SAND--89-0816

DE91 002586

SAND89-0816
Unlimited Release
Printed October 1990

Distribution
Category UC-515

Optimal Sensor Fusion for Land Vehicle Navigation

J. Dan Morrow
Advanced Technology Division 5267
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

Position location is a fundamental requirement in autonomous mobile robots which record and subsequently follow x,y paths. The Dept. of Energy, Office of Safeguards and Security, Robotic Security Vehicle (RSV) program involves the development of an autonomous mobile robot for patrolling a structured exterior environment. A straight-forward method for autonomous path-following has been adopted and requires "digitizing" the desired road network by storing x,y coordinates every 2m along the roads. The position location system used to define the locations consists of a radio beacon system which triangulates position off two known transponders, and dead reckoning with compass and odometer. This paper addresses the problem of combining these two measurements to arrive at a best estimate of position. Two algorithms are proposed: the "optimal" algorithm treats the measurements as random variables and minimizes the estimate variance, while the "average error" algorithm considers the bias in dead reckoning and attempts to guarantee an average error. Data collected on the algorithms indicate that both work well in practice.

Acknowledgements

I would like to thank Frank Wunderlin for writing the software to implement these algorithms, and collecting all the data with the Robotic Security Vehicle. I acknowledge John Harrington for his contributions to the ideas presented in this paper. I thank Steve Roehrig for his many valuable and constructive comments. And, finally I thank Connie Vanderburg for formatting and printing the many revisions of this report.

TABLE OF CONTENTS

FIGURES

Optimal Sensor Fusion for Land Vehicle Navigation

Introduction

The Advanced Technology Division 5267 at Sandia National Laboratories has developed a number of prototype mobile robotic systems for several different applications. A current project for the Dept. of Energy, Office of Safeguards and Security (OSS), involves developing a robotic security vehicle capable of autonomously patrolling a structured exterior environment.

For the program we have adopted a simple strategy to accomplish autonomous road-following which requires "mapping" the roads. Thus, there are a discrete number of roads which the vehicle can traverse autonomously, and these are stored as lists of x,y points. In order to map the roads, a position location system is needed to determine the x,y position of the vehicle.

It is the position location problem which this paper addresses. First the position location problem is defined. The algorithms developed to solve the problem are then discussed and collected data on their performance is presented.

Problem Description

Position location is a fundamental requirement in mobile robotics and is critical to an autonomous mobile robot which needs to plan and execute a path to a desired destination. For this paper, the position is defined as a 2-Dimensional position consisting of x,y and heading. In addition, experience has shown that autonomous operations such as path following require a fairly accurate (+/- 0.5 m) determination of x,y position. For our applications in physical security, the area of coverage can be fairly small (5-30 km square).

There are many methods of determining, or measuring, position. These range from navigational satellites, radio beacon triangulation systems, laser triangulation systems, dead reckoning with a compass and odometer, to inertial navigation systems. To minimize costs, the system that Sandia has been using on the OSS Robotic Security Vehicle project is a combination of radio beacon and dead reckoning.

Dead reckoning uses odometer and heading information to compute a new position from the robot's old position. This method must be initialized to some global position and then propagates the global position from the heading and distance travelled. This method is attractive because it is self-contained on-board the vehicle and is inexpensive. Unfortunately, the dead reckoning errors accumulate and the position estimate "drifts" with distance travelled. This drift error is a low frequency error, meaning that repeated readings at the same location will yield very nearly the same result, but moving off that location and returning will yield different results. Another attribute of a low frequency error is that the error between the robot's true position and the estimated position will steadily and smoothly diverge.

Because of the drift associated with dead reckoning error, an external reference system is often used to remove that error. A radio beacon system consisting of two transmitters and a receiver is one such system. By triangulating off known locations of the transmitters, an absolute position can be computed.

