bmatrix} h_0 \\ h_1 \\ \vdots \\ h_{\frac{6t}{T}-1} \end{bmatrix} = I_0 g\left(\begin{bmatrix} r_0 \\ \theta \\ \phi \\ v_0 \end{bmatrix}, t\right)$$

Equation 2.10

The subscripts on h indicate both a sensor number and a time index. Variables h_0 through h_5 are the noiseless outputs for sensors 0 to 5 respectively. Variable h_6 corresponds to the output of sensor 0 at time T, the integration period, variable h_7 corresponds to the output of sensor 1 at time T, and so on.

The linear effect of intensity arises from the manner in which the sensors produce measurements. Essentially, they are counting photons that impact their surfaces. When source intensity is doubled, twice as many photons are emitted by the source, and subsequently twice as many arrive at the detector array.

This phenomenon also accounts for why, when developing the atmospheric model, the choice of a $1m^2$ 3000K blackbody was unimportant. A blackbody of that size and at that temperature emits approximately 40 Watts/sr. This is far brighter than any missile that might be fired at an aircraft, but since intensity is linearly related to sensor output, we could do the calculations to scale the sensor measurements into the appropriate range at a later time.

This effect has important consequences for our cost function. Since intensity is an unknown, and can vary continuously over a wide range, an algorithm which searched the state space over all six variables (r_0, θ, ϕ, v_0, I_0, and the atmospheric model f) would have to expend a large amount of time iterating over the possible values of I_0. However, we can exploit the linear relation between intensity and sensor output as follows:

We are looking for the point in state space which corresponds (via our atmospheric modeling function f) to the point in measurement space closest to our noisy measurements. The intensity I_0 scales all sensor measurements by the same factor. Geometrically, this is equivalent to changing the length of the measurement vector. Since intensity is unknown, we can perform our search by specifying only five of the state variables - namely r_0, θ, ϕ, v_0, and f - and then determining the intensity for which the distance to the noisy measurement vector is minimized.

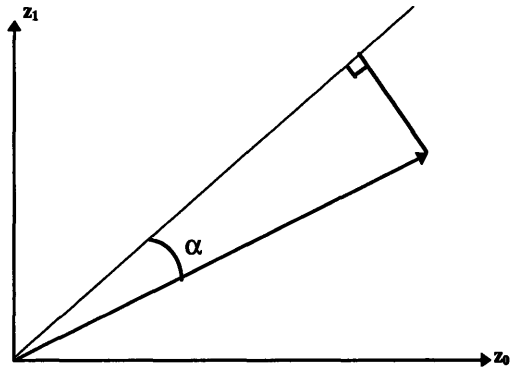

Figure 2.4 - Modeled and Actual Sensor Vectors in Two Dimensional Measurement Space

Consider Figure 2.4, where we once again simplify to a two-dimensional measurement vector. The line extending from the origin represents the possible model outputs for a particular choice of range, azimuth, elevation, and velocity, with intensity completely unspecified. It is, in essence, a family of measurement vectors. Since the shortest distance from a point to a line is perpendicular to the line at the point of intercept, this intercept point represents a minimization over I_0 of Equation 2.10.

Clearly, since the length of the measurement space vector is essentially completely unconstrained, the state vector for which Equation 2.10 is minimized will correspond to the modeled measurement vector which makes the minimum angle with the measured, noisy vector. Again consider Figure 2.4. For any other state vector with a smaller value of α, an intensity could be chosen such that the distance between the modeled measurement vector and the noisy measurement vector is less.

Recall that

$$h \cdot z = |h||z| \cos\omega$$

where h and z are two arbitrary vectors, $|h|$ is the magnitude or length of h, $|z|$ is the magnitude or length of z, and ω the angle between the two. This operation, known as the *dot product* or *inner product*, is computationally inexpensive, being merely the sum of the elementwise products of the two vectors. We are interested in minimizing ω, which can be expressed as

$$\omega = \cos^{-1}\left(\frac{h \cdot z}{|h||z|}\right) \qquad \textbf{Equation 2.11}$$

In our case, h is the model output vector for a given state vector $\mathbf{x}$, (i.e. $f(\mathbf{x},t)$) at some arbitrary intensity. The variable z is our noisy measurement vector. Since the inverse cosine is a monotonic function, and since all sensor outputs are positive, we can simplify our search to maximizing the operand of $\cos^{-1}$ in Equation 2.11 above.

If our system does not need to output an estimate of intensity - and it does not if we wish only to estimate TTI and AOA - we can further simplify the calculations required by prenormalizing our sensor model functions g. Normalizing the noisy sensor measurement vector is relatively inexpensive, as it must only be performed once for each search of the state space.

Mathematically, we have produced a cost function Q, that yields a measure of the distance between our truncated state vector $\mathbf{y}$ (r_0, θ, ϕ, v_0, g) and noisy measurement vector Z_k.

$$Q(\vec{y}, Z_k, k) = \sum_{s=0}^{5} \sum_{j=0}^{k} z'_{s,j} \, g_s'(\vec{y}, jT) \qquad \textbf{Equation 2.12}$$

The prime mark on the variable z and the function g indicate that they have been normalized. An unnormalized version, which would allow the system to output an estimate of intensity, would be

$$Q(\vec{y}, Z_k, k) = \sum_{s=0}^{5} \sum_{j=0}^{k} z'_{s,j} \frac{g_s(\vec{y}, jT)}{|g(\vec{y}, jT)|} \qquad \textbf{Equation 2.13}$$

This is slightly more computationally expensive, requiring a magnitude computation for each integration period processed, but it allows the system to estimate intensity. This would be accomplished by determining the relative lengths of the model vector which maximizes the cost function (representing a 40 W/sr source) and the noisy sensor measurement vector. That is,

$$I_{est} = 40.0 \frac{|Z_k|}{|\psi(\vec{y}_{opt}, k)|} \text{ [W / sr]} \qquad \textbf{Equation 2.14}$$

where I_{est} is the estimated intensity, $|Z_k|$ is the magnitude of the $6k$-dimensional series of measurement vectors, $\mathbf{y}_{opt}$ is the optimum truncated state vector, and Ψ is the $6k$-dimensional series of noiseless modeled sensor outputs.

2.2.3.3 Advantages

This version of the cost function has several advantages. Unlike the Bayesian cost function, $Q(\bullet)$ has a range of [0,1]. Computational precision is not an issue. This algorithm is also less computationally expensive: only $6k$ multiplications and $6k$ additions are required for each search point, compared to $6k$ multiplications, $6k$ additions, and $6k$ natural logarithms for the Bayesian formulation.

The real advantage, however, is realized in the collapse of the search space. Since intensity can vary over a large range (approximately two orders of magnitude for typical missiles [3]), a relatively large number of intensity values would need to be evaluated for a purely iterative search of trajectory state space. Using MSDM, we escape the limitations of an iterated search across intensity while requiring less computational power than would a minimum squared error approach. In other words, rather than being able to only predict one of a finite number of intensities (e.g. n times some base number if our iteration uses evenly spaced search points), while with MSDM, a "best fit" is selected from a continuous range.

2.3 Search Method

Two of the three requirements for obtaining a working estimation system are now ours: a model of the environment, and a cost function for telling us how close we are to a

solution. The third - a method for searching our state space for an optimum solution - is presented below.

2.3.1 Luettgen's Approach

As described above, Luettgen uses an EKF approach to TTI estimation. The reader is again referred to references [5], [6], and [7] for more information on how the EKF makes state estimates.

Luettgen does not, however, use an EKF for AOA. For his AOA estimation, he turns back to his atmospheric model. He notes that the ratio of the outputs of the two side-lobe NLOS sensors[6] is relatively insensitive to the distance between the detector array and the source, and completely insensitive to source intensity. He performs a polynomial fit to this ratio as a function solely of AOA (under the assumption that range and intensity have no effect). Inverting this function, he has a polynomial which yields AOA given only two measurement functions.

2.3.2 Failure of linear approach

Ideally, we would like to use some sort of gradient ascent to locate our optimum state estimate. That is, we would make some initial estimate of the state, then find the direction in state space which decreases the angle in measurement space between the modeled measurement vector corresponding to our current state estimate, and the actual, noisy measurement vector. We could then repeat this process until some criterion had been met.

Restated, this is:

1. Make some initial guess $y_0 = [r_0, \theta, \phi, v_0]$, recalling that I_0 has been factored out.
2. Calculate the corresponding modeled measurement vector b_0 using each g, our model functions from Equation 2.10 above. That is, $b_0 = g(y_0)$ for each atmosphere we consider possible.
3. Compute the angle our modeled measurement vector b_0 makes with our measurement vector Z_k, using our cost function $Q(\bullet)$ from Equation 2.13 above.
4. Determine the gradient of $Q(\bullet)$ with respect to the state vector s_0, and update our guess by moving in that direction in state space. That is, $s_{n+1} = s_n + c_0 \nabla Q$, where c_0 is some convergence constant.
5. Repeat until Q reaches some threshold.

---

6 That is, the two sensors whose boresights are 90° offset from the LOS sensor - the "left" and "right" sensors.

When employed, however, this method seldom converges. The reason for this is evident if one examines the cost function.

$$Q(\vec{y}, Z_k, k) = \sum_{s=0}^{5} \sum_{j=0}^{k} z'_{s,j} \frac{f_s(\vec{y}, jT)}{|f(\vec{y}, jT)|}$$

We are seeking the global maxima (w.r.t. the vector **y**) of this function given some Z_k. Each individual term of the summation could also have one or more maxima, which would correspond to the state vector most likely given only that particular sensor measurement $z_{s,j}$. Furthermore, the noise in the system nearly guarantees that these maxima will not occur in the same place. Clearly, when we add many of these individual functions together, the resulting function $Q(y, Z_k, k)$ will have abundant local minima. Therefore, we must look for a search method other than gradient descent.

2.3.3 Global Characteristics Suggest Oct Tree

While over a small region the local maxima are abundant, the cost function is still globally smooth. That is, a state vector that is far from the true state will nearly always evaluate to a lower value of Q than one that is close to the true solution.

This combination of properties - global smoothness and local bumpiness - suggest that we need a search method whose resolution will increase as we draw closer to our final estimate. One such method that works quite well for our case is known as an Oct Tree search.

Roughly speaking, the Oct Tree is the multidimensional counterpart of the binary search. We begin by searching the state space at a very rough resolution - on the order of 100 points evenly spaced throughout the entire region. We then narrow our search to examine those points in the neighborhood of the point which yielded the maximum value in the original search. This process is repeated, doing consecutively finer grained searches, for a set number of iterations, at which time the resolution has reached some specified value.

This method has several desirable characteristics, outlined below.

2.3.3.1 Adjustable Convergence Time

Since the number of iterations can be arbitrarily chosen, a good balance between convergence time and desired resolution can be found. If accuracy in the final answer is of more importance than is keeping the computational requirements low, many iterations can be done, with the final search being performed at a very fine resolution. Conversely, if the need for an answer in real time drives the system design, and only limited computational power is available, accuracy can be sacrificed to speed by simply stopping the search sooner.

2.3.3.2 Separable TTI/AOA

We have noted previously that intensity is not easily distinguished from range. A near, dim object produces a measurement vector very similar to that of a far, bright source. Thus, it would be unwise to narrow our search over range and velocity in the first stages of the search. If, for example, we had a source 3 km from the detector, our first, low-resolution search might find the cost function maxima to be at some azimuth and elevation close, but not equal to the true values, and at 5 km range. In this case, the value at 3 km might be only slightly lower, and then only because of the initial inaccuracy in estimating AOA.

For this reason, a slightly modified version of the Oct Tree search is employed. For all iterations but the last, range and velocity are searched only at low resolution. Because AOA estimation accuracy is little affected by range, this allows the system to make a fairly good guess of azimuth and elevation before attempting to find the correct range and velocity. Having an accurate estimate of AOA before attempting to discern TTI keeps us from narrowing our search to the wrong range/velocity region.

In effect, this approach separates the TTI and AOA estimation problems. Although the same method is used for both, AOA is obtained first, and used to estimate TTI. In a real time system, AOA might be provided first to a directional countermeasures system. When TTI becomes available some time later, additional measures can be taken based on the new information it provides.

3. Analysis

The models, cost function, and search method presented in the previous section comprise the elements of the algorithm we will use to warn our theoretical pilot of the theoretical incoming missile. Having developed the system, we now examine how well it works.

3.1 Test Parameters

It would be desirable to test the system for every possible combination of engagement. However, computational constraints did not allow this. Even running on six separate DEC 4000-90 workstations, and even over the limited number of engagements chosen to comprise the test matrix, results took nearly three weeks to generate.

The rationale for the selection of test parameters is given below. For each engagement, 200 separate runs were made, in order to fully characterize the effect of noise on the system.

3.1.1 Duration/Integration Period

A standard integration period of 100 ms was chosen to coincide with the integration period used by existing algorithms developed at LIRIS. Rather arbitrarily, one second was chosen as the duration over which the algorithm would operate. Thus, with a six-sensor system, the measurement space was sixty-dimensional.

3.1.2 Range

Initial ranges of 5000m, 3000m, and 1500m were chosen as being representative ranges of modern-day heat-seeking surface-to-air missiles (SAMs). For shots much closer than 1.5 km, engagement times are so short that trajectory estimation would be of little value. By the time an estimate was made, the pilot would have little time to react. For ranges beyond 5000m the probabilities of successful engagement by currently employed SAMs are low.

3.1.3 Azimuth/Elevation

For an aircraft in level flight, clearly few attacks will come from straight above. Because of the proportional navigation that most missiles employ (see section 2.1.2.2.2), very few will come from straight below. Therefore, elevations of +15°, 0°, -15°, -30°, and -60° were chosen, where negative values indicate elevation below the horizontal plane.

As discussed above, the atmosphere is symmetric in azimuth. This fact, when combined with sensor geometry meant that values of 0°, 22.5°, and 45°, measured from the reference axis of the aircraft, were sufficient to span the azimuthal gamut.

3.1.4 Intensity

Typical missile intensities range from 5 mW/sr to 500 mW/sr [3]. Therefore, 5 mW/sr, 50 mW/sr, and 500 mW/sr were chosen to span this range.

3.1.5 Velocity

Typically, missile velocities range from 300 to 900 m/s. Aircraft typically travel at 100 to 500 kts, which translates to approximately 50 and 250 m/s. For a head-on shot, this gives a maximum closure velocity of 1150 m/s. The minimum closure velocity occurs when a slow missile chases a fast aircraft from directly behind. For our numbers, this means a 50 m/s closure rate.

At 50 m/s the missile is unlikely to be able to catch the aircraft. We concern ourselves here with only those cases where the aircraft is in danger of being struck, so we choose 250, 750, and 1150 m/s as our engagement velocities.

3.1.6 Atmospherics

Limited by computational constraints, we chose to search only two atmospheres. The first atmospheric model used a low extinction coefficient and scattering model A. The other atmosphere was generated using a high extinction coefficient and scattering model D. This combination represents a wide spread in atmospheric conditions. We chose to vary the atmospheric parameters in this manner in order to demonstrate that algorithm performance is relatively insensitive to variations in the scattering and absorption coefficients.

3.2 Results

As stated above, for each possible combination of test parameters, the algorithm was run on 200 distinct sensor measurement vectors. That is, although our ideal, noiseless sensor measurement vector was the same, 200 different samples of the system noise function were chosen, in order to give a good idea of the sensitivity of the algorithm to the random parameter of sensor measurement noise.

3.2.1 TTI

Performance of the algorithm for TTI estimation was in general fair, although strongly influenced by several factors explained below. First we examine the performance of the EKF estimator employed by Luettgen.

3.2.1.1 EKF Estimator Performance

As explained previously, the EKF estimation system employed by Luettgen attempted only to estimate TTI in the situation where the aircraft's AOA was known to be coincidental with the sensor boresight. Nevertheless, they provide an interesting basis for comparison.

The reader is referred to [13] for Luettgen's full results, but the following summary will suffice for our purposes here: At 1500 to 2000 meters with a slow closing velocity, RMS errors in TTI were no lower than 5 seconds.

3.2.1.2 MSDM Estimator Performance

In comparison, the MSDM estimator performed slightly better. Referring to Table 3.1, we see that for a 1500 meter shot, at slow closing velocities, even the lowest intensity from our test set results in a RMS error of approximately 4.3 seconds.

Table 3.1 - TTI RMS Error for Low Extinction Atmosphere

		Velocity		
Intensity	***Range***	250 m/s	750 m/s	1150 m/s
5 mW/sr	1500 m	4.318	0.430	0.084
50 mW/sr	1500 m	3.733	0.132	0.039
500 mW/sr	1500 m	3.380	0.100	0.029
5 mW/sr	3000 m	15.064	2.883	1.712
50 mW/sr	3000 m	4.257	2.185	1.455
500 mW/sr	3000 m	1.331	1.371	1.298
5 mW/sr	5000 m	16.847	10.309	4.517
50 mW/sr	5000 m	17.093	1.738	0.836
500 mW/sr	5000 m	4.415	0.520	0.287

Table 3.1 and Table 3.2 summarize the results of the MSDM algorithm's estimates. Examining the data, we see that range, intensity and velocity all have a strong effect on the accuracy of the estimation. This is hardly surprising in the case of range and intensity, as

Table 3.2 - TTI RMS Error for High Extinction Atmosphere

		Velocity		
Intensity	***Range***	250 m/s	750 m/s	1150 m/s
5 mW/sr	1500 m	1.506	0.210	0.059
50 mW/sr	1500 m	0.844	0.113	0.036
500 mW/sr	1500 m	0.553	0.111	0.037
5 mW/sr	3000 m	10.840	2.683	0.906
50 mW/sr	3000 m	5.046	1.045	0.445
500 mW/sr	3000 m	2.244	0.590	0.151
5 mW/sr	5000 m	16.432	7.834	5.539
50 mW/sr	5000 m	13.829	4.337	2.361
500 mW/sr	5000 m	11.432	2.355	1.096

the signal-to-noise ratio is also a direct function of these parameters.

That velocity should have such a strong impact on estimation accuracy makes sense on consideration. In a noisy environment, a 250 m/s closing velocity would imply that the sensor input vector was not changing much as a function of time. On the other hand, the measurement at $t = 1$ second for a 1150 m/s engagement would be markedly different from the measurement at $t = 0$. In other words, the rate of increase of the sensor inputs cue the algorithm more accurately to the missile velocity, resulting in a lower-error estimate of TTI.

Interestingly, we see that there is no clear correlation between atmospheric conditions and estimator performance. In some cases, the estimator performs better for low-extinction atmospheres, while in other cases, better estimation is obtained for trajectories in high-extinction engagements.

One possible explanation is that the model may be a poor fit in regions where one or the other atmosphere is estimated with high error. For example, the averaged model might be a worse fit to the actual atmospheric transfer function for low extinction than for high extinction, in the mid-range, low intensity, slow shot regime. Further work might include an examination of which atmospheric conditions are well-fitted by the averaged atmospheric model.

Table 3.1 and Table 3.2 make the limitations of MSDM TTI estimation readily apparent. At 3000 meters initial range, a 250 m/s closing velocity implies a 12-second flight time. The RMS errors for low-intensity engagements of this type are in the range of 10-15 seconds, clearly more than is acceptable. At the same time, the 1-2 second RMS errors corresponding to high-intensity engagements indicate that usable estimates of TTI may be estimated by the algorithm when intensities are high enough.

It is clear, then, that the performance of the TTI estimation varies over a wide range. The data presented in Table 3.1 and Table 3.2 outlines what one could expect for various combinations of engagement parameters.

Since the presence of measurement noise ensures that our TTI estimate will not be perfect all of the time, the question of confidence measurement arises. That is, given some output of the algorithm, can we say anything about how sure we are that the estimate is correct? Bayesian analysis similar to the type presented in section 2.2.2 could yield such a measure. Although it is beyond the scope of this thesis to perform such an analysis, nevertheless a few distribution graphs are presented here.

Figure 3.1 to Figure 3.4 (shown on pages 40 to 43) require some explanation. Each depicts the histogram of predicted TTI values, representing an approximation to the probabilistic distribution of TTI estimates given a set of engagement parameters. All engagement parameters save intensity are listed at the bottom of the page, and apply to all histograms on that page. The four histograms represent engagements with source

intensities of 5 mW/sr, 50 mW/sr, and 500 mW/sr, with the fourth histogram showing the aggregate distribution over all intensities.

The legend at the top of each histogram indicates the following:

- **RMS Err:** The root-mean-squared error of TTI estimates for the set of engagement parameters corresponding to that histogram.
- **Inten:** The source intensity for that histogram
- **Inf TTI:** If a velocity of zero is predicted by MSDM, the corresponding TTI estimation is necessarily infinity. These estimates are discarded, and this statistic indicates how many times this occurred for the given plot.
- **True TTI:** The zero-error TTI estimate, indicated by a dashed vertical line on the plot.
- **Vel std:** The standard deviation of velocity estimates.
- **Vel mean:** The mean of the velocity estimates.
- **Rng std:** The standard deviation of the range estimates.
- **Rng mean:** The mean of the range estimates.

The column of statistics on the right is intended to give some sense of the primary source of error in the estimate of TTI. Since TTI is a ratio of two estimates (namely those of velocity and range), the distribution of TTI estimates does not provide information about the accuracy of range and velocity estimates. It would bear examination if the range estimate were highly variable while the velocity error was low (or vice versa), but we see from the graphs that this is not the case. In general both range and velocity are estimated with proportional amounts of error.

Four cases are presented. The first two, with r_0 = 3000m, v_0 = 750 m/s, represent a baseline case. Figure 3.1 shows this configuration for a high extinction atmosphere, where Figure 3.2 shows this configuration for a low extinction atmosphere. It is interesting to note that both atmospheres display some bimodality in the distribution of TTI estimates. This points to the presence of local minima in the sensor modeling function.