The radio beacon position estimate is not a function of distance travelled, but is corrupted by "random" high frequency noise. This means that repeated measurements at the same position give different values, resulting in a large variance in the measurement. The system is fairly stable so that measurements at a given location that are spread out in time will give very similar results within the "jitter" described above; i.e. this measurement does not drift with time or distance travelled. Data collected on our radio beacon system showed a 6 sigma spread in x and y values of about 1 meter with a normal distribution shape; see Appendix B for a summary of the data. Note that a 1 meter 6σ spread is too large for autonomous path following.

Ideally, an on-board position location system such as dead reckoning is preferred for a mobile robot because it requires no outside support hardware and is more covert and secure. Unfortunately, dead reckoning also drifts with distance travelled and therefore must be "corrected" by an external source. Conversely, the radio beacon system gives an accurate position estimate which is corrupted by a zero-mean, high-frequency noise. This results in a position estimate which "jumps around" when repeated readings are taken at the same location. For recording a path, this characteristic can make the path look very jagged when it is in fact very smooth. Thus, combining the radio beacon measurement with the dead reckoning measurement (relatively "smoother") will smooth out the jaggedness (lower the variance) in the position location estimate.

In an effort to study how these two position estimates should be combined, two fusion algorithms were developed.

Fusion Algorithms

The first approach to a fusion algorithm was a "Kalman-filter" type approach which modelled both measurements of position (dead reckoning and radio beacon) as unbiased, normally-distributed random variables. The variance of the linear combination was minimized to yield the weighting factors on each measurement.

The second approach was undertaken to account for an inaccurate assumption in the first approach: namely, that the dead reckoning measurement gives an unbiased measurement of position. Looking at traces of x,y position as recorded by radio beacon and dead reckoning measurements show clearly the dead reckoning measurements diverging from the radio beacon measurement. Thus, in an effort to account for the bias in dead reckoning, a second fusion algorithm was developed which models the dead reckoning measurement as a biased variable with no variance. The bias term is a linear function of distance travelled.

For both fusion algorithms, equations were developed which related the measurement fusion weighting factor to the distance travelled between fusions and the dead reckoning drift rate. For the first algorithm, both dead reckoning and radio beacon measurements are assumed to be unbiased estimates of the true position with variances. For the second algorithm, the radio beacon measurement was modelled the same, but the dead reckoning measurement was modelled as a biased estimate with no variance.

Basic Fusion Equation

The basic fusion equation for combining the two measurements is:

$$z = \alpha*(DR) + (1-\alpha)*(RB)$$

where z is the fused estimate of the x (or y) coordinate, (DR) is the dead reckoning measurement, (RB) is the radio beacon measurement, and α is the fusion weighting factor. Because it is anticipated that the x,y coordinate variances will be approximately equal, the same value of α is used for both x and y coordinates.

Using the basic fusion equation described above, the navigational procedure is as follows:

1) The x,y position is initialized to some global position. This can be done by a variety of methods, one of which is to use the radio beacon system.
2) The dead reckoning calculation updates the position quite frequently (exactly how often depends on computational time available, and resolution of the odometer). The change in distance since the last dead reckoning update is used to determine when a new update should be performed. When this change in distance exceeds some minimum value, changes in x and y are computed based on the change in distance and vehicle heading. These changes in x and y are added to the current estimate of x,y position. In this way the x,y position is propagated by dead reckoning.
3) Occasionally, a radio beacon measurement of x,y position becomes available. This is assumed to be much less frequent than the dead reckoning updates. When a radio beacon measurement is available, there are now two estimates of position: the radio beacon measurement, and the dead reckoning propagation measurement.
4) The fusion algorithm is used to combine these measurements to arrive at a "best" estimate of position. The basic fusion equation is simply a linear combination of the measurements. The number α determines the relative weighting of each measurement in the combination.
5) When the measurements are combined by the fusion algorithm, the resulting x,y "best" estimate overwrites the dead reckoning x,y location. In this way, the dead reckoning is always propagating the best estimate of current position. Also, by overwriting the dead reckoning x,y estimate, the accumulated error in dead reckoning is removed.
6) In a sense, the "dead reckoning" measurement really contains both dead reckoning propagation terms, but also previous radio beacon measurements as well. It is simply the current position estimate which is propagated by dead reckoning and corrected by the radio beacon measurements.

The dead reckoning is always propagated from the fused estimate since this represents the best estimate of current position available. Thus, (DR) is overwritten with z, the fused estimate. Then (DR) is propagated with dead reckoning calculations until another radio beacon measurement, (RB), is available. It is important to note that because (DR) is overwritten with the latest fusion, (DR) contains "parts of" previous (RB) measurements. Thus, (DR) is not strictly dead reckoning but contains previous (RB) data as well. (DR) is referred to as the "estimate" or the "position estimate," and represents the current best estimate of location.

Assuming two normally distributed variables, (DR) and (RB), the fused estimate variance is given by:

$$P = \alpha^2 * P_{dr} + (1-\alpha)^2 * P_{rb}$$

Where P_{dr} is the dead reckoning measurement variance, and P_{rb} is the radio beacon measurement variance.

Just as (DR) is overwritten with the fusion result, so P_{dr} is overwritten with P, the fusion estimate variance. Thus, the "dead reckoning" variance always contains the variance of the current estimate of positon.

Fusion Algorithm #1: Optimal

The first fusion algorithm for combining a dead reckoning estimate of an x or y coordinate with the radio beacon system estimate of the x or y coordinate assumes that these estimates are unbiased estimates of the coordinate with normal distributions. While the radio beacon measurement is assumed to have a constant variance, the dead reckoning measurement variance increases with distance travelled. The starting point for this algorithm is the relationship between the dead reckoning variance and the dead reckoning drift rate.

For a normally-distributed measurement, the probability distribution of measurements is the familiar "bell"-shaped curve. While theoretically the probability distribution extends to $+\infty$ and $-\infty$, the practical limits are usually taken at $\pm 3\sigma$, since this includes 99.8% of a normal distribution population. The $\pm 3\sigma$ spread is referred to here as the "range" of the measurement since it indicates what difference between the maximum and minimum measurements can be expected. So the range of the measurement (R) can be written as the following function of standard deviation, σ:

$$R = 6\sigma$$

The range can sometimes be a more "intuitive" indicator of the spread in a measurement than variance (or standard deviation). Often a measurement is described as Y $\pm$ x, where x is usually referred to as the "uncertainty" of the measurement. As the measurement becomes less accurate, or more uncertain, x increases. So the range can be interpreted as the uncertainty in a measurement.

Since we know that the dead reckoning propagation of position drifts with distance travelled, the uncertainty (or range, R) of the estimate also grows with distance travelled, so let:

$$\Delta R = Q \, \Delta s$$

where ΔR is the change in range (or uncertainty), Q is the dead reckoning drift constant, and Δs is the distance travelled since the last radio beacon update.

So we will increase the uncertainty of the estimate as we propagate : with dead reckoning according to $\Delta R = Q \, \Delta s$. Now we have to relate the ΔR change to a change in variance.

We start out with the relationship between range, R, and the standard deviation, σ:

$$R = 6\sigma,$$

and solve for σ:

$$\sigma = R/6.$$

Next, recall that variance and standard deviation are related as follows:

$$P = \sigma^2,$$

so we can write variance as a function of R:

$$P = (R/6)^2$$

We can differentiate this equation to relate dP to dR:

$$dP = (R/18) \, dR.$$

Substituting the expression for R as a function of P:

$$R = 6 \sqrt{P},$$

gives:

$$dP = \frac{\sqrt{P}}{3} \, dR.$$

Finally, noting that $dR = Q \, ds$, we can write:

$$dP = \frac{\sqrt{P}}{3} \, Q \, ds.$$

This equation allows us to propagate the variance as the position estimate is propagated with dead reckoning calculations. When the radio beacon measurement is made, there are two position estimates: one from the radio beacon, and one from dead reckoning propagation. Each of these measurements has a variance associated with it and so the fused result of the two also has a variance.