The next two graphs are intended to demonstrate a "best case" and "worst case" scenario. In Figure 3.3, a close-range, high-velocity engagement is show, whereas in Figure 3.4, the situation is just the opposite - low closure rate and extreme range. In both cases we have displayed results for both atmospheric models aggregately. In each case the distribution of TTI estimates is as we would expect: wide for the long-range case, narrow for the close-in shot.

Figure 3.1 - Mid-range TTI estimation distribution for high extinction

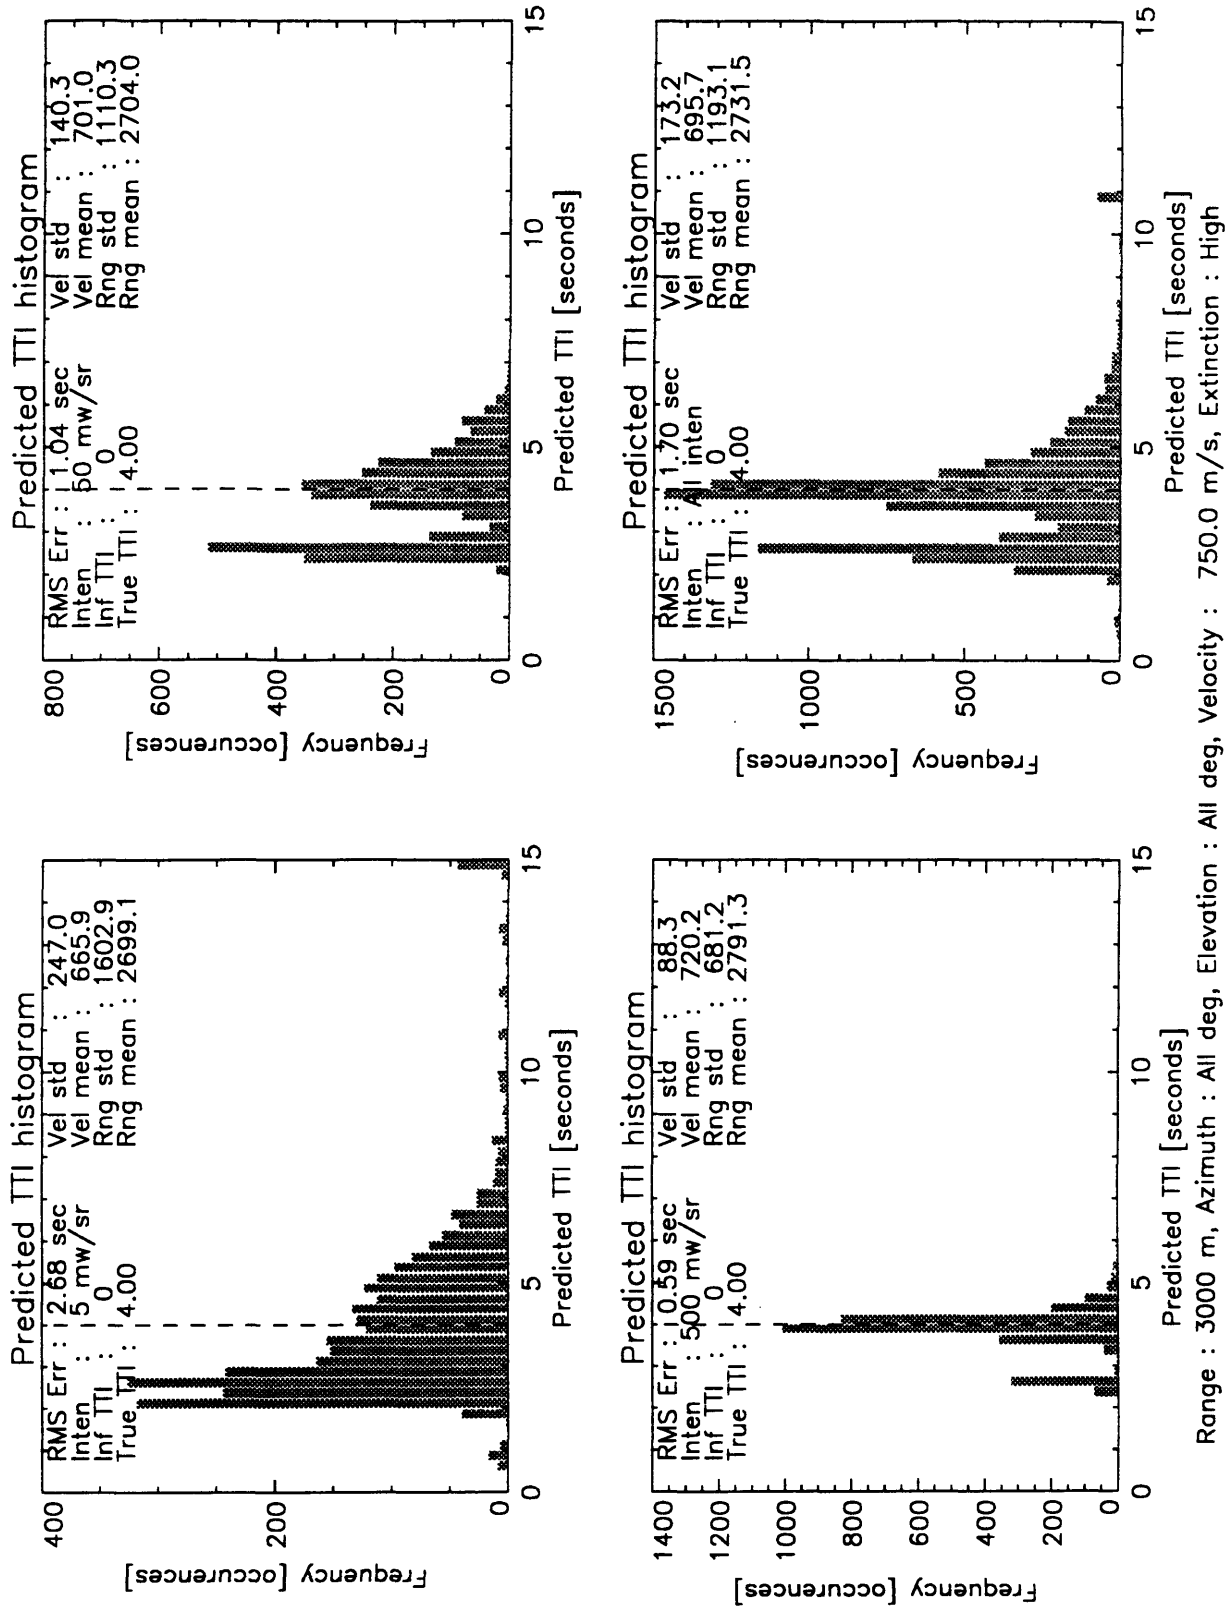

Figure 3.2 - Mid-range TTI estimation distribution for low extinction

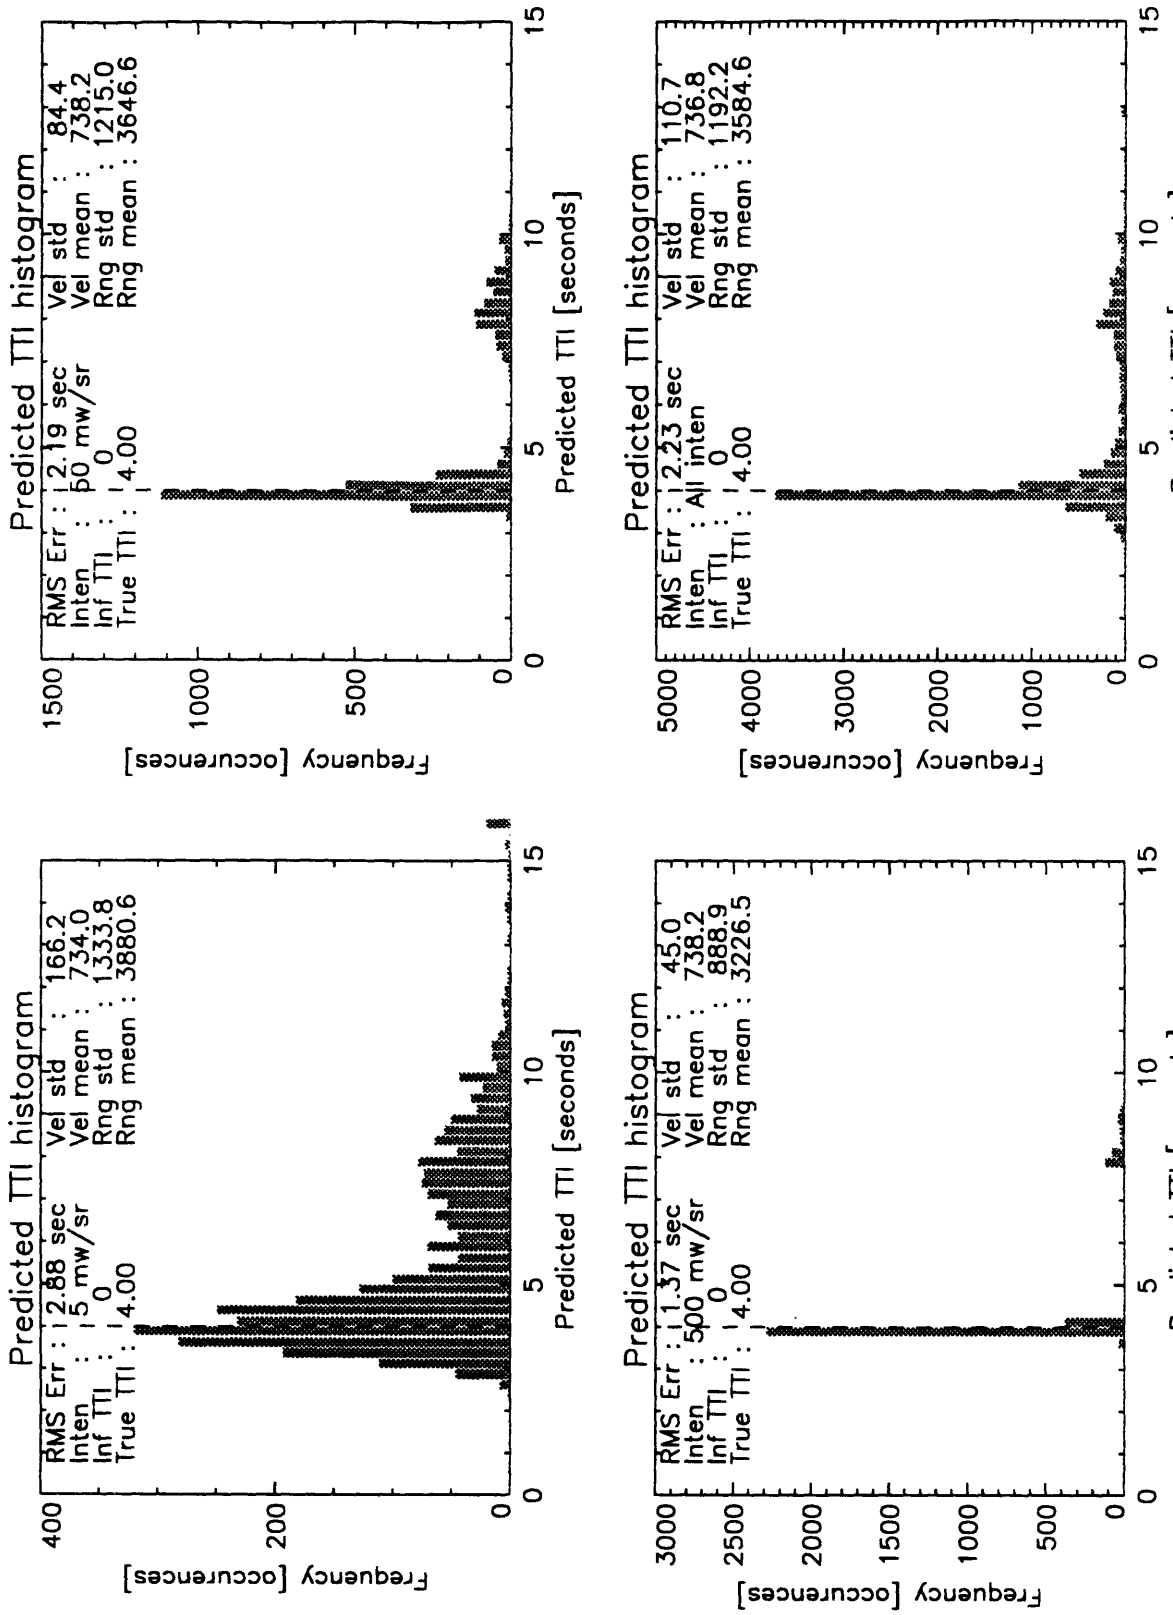

Figure 3.3 - High SNR TTI estimation distribution

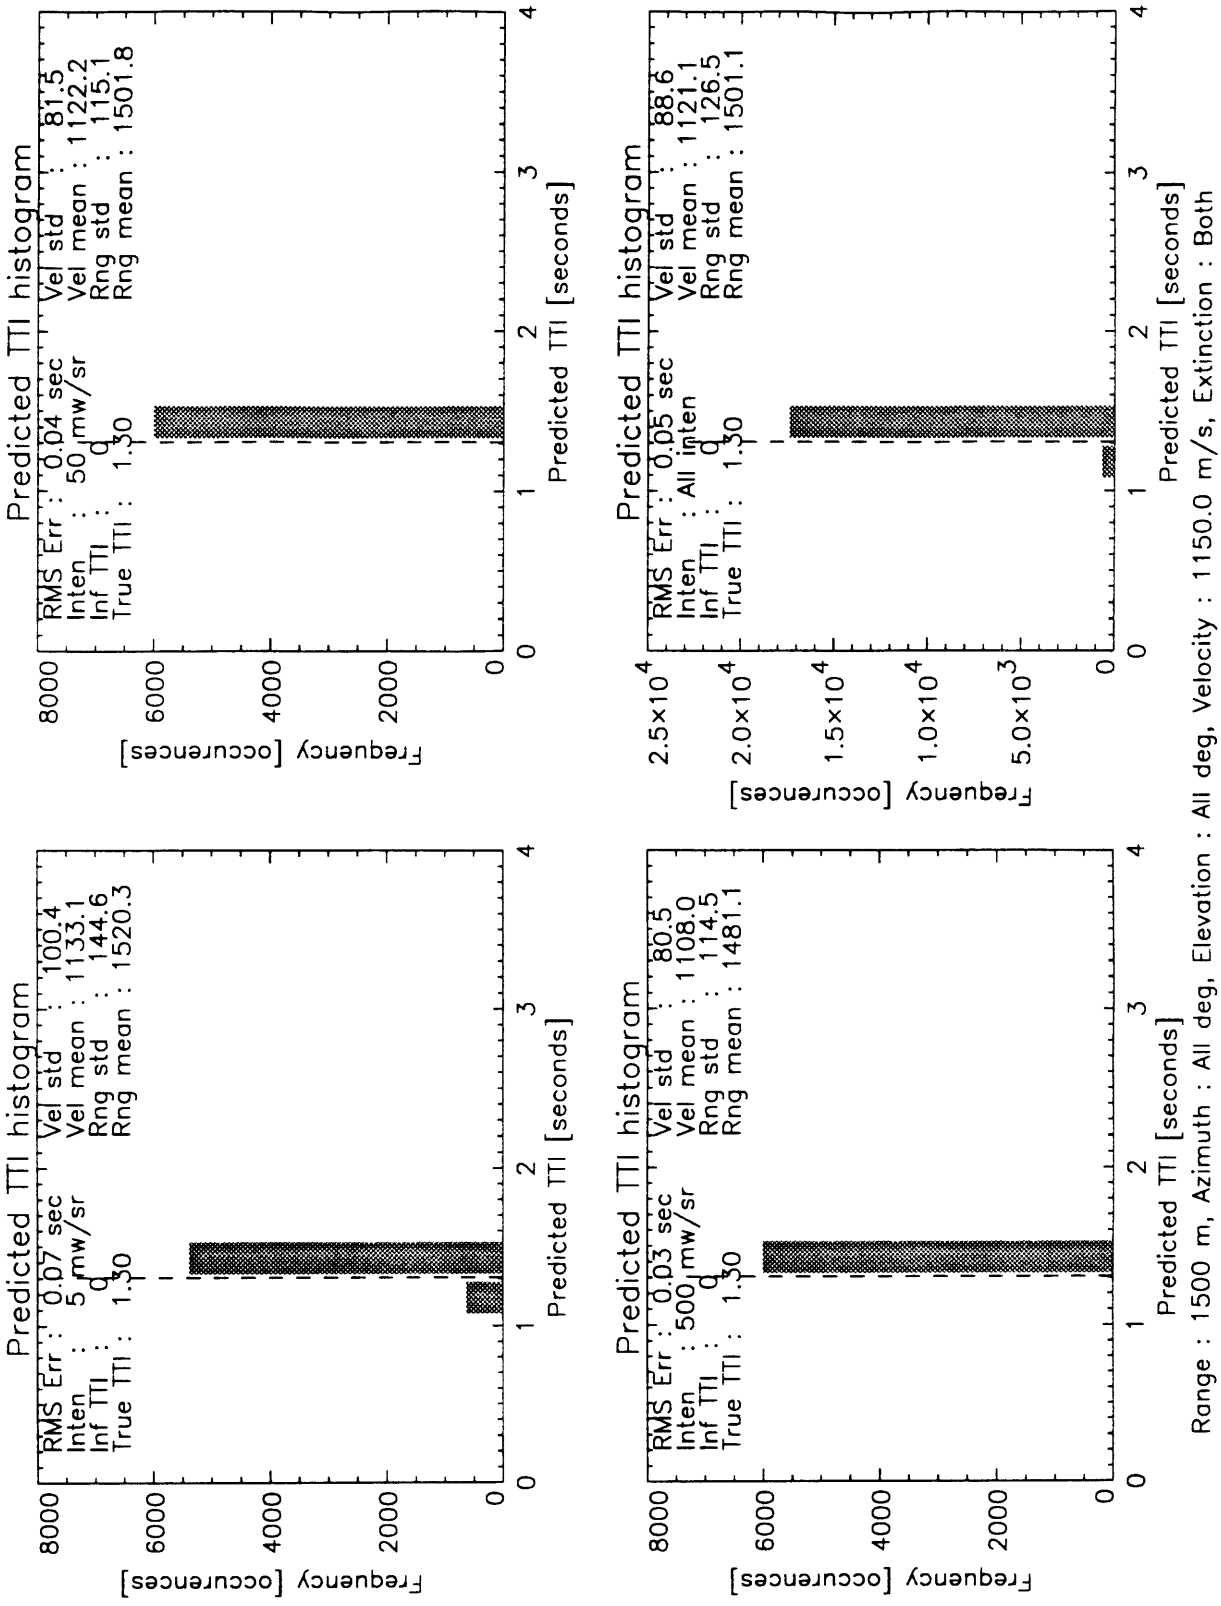

Figure 3.4 - Low SNR TTI estimation distribution

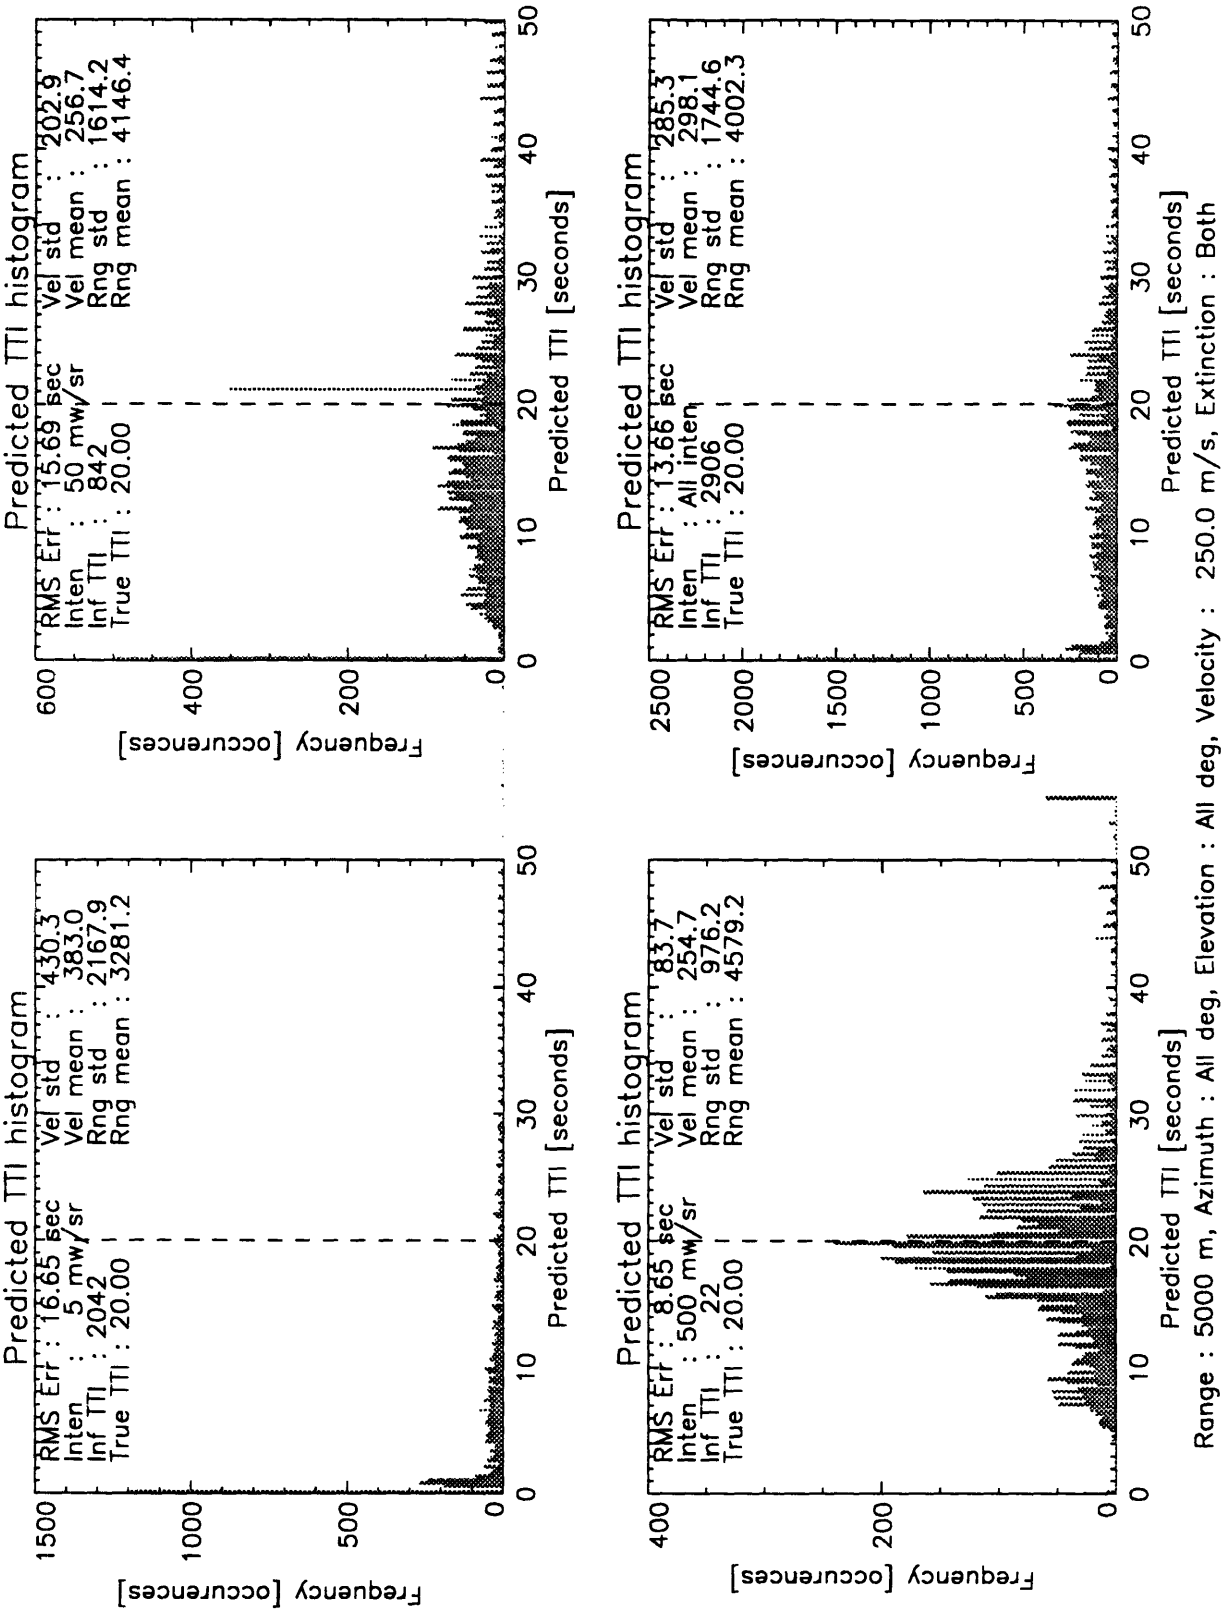

3.2.2 AOA

Angle of arrival estimation proved to be fairly accurate, and in general was influenced by the same factors as TTI estimation. Again, we saw that MSDM performed somewhat better than the EKF method employed by Luettgen.

3.2.2.1 EKF Estimator Performance

The reader is referred again to [13] for a complete treatment of Luettgen's AOA estimation results, but in general RMS errors ranged between 3 and 15 degrees. These estimates were taken for a source at 2000m, closing at 100 m/s, with a source intensity of between 1 and 3 mW/sr. This is a one-dimensional problem: elevation is considered fixed and known at 0 degrees.

The MSDM algorithm, when performed on an engagement with 3000m range, at 250 m/s closure rate, with a source intensity of 5 mW/sr, produced an estimate with an RMS error of between 3 and 8 degrees, depending on the atmospheric model used. This is, however, the two-dimensional error. That is, error was determined by the formula

$$\theta_e = \cos^{-1}(\vec{e} \cdot \vec{p})$$

where θ_e is the error angle, **e** represents the three-dimensional directional estimate vector, normalized to unit length, and **p** represents the three-dimensional true direction vector, also normalized to unit length.