The variance of the "fused" estimate can be minimized by setting $dP/d\alpha = 0$ to give:

$$\alpha = \frac{P_{rb}}{P_{rb} + P_{dr}}$$

Where P_{rb} is the radio beacon measurement variance, and P_{dr} is the dead reckoning variance before fusion. After fusion, the dead reckoning variance is over-written with the smaller fused estimate variance. This is very similar to the optimal fusing in Kalman filtering and is described in [1].

This fusion algorithm differs from Kalman filtering in that a dynamic model is not used to propagate the measurement in between updates. Instead, we assume that the dead reckoning calculation (propagation of position) is done much more often than the radio beacon measurement is taken. Thus, we do not have a dynamic model, but rather two sets of measurements, one more frequent than the other. In a sense, the dead reckoning calculation acts like the dynamic model since it propagates the position estimate in between updates with the radio beacon measurement.

The variance of the estimate grows when propagated with dead reckoning in between radio beacon updates and then falls when the radio beacon information is fused with the dead reckoning propagation. The result is a "saw-tooth" curve of the position estimate variance with distance travelled (see Figure 1). This algorithm is easily implemented by propagating the position estimate variance and using it to compute α.

Figure 1 has three plots of variance vs distance. Each curve varies the distance between radio beacon updates: 0.25m, 1.0m, and 2.0m. The variance is propagated from an initial value and reaches a steady-state condition after only about 15m. The variance grows steadily as the position estimate is updated with dead reckoning information, but drops abruptly when fused with the radio beacon information. Note that the higher the radio beacon update rate, the lower the steady-state variance values and the smaller the variance drop when fused. Frequent radio beacon updates provide more information for an improved estimate (lower variance), and allow less time for the variance to increase during dead reckoning propagation.

Figure 2 shows the propagation of the fusion weighting factor, α, with distance travelled for the same 3 radio beacon update rates used in Figure 1. Note that the values of α are only used during fusion with the radio beacon data, and at no time during dead reckoning propagation. The higher the radio beacon update rate, the higher the value of α used. The value of α is the weighting on the dead reckoning measurement during fusion, so a higher value of α indicates more weight on dead reckoning and less on the radio beacon. For frequent radio beacon updates, each individual radio beacon measurement can be weighted less because of the quantity of measurements that will be taken. Furthermore, a frequent radio beacon update rate implies the dead reckoning variance growth between updates is small, so the dead reckoning position estimate drifts little between radio beacon updates and is quite accurate.

The actual performance of the optimal algorithm will be discussed later in the results section after the next algorithm is discussed. The second algorithm considers the fact that, in reality, dead reckoning is not an unbiased estimate of position. Banta [2] reports that many of the error sources in dead reckoning are not zero-mean. For example, the dead reckoning calculations used here assume a 2-D world; any hills or valleys will bias the dead reckoning "long."

Thus, a dead reckoning measurement gives a biased estimate of position. The "optimal" weighting algorithm described above minimizes the variance of the fused position estimate at the cost of propagating the dead reckoning error. The next algorithm considers the bias in the dead reckoning.

Fusion Algorithm #2: Average Error

The second fusion algorithm assumes a similar radio beacon to the optimal algorithm assumption, while the dead reckoning is modelled as a biased estimate of position with no variance. Thus, the drift in the dead reckoning is accounted for as a bias to be added to the dead reckoning estimate instead of a variance. So the radio beacon measurements have a variance but no bias, and the dead reckoning measurements have a bias but no variance.

The fused estimate has some amount of bias due only to the dead reckoning term as well as some variance due to the radio beacon term associated with it. We assume that the variance associated with the radio beacon measurement is far larger than the variance associated with the dead reckoning measurement, and therefore neglect any variance in the dead reckoning measurement. This is done because the dead reckoning is "smooth" compared with the radio beacon.

Again, let the fusion equation be written as:

$$z = \alpha*(DR) + (1-\alpha)*(RB)$$

where z is the fused estimate of the x (or y) coordinate, (DR) is the dead reckoning measurement, and (RB) is the radio beacon measurement.

While the optimal fusing algorithm had only one criterion- minimizing the variance of the fused estimate, this second algorithm has two criteria: 1) minimize the variance of the fused estimate, and 2) minimize the error in the fused estimate. Unfortunately, these goals are mutually exclusive. Since the radio beacon measurement has no error and the dead reckoning measurement has an error, minimizing the error in the fused estimate implies using only the radio beacon measurement ($\alpha=0$). Since the dead reckoning measurement is assumed to have zero variance (compared to the radio beacon measurement), minimizing the variance of the fused estimate implies using only the dead reckoning measurement ($\alpha=1$). Neither of these extremes is acceptable because they result in either zero variance and full error, or zero error and full variance.

An important result of fusing the dead reckoning and radio beacon measurements as described above with a weighting factor, α, is that the steady-state error is bounded for a constant Δs even with a constant non-zero bias in the dead reckoning measurement. The worst case dead reckoning bias is a constant which results in an error which always grows. It is shown in Appendix A that the minimum and maximum steady state errors for a constant drift rate (Q) and a constant Δs can be written as:

$$e_{min} = Q\ \Delta s\ \alpha/(1-\alpha)$$

$$e_{max} = Q\ \Delta s\ 1/(1-\alpha)$$

and the average error as:

$$e_{avg} = Q\ \Delta s\ (1+\alpha)/(1-\alpha)$$

where Δs is the distance travelled between radio beacon fusion updates.

What is desired is an expression which relates the fusion weighting factor, α, to parameters associated with the dead reckoning measurement (drift rate Q, distance travelled since last fusion, Δs), and the radio beacon measurement (variance, P_{rb}). Since minimizing both the variance and the error of the estimate is not possible, some other method is required. This is a classic "trade-off" problem where maximizing one performance criteria minimizes the other.

The method chosen here is to compute the fusion factor α for a particular steady-state average error, and then examine the resulting variance. Two equations have been derived which relate the variance and error of the estimate to the measurement parameters.

The equation for the average steady-state error, when solved for α, gives the following expression:

$$\alpha = \frac{e - Q\Delta s/2}{e + Q\Delta s/2}$$

where e is the average steady-state error, Q is the constant bias dead reckoning drift rate, and Δs is the distance travelled between radio beacon updates. This provides a relationship between the error and the fusion factor α.

Note that since $\alpha>0$,

$$e > Q\Delta s/2$$

which implies that for a given dead reckoning performance (Q) and available update rate (Δs), there is a limit to how small an average error can be.

In general, this should be a very conservative relationship which over-estimates the error since it assumes the worst-case-- constantly increasing error. The drift rate is very sensitive to the compass calibration, the operating environment, calibration of the odometer, and many other factors. In general, the bias will at times add and other times subtract from the accumulated error.

The relationship between the variance and the fusion factor α is given by:

$$P = (1-\alpha)^2 P_{rb}$$

where P is the variance of the estimate, and P_{rb} is the variance of a radio beacon measurement.

These two equations give relationships between the fusion factor α and the dead reckoning parameters, drift rate Q and distance between updates Δs, and the variance of the radio beacon measurement. In particular, the fusion factor α can be varied with Δs, accounting for the additional error accumulation if radio beacon updates are delayed. Also the effect of α on the variance of the estimate can also be determined. These two equations can be used to modify α to satisfy an error criterion or variance criterion, but not necessarily both.

In order to view the α vs Δs relationship for various values of dead reckoning drift rate Q, the above equation was used to generate a family of curves for a constant average error=0.5m. These curves are plotted in Figure 3. As expected, the value of α decreases as Δs increases, indicating that the dead reckoning is weighted less as the accumulated error in it increases. Large values of Δs would most likely result in an unsatisfactory result since a large Δs implies that much error is accumulated between radio beacon updates and therefore there would be a large difference in the position before and after fusion.

Figure 4 has two average error algorithm curves of α vs Δs for different values of the average error parameter (0.1m and 0.5m). In addition, it has the optimal algorithm steady-state α vs Δs so that the two algorithms may be compared. The algorithms actually give similar α vs Δs curves for Q=5% and e=0.5m. Thus, we would expect similar performance from them. Tightening up the error criterion to 0.1m, however, gives a much lower α value for the average error algorithm, resulting in a much lower weight on the dead reckoning measurement.