Thus, MSDM was able to estimate both azimuth and elevation with a lower error than the EKF operating under similar conditions.

Table 3.3 - AOA RMS Error for Low Extinction Atmosphere

		Velocity		
Intensity	***Range***	250 m/s	750 m/s	1150 m/s
5 mW/sr	1500 m	2.188	2.331	2.687
50 mW/sr	1500 m	1.564	2.113	2.614
500 mW/sr	1500 m	1.419	2.371	2.637
5 mW/sr	3000 m	3.665	3.716	3.723
50 mW/sr	3000 m	1.496	1.834	2.011
500 mW/sr	3000 m	0.701	1.043	1.453
5 mW/sr	5000 m	8.467	8.034	7.871
50 mW/sr	5000 m	2.733	2.650	2.604
500 mW/sr	5000 m	1.110	1.063	1.032

Table 3.4 - AOA RMS Error for High Extinction Atmosphere

		Velocity		
Intensity	***Range***	250 m/s	750 m/s	1150 m/s
5 mW/sr	1500 m	2.377	1.977	2.461
50 mW/sr	1500 m	1.669	1.231	2.178
500 mW/sr	1500 m	1.435	1.010	2.203
5 mW/sr	3000 m	7.276	6.205	5.436
50 mW/sr	3000 m	2.484	2.482	2.479
500 mW/sr	3000 m	1.258	1.531	2.054
5 mW/sr	5000 m	25.431	21.661	18.343
50 mW/sr	5000 m	9.295	8.242	7.474
500 mW/sr	5000 m	2.972	2.803	2.702

3.2.2.2 MSDM Estimator Performance

Table 3.3 and Table 3.4 summarize the performance of the MSDM AOA estimation under low extinction and high extinction respectively. As with TTI estimation, we observe a direct correlation between decreasing range/increasing intensity and increasing estimation accuracy, and presumably for the same reason. That is, the lower signal levels produced by a dimmer or more distant source lead to a decreased SNR. The increased distortion in the sensor measurement vectors results in a less accurate estimation of AOA.

There are two interesting facts which become evident when examining this data. First, results from runs in the high extinction atmosphere had mostly higher estimation errors than the corresponding scenarios in a low extinction error. One possible explanation for this is that the higher extinction coefficient, combined with a higher degree of scattering, tends to make the reading for each sensor more similar. That is, in an environment with very high levels of scattering, we would expect each sensor, even NLOS sensors, to read approximately the same level of incident radiation. Think of the problem of trying to tell what direction the sun is in on a very foggy day, and one can appreciate why AOA accuracy might fall off.

The second thing that we note is that AOA accuracy does not in general increase for increasing velocity. In some cases, low-speed engagements result in the highest accuracy for a given range and intensity, while in others medium-speed produces the best results. High-speed engagements might also claim the lowest RMS error, although this is more rare.

This is a quite unexpected result. Further clouding the issue is the fact that at some ranges, the velocity of minimum error changes with source intensity. Unfortunately, the three velocities and three intensities do not provide a sufficiently detailed picture to draw

conclusions about what is driving the error rate. Future work could be focused on further analyzing this phenomenon.

We can, however, draw some broad conclusions about the performance of AOA estimation. It appears that, for low enough SNR, AOA performs quite poorly. As SNR increases, AOA performance improves, until reaching a minimum of around 2 degrees, although or some cases, the RMS error dips as far as 0.7 degrees.

Included here as well, in Figure 3.5 to Figure 3.8, are error distribution graphs for AOA estimation. Note that these are not AOA *value* histograms, as was the case for the TTI distributions in Figure 3.1, Figure 3.2, Figure 3.3, and Figure 3.4, but rather *error* distributions. Labeling in all other respects is the same as for TTI.

We have again presented, as a sort of average-case scenario, two sets of histograms with 3000m initial range and 750 m/s velocity, across all AOA combinations. The first, Figure 3.5, is for this trajectory in a high extinction atmosphere, while Figure 3.6 shows the same trajectory in a low extinction environment. Figure 3.7, representing the high SNR case, shows a trajectory with 1500m initial range and 1150 m/s closing velocity, for all AOA and both extinction cases. Figure 3.8 gives the distribution for a low SNR case, with 5000m initial range and a closure rate of 250 m/s.

Figure 3.5 - Mid-range AOA estimation distribution for high extinction

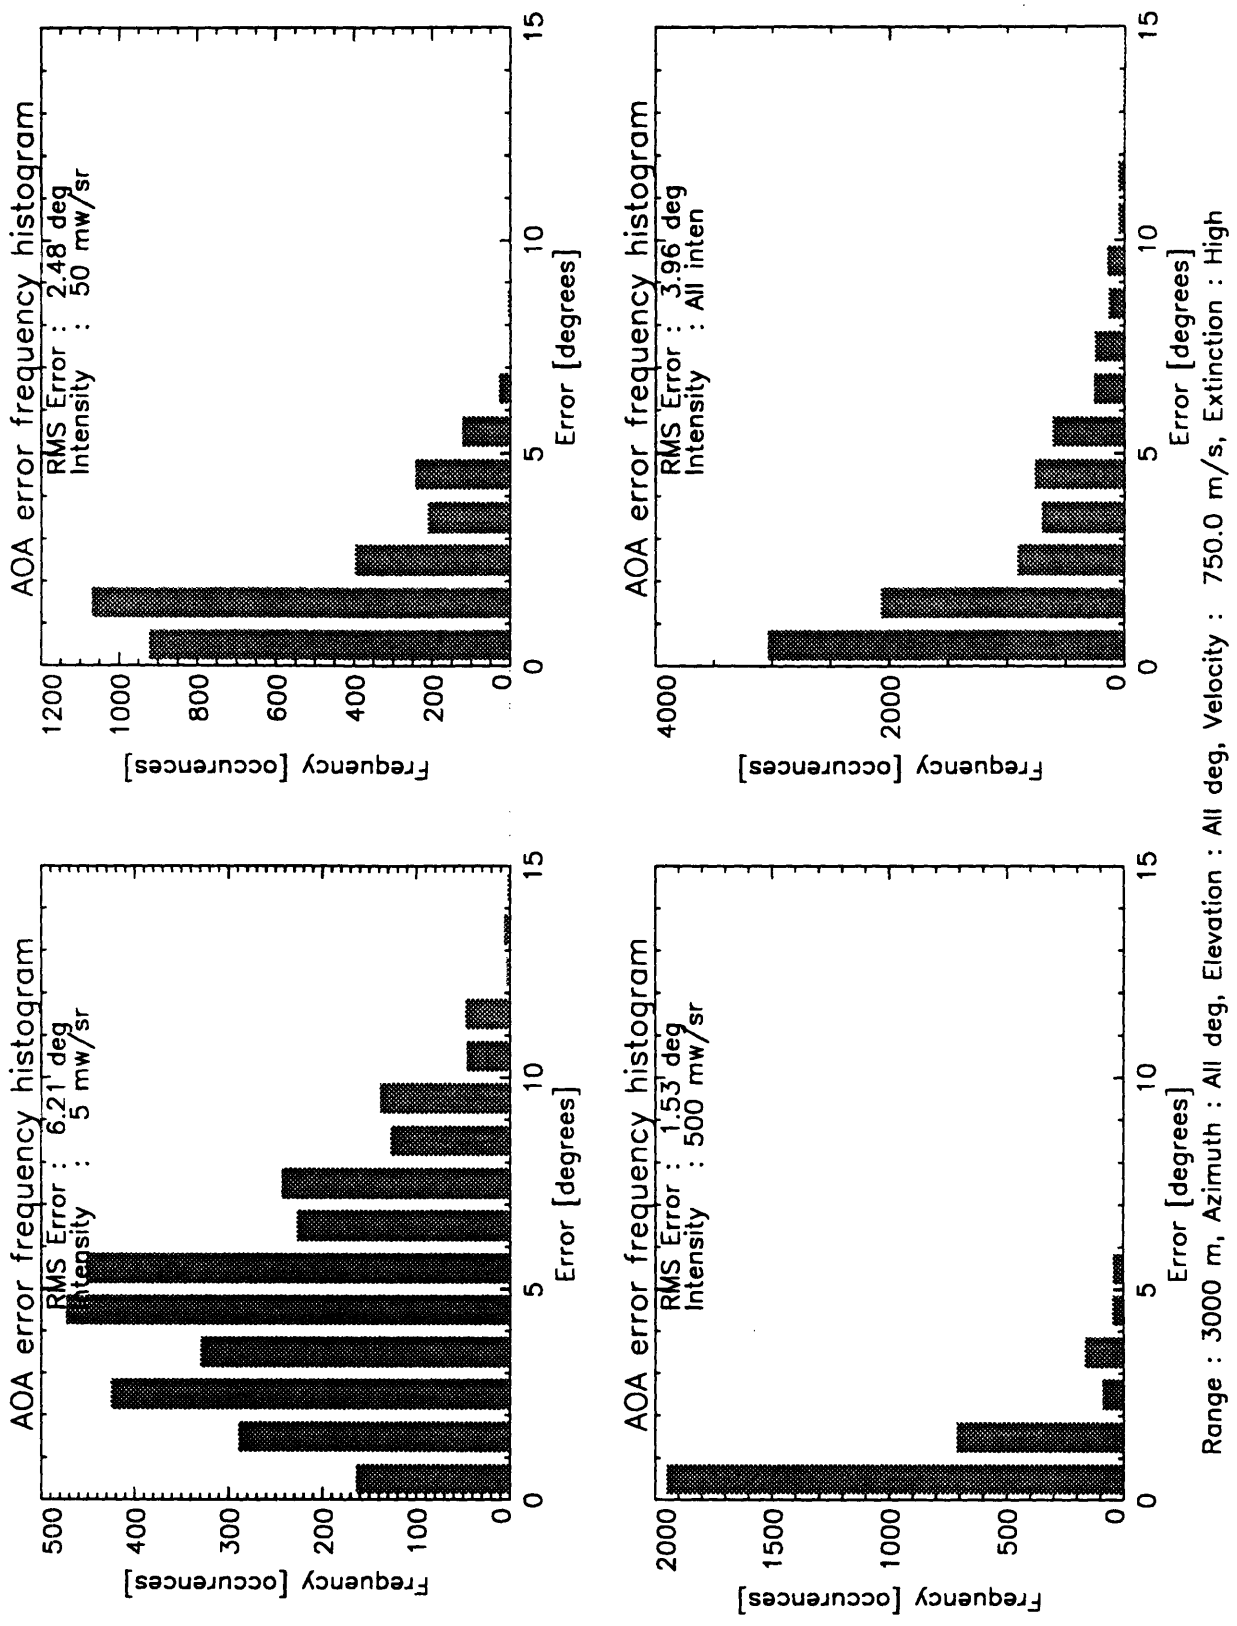

Figure 3.6 - Mid-range AOA estimation distribution for low extinction

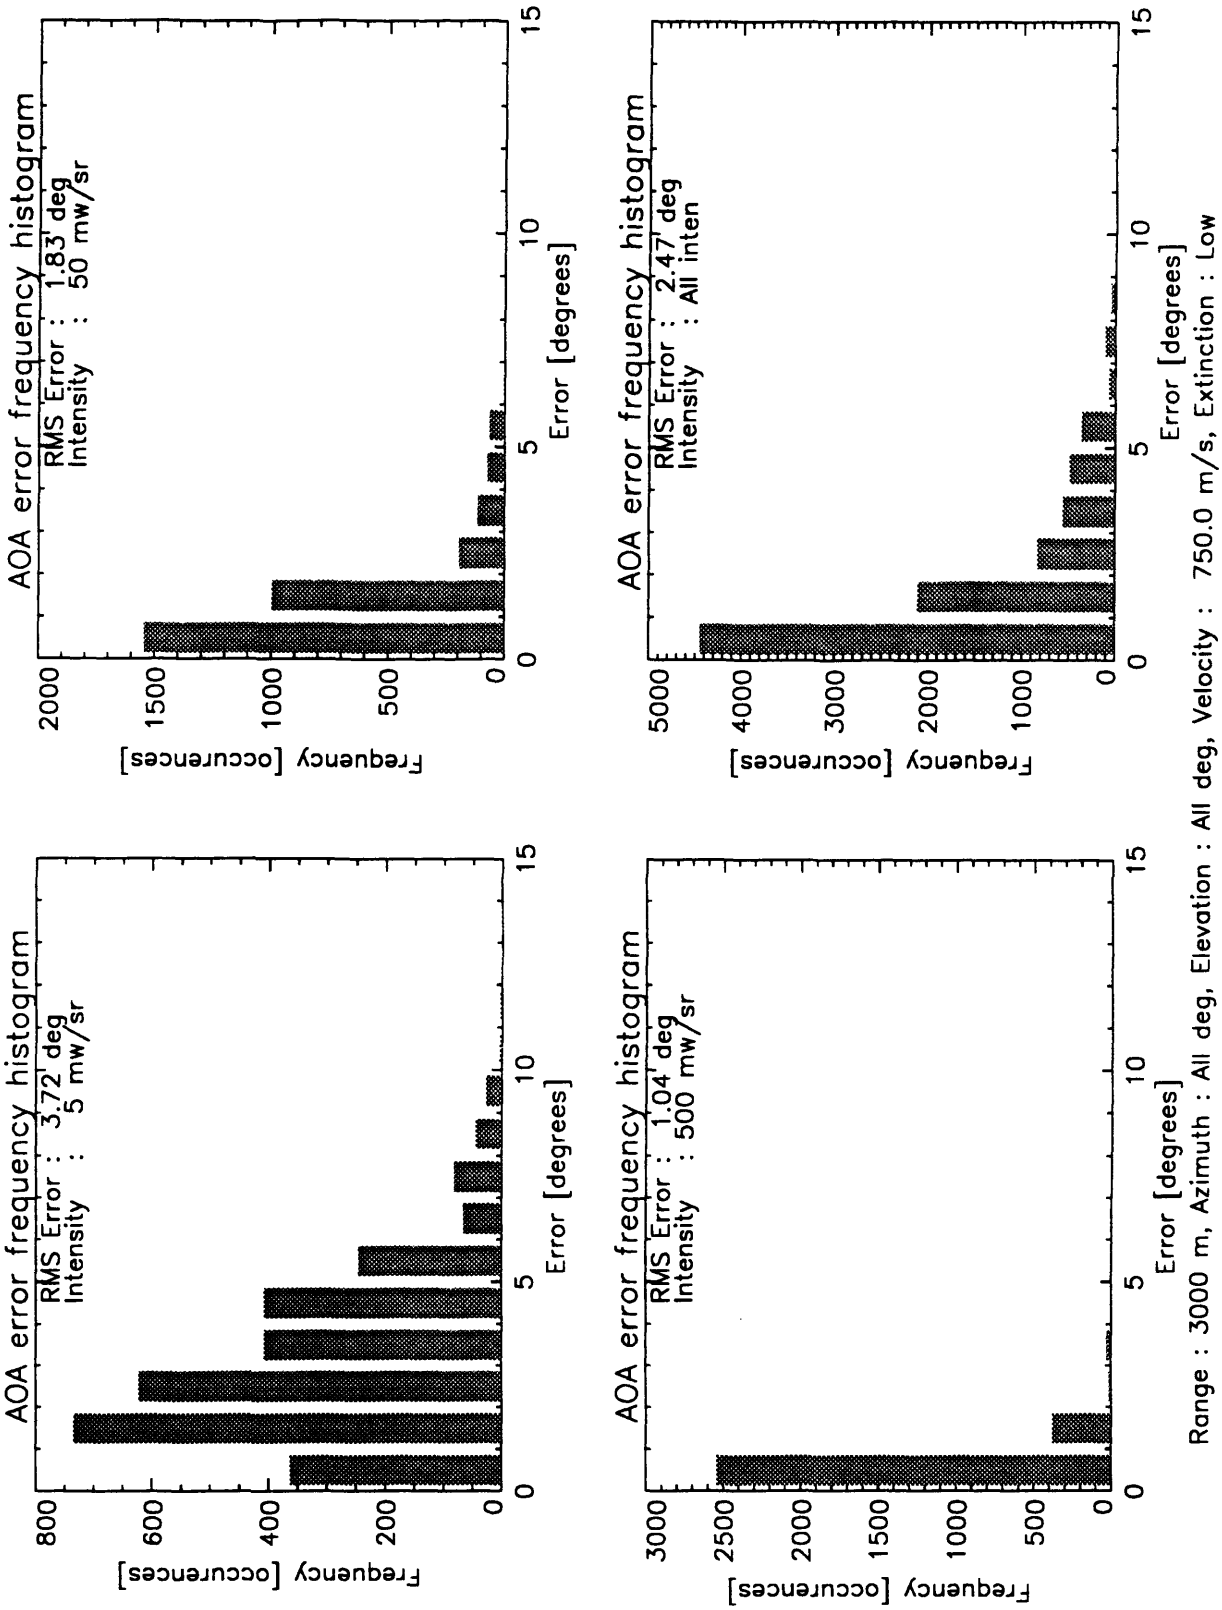

Figure 3.7 - High SNR AOA estimation distribution

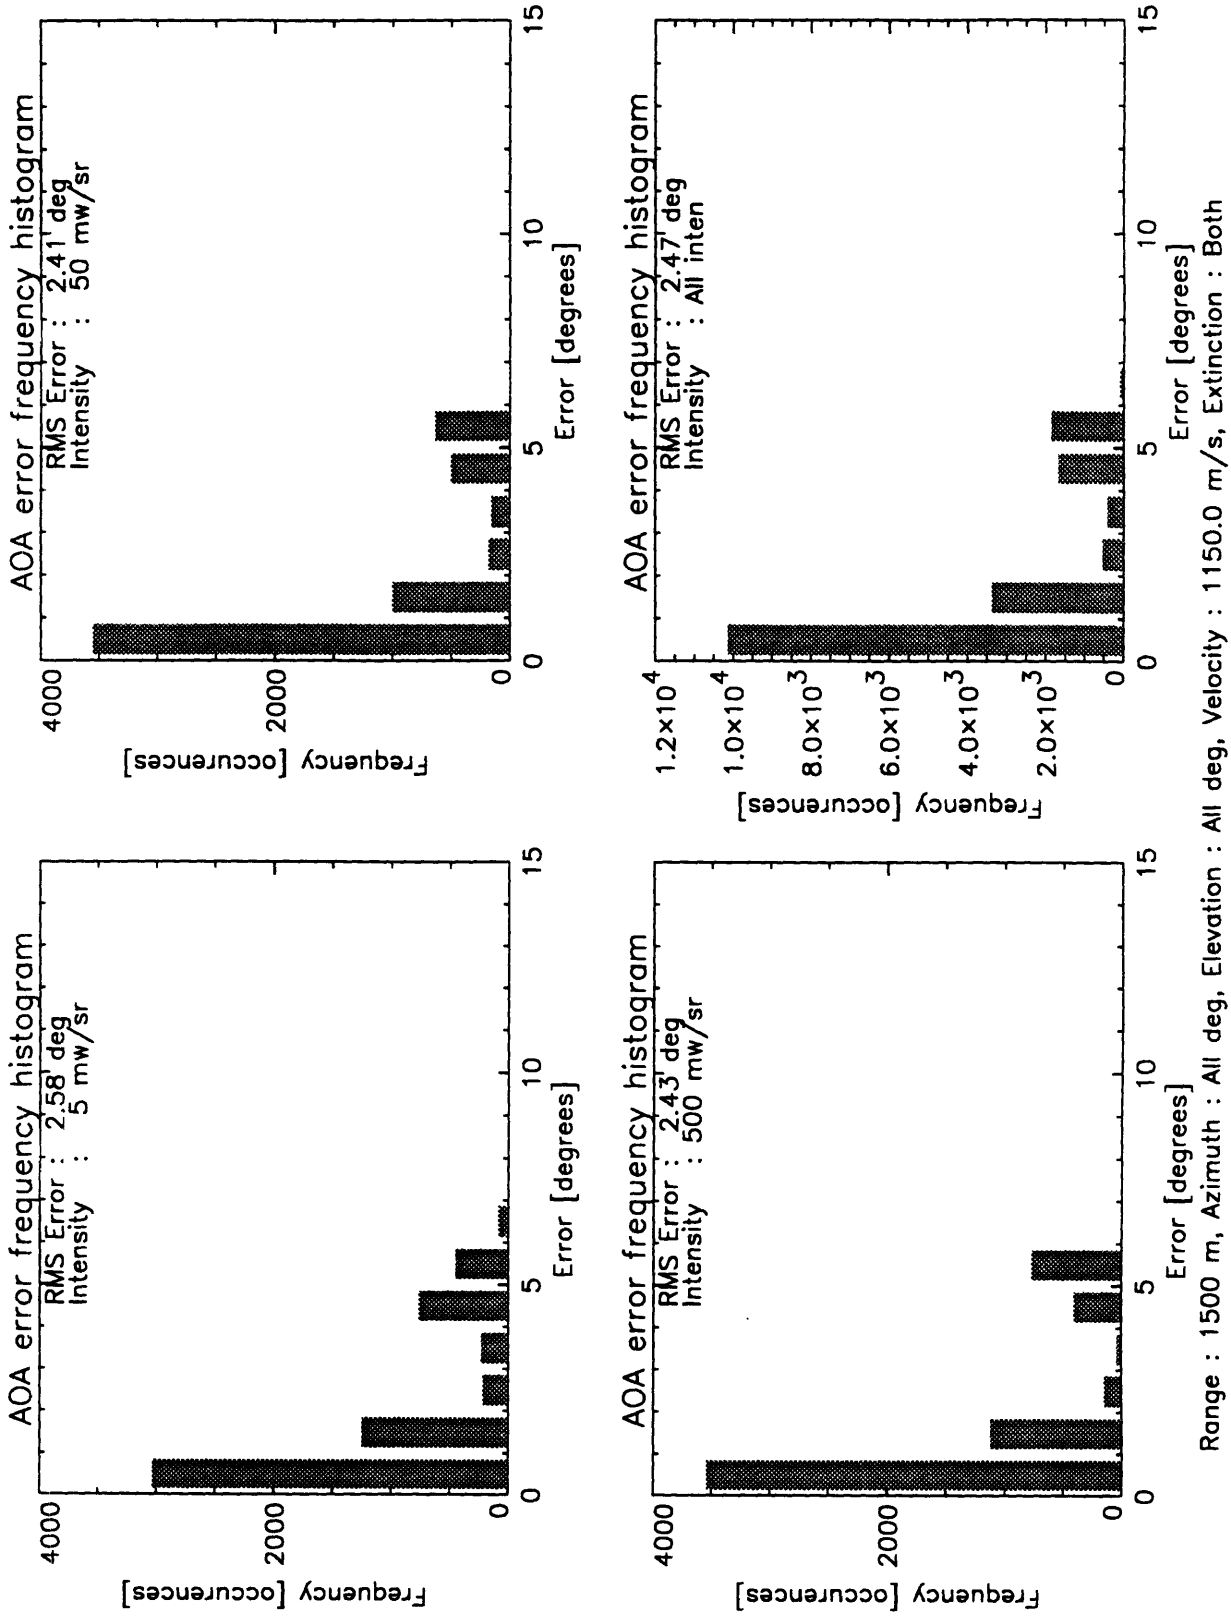

Figure 3.8 - Low SNR AOA estimation distribution

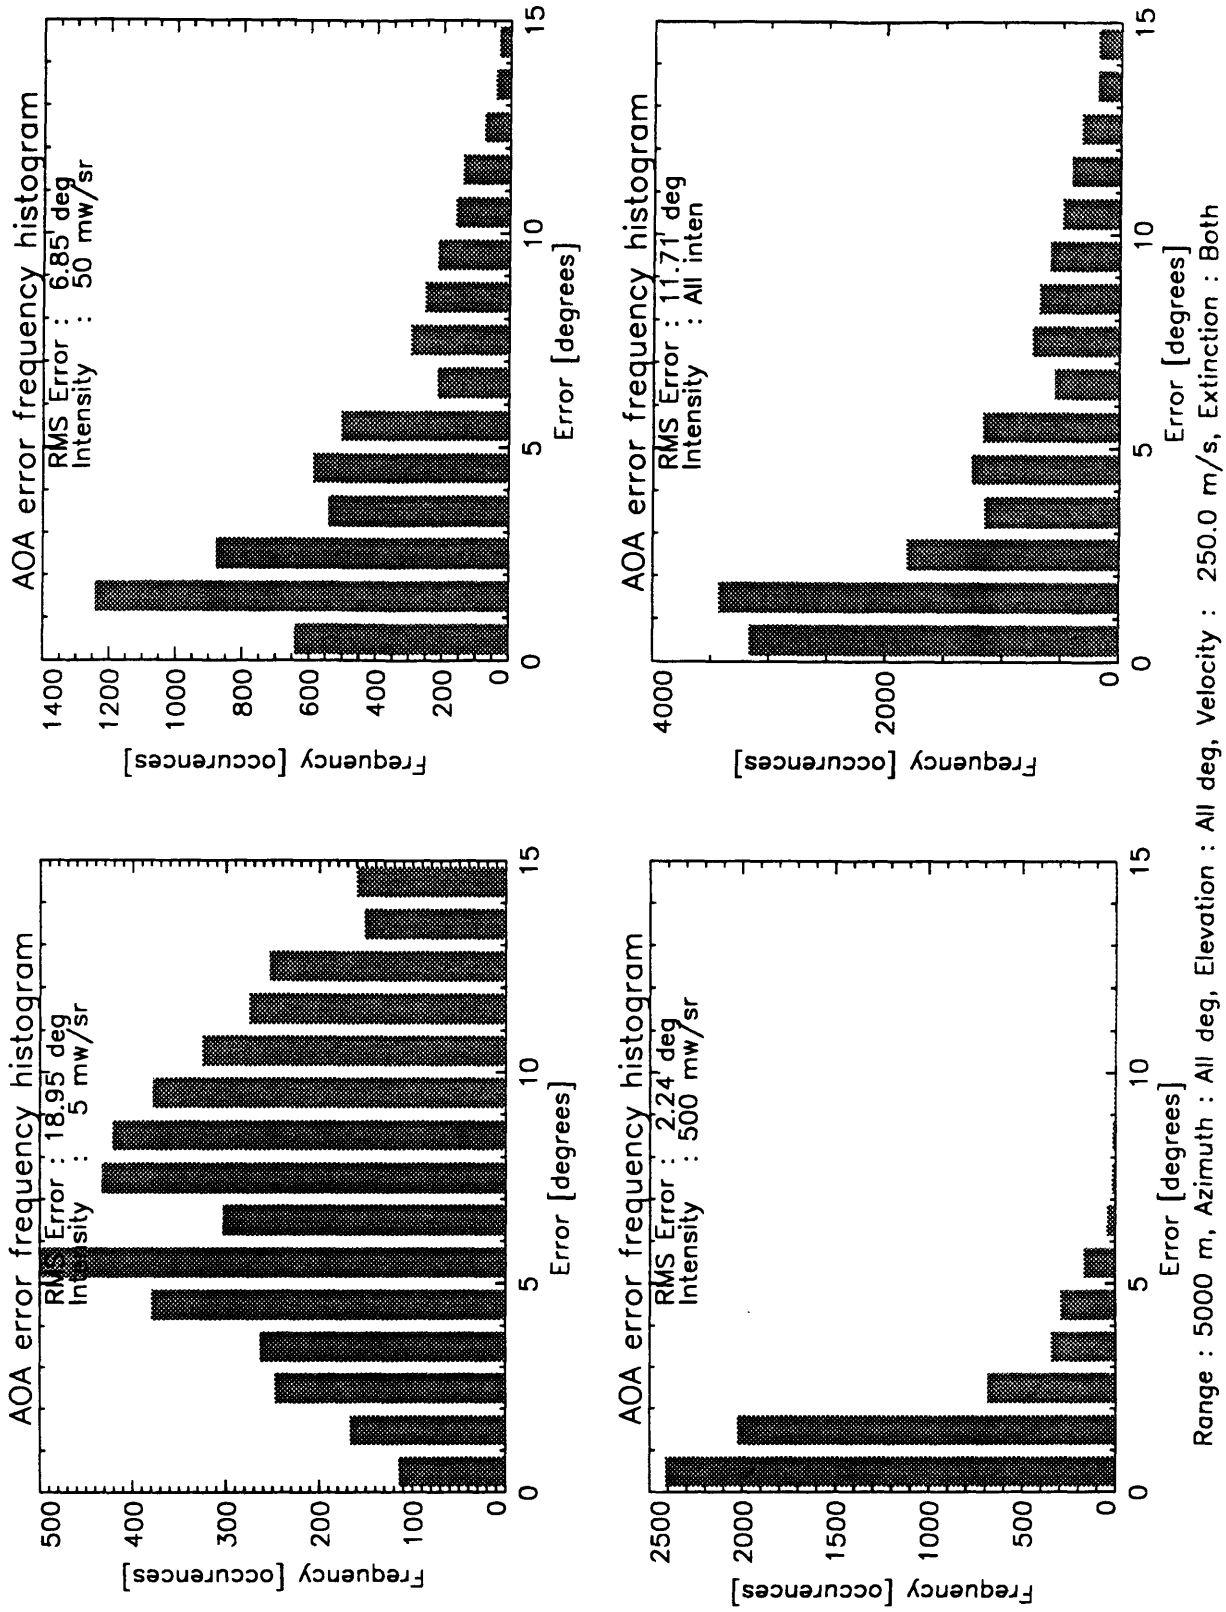

4. Conclusions

4.1 The Real World

Although the algorithm presented here performs reasonably well when applied to model data in controlled situations, there are several issues that would need to be addressed when implementing a real-world system. These issues are presented below, and suggest directions for further research.

4.1.1 Atmospheric modeling

Although atmospherics do not seem to make as big a difference as was thought at the outset of this project, a real implementation would have to be able to deal with more than just the low and high absorption coefficients. Similarly, we have not addressed the effects of precipitation and cloud cover. Anything that affects the way radiation travels from the source to the sensors will have an impact on the accuracy of the estimation system.

We have further assumed that the sensors are at 5000 meters altitude. In real life, aircraft are likely to fly at a variety of altitudes, and model data would have to be generated for these configurations as well.

This raises the issue of ground modeling. The current implementation of the TWES models the ground as a perfect absorber. In reality, this is not the case. When an aircraft is flying at low altitudes, the sensor output is strongly affected by the optical characteristics of the earth. A more sophisticated model would need to be developed that would take the effects of reflection and absorption into account. Further, as the information will generally not be available from on-board instruments, the algorithm must either be made insensitive to changes in ground characteristics or be altered to estimate them as well. This may well cause the problem to be so loosely constrained as to be unsolvable.

Clearly, as we introduce more factors into the model, we must deal with the issue of storage. For the assumptions made in this thesis, sufficient memory was available to simply memorize our atmospheric model as a multidimensional lookup table. This will no longer be an option when additional factors are taken into consideration. Fortunately, there exists a wealth of literature dealing with surface fitting that could be applied to this issue.

4.1.2 Trajectory modeling

In addition to developing a more sophisticated atmospheric/sensor model, a more robust model of the flight trajectory would be needed. We have assumed that both the missile and the aircraft are following simple paths at constant velocities. What would be the impact on TTI and AOA estimation accuracy if one or the other were to change course or direction?

One possible solution to this dilemma would be to shorten the duration over which the algorithm examines the engagement (i.e. use fewer sensor measurements), but this harkens

back to the issues raised concerning the EKF: the more sensor measurements the algorithm has available to examine, the more likely it is to arrive at the correct result. It is worthy of some analysis to see if a more accurate model of flight trajectories might solve the problem.

4.1.3 Radiative characteristics

Two further assumptions we have made are that the missile is an isotropic source, and that the intensity of the source remained constant throughout the period of interest. In reality, a missile is quite clearly a non-isotropic emitter. Fortunately there are several good models available for this phenomenon that could be employed when implementing an actual estimation system.

The second issue, temporal intensity variations, poses a much larger problem. The assumption that intensity is a constant lies at the heart of the MSDM method. Introducing temporal variations completely invalidates this approach, unless they are constrained in such a way that they can be factored out. As mentioned previously, it will probably be possible to detect that such changes are occurring. If a change in intensity were detected, the estimate from our estimation algorithm would have to be discarded, as the constant intensity assumption is central to its operation.

If temporal variations in intensity could not be detected it would not be possible to use MSDM, and we will need to return to an MSE, EKF, or Bayesian approach.

4.1.4 Other

There were other directions of research suggested in the Analysis section of this thesis. These included an examination of under what conditions the atmospheric approximation was a good fit to the original eight modeled atmospheres, and a more detailed look at why AOA accuracy does not always increase with increasing source velocity.

4.2 Summary

In the introduction to this thesis, we stated our goal of providing an estimate of missile trajectory as a means to the end of avoiding missile interception. At that time, we also laid out what qualities we desired our solution to have. These were accuracy, robustness, and real-time execution.

4.2.1 Accuracy

The accuracy of TTI estimation was found to vary quite dramatically - from completely unusable at long ranges for slow shots to quite accurate in close-fast scenarios. This is, however, just the way we would like it to be. Missiles moving at slow speed at a great range are not as immediate a threat as those which are close in and quickly approaching. In this second category of threats, we find that the MSDM algorithm is quite good at accurately approximating the TTI.

AOA estimation was in general good, with most trajectories studied being predicted to within 3 degrees RMS, even at long ranges. Given that the sensors in this analysis are non-imaging, and have quite a wide FOV, this is rather impressive. And although this may or may not be sufficient for directed countermeasures devices, it is far beyond the level of accuracy needed to provide pilot cueing. Very few humans can point to a particular compass heading within three degrees, but the directional indication that the MSDM AOA estimation can provide to the pilot will almost always be accurate to within the range necessary to provide "over there" warning.

4.2.2 Robustness

As stated in the introduction:

> The system must be able to handle a wide variety of threats, from slow to fast, near to far, bright to dim. An enemy will seldom provide his target with information about when, where, what and how he is going to shoot, so the system must be able to accurately predict TTI and AOA over a range of engagement parameters.

As stated in the previous paragraphs, this objective was achieved quite well for AOA, and somewhat less so for TTI in low SNR conditions.

4.2.3 Real-time

Although the system as implemented is far from real-time, requiring 20 seconds of CPU time on a DEC Alpha station to process one second of sensor data, there is reason to believe that a real-time implementation could be achieved.

First, the nature of the oct-tree computation is highly parallelizable. Since the search performs essentially the same computation on several different points in the state space, then compares the results to determine the next iterative step, it could be implemented such that each state-space point is computed simultaneous with the others for a given step of the iteration. This alone would provide a speed increase.

Further, we note that the oct-tree computation iteratively refines its estimates of AOA, computing TTI in the last iteration. The intermediate values of AOA could be made available to the pilot. As stated above, the three-degree RMS accuracy produced by the estimator, although desirable for TTI estimation, is more than sufficient for pilot cueing. An intermediate value of lower resolution would be available far earlier, and could be of sufficient accuracy for the purposes of early warning.

4.3 The Last Word

The Measurement Space Direction Matching method, when combined with the oct-tree search of state space, provides a balance of speed and efficiency that can quickly and accurately estimate Time to Intercept and Angle of Arrival for closing missiles. Although much work remains to be done before a system using these methods could be fielded, this work represents a foundation on which further efforts can be built.

Bibliography

1. S. L. Angle, "An Evaluation of the All-stars Time-to-Intercept (TTI) Algorithm," LIRIS IR&D Report, Aug. 1990.
2. *Atmospheric Propagation Codes of DoD Systems.* Department of Defense Spectral Reflections Technote, April 1991.
3. M. Bulpett, "Threat Warning Missile Database," LIRIS Technical Document, June 1992.
4. B. Carnahan, H. A. Luther and J. O. Wiles, *Applied Numerical Methods.* John Wiley & Sons, Inc., 1969.
5. F. E. Daum, "Exact Finite Dimensional Nonlinear Filters," *Proceedings of 24th Conference on Decision and Control.* Dec. 1985.
6. F. E. Daum, "Exact Nonlinear Filters, Exponential Densities and Separation of Variables," *Allerton Conference on Communication, Control, and Computing.* Oct. 1985.
7. F. E. Daum, "Exact Finite Dimensional Nonlinear Filters for Continuous Time Processes with Discrete Time Measurements," *Proceedings of the 1984 IEEE Conference on Decision and Control.* Dec. 1984.
8. A. Drake, *Fundamentals of Applied Probability Theory.* McGraw-Hill Book Company, 1967.
9. R. V. Hogg and E. A. Tanis, *Probability and Statistical Inference.* Macmillian Publishing Company, 1983.
10. E. Huthchinson, "The Kalman Filter Applied to Aerospace and Electronic Systems," *IEEE. Transactions on Aerospace & Electronic Systems.* Jul. 1984.
11. J. S. Lim, *Two-Dimensional Signal and Image Processing.* Prentice Hall, 1990.
12. L. Ljung and T. Soderstrom, *Theory and Practice of Recursive Identification.* MIT Press, 1983.
13. M. Luettgen, "Trajectory Estimation of an Optically Radiating Source," M.S. Thesis, Dept. of EECS, M.I.T., May 1990.
14. E. McCartney, *Optics of the Atmosphere.* John Wiley, 1976.
15. A. Nehorai and M. Morf, "Estimation of time difference of arrival for multiple ARMA sources by pole decomposition," in Proc. IEEE CDC, Orlando, FL, pp. 1000-1002, 1982.
16. A. V. Oppenheim and R.W. Schafer, *Discrete-Time Signal Processing.* Prentice Hall, 1989.

17. A. Papoulis, *Probability, Random Variables and Stochastic Processes.* McGraw Hill Book Co., 1984.

18. W. S. Ross, "Point-Source Optical Propagation in a Multiple Scattering Medium," Ph.D. Thesis, Dept. of EECS, M.I.T., Aug. 1980.

19. W. S. Ross, W. Jaeger, J. Nakai, T. Nguyen and J. Shapiro, "Atmospheric Propagation of Ultraviolet Radiation through the Lower Atmosphere," Aeronautical Research Associates of Princeton Report No. 374b, November 1978.

20. E. B. Saff and A. D. Snider, *Fundamentals of Complex Analysis for Mathematics, Science and Engineering.* Prentice Hall, 1993.

21. G. C. Schmidt, "Designing Nonlinear Filters Based on Daum's Theory," *Journal of Guidance, Control, and Dynamics.* Vol. 16, No. 2, pp. 371-76.

22. D. H. Staelin, A. W. Morgenthaler, J. A. Kong, *Electromagnetic Waves.* Prentice Hall, 1994.

23. H. C. Van de Hulst, *Light Scattering by Small Particles.* John Wiley, 1957.