Discussion of Results

Data was collected to characterize the parameters in both the dead reckoning measurements and the radio beacon measurements. The radio beacon data is summarized in Appendix B. The dead reckoning drift rate appears to be approximately 5% of distance travelled based on empirical measurements with the radio beacon serving as the "true" position measurement.

A path was driven and the following data was recorded: raw radio beacon measurement, raw dead reckoning measurement (propagated from the initial point with no fusion with radio beacon), the "optimal" fusion estimate, and the average error fusion estimate. This data was collected for two different radio beacon update rates: 0.2s and 1s. The path was driven at a relatively constant speed of 5 mph (about 2.2 m/s), so the corresponding distance between updates was approximately 2.2m for 1s and 0.44m for 0.2s.

Figure 5 shows the α value as propagated during actual runs at different radio beacon rates (0.2s and 1s) for both algorithms. The optimal algorithm α propagation will be discussed first, followed by the average error algorithm α propagation. For the 1s update rate (Δs=2.2m), the measured value of "steady-state" optimal α (0.69) agreed well with theoretical predictions (Figure 2). The 0.2s update rate (Δs=0.44m) value of optimal α (0.88) also agreed well with theoretical predictions. For the average error algorithm, values of α for both update rates agreed well with the theoretical predictions in Figure 3.

Figures 6 and 7 show x,y traces of position location during runs of update rates 1s and 0.2s respectively. Each plot has 4 different curves for position location: dead reckoning only, radio beacon only, optimal fusion algorithm, and average error fusion algorithm. Both plots are "snapshots" in the middle of longer runs; these "snapshots" are small enough to show the necessary detail. Both plots clearly show the dead reckoning position estimate diverged from the radio beacon measurement. As the radio beacon measurement is roughly the "true" position, it is clear that a dead reckoning position location system alone accumulates much error.

Figure 6 clearly shows the jaggedness in the fusion algorithms when the data is fused. The size of the "sawtooth" in the fusion curves is important when using the information for path-following; too large of a sawtooth will result in unacceptable oscillations while path-following. These "sawteeth" are noticeably absent in Figure 7, which has the higher update rate. The higher update rate does not allow much error to accumulate, and so the fusion estimates are always fairly near the "true" radio beacon measurement.

Figure 7 also clearly shows the "jaggedness" of the radio beacon measurement trace. It is this roughness which the optimal algorithm seeks to minimize. Notice that both fusion algorithm traces in Figure 7 are much smoother than the radio beacon trace.

The last item to note in Figures 6 and 7 is that the optimal algorithm trace is closer to the radio beacon trace than the average error algorithm trace. This is expected since the optimal algorithm weighted the radio beacon measurement more (lower α) for both update rates. Both algorithms give a position estimate trace that is close to the radio beacon measurement trace, and smoother than the radio beacon measurement trace.

Summary/Conclusions

This paper has focused on methods to fuse two position location measurements, dead reckoning and radio beacon, to derive a "best" estimate of position. The first algorithm, called the "optimal" algorithm, treats both measurements as random variables and seeks to minimize the variance of the estimate. The second algorithm considers the biases in dead reckoning and attempts to guarantee a maximum average error. Both algorithms vary the fusion weighting with distance travelled between radio beacon updates in order to account for the error accumulation in dead reckoning. Collected data show that both algorithms work well in practice by removing error from the position estimate and providing a smoother position estimate x,y trace.

The equations developed provide a tool to investigate the trade-offs in accuracy, variance, and update rate of two or more position location systems. GPS, laser triangulation, INS, and/or other types of position location systems can be studied with these techniques. In particular, by combining several lower-cost, medium-accuracy position location systems, the desired accuracy may be achieved. Using multiple systems also increases the fault-tolerance of the system; losing one system may lower the accuracy, but position location capability will be retained.

Future work will involve exploring in more depth some of the simplifying assumptions made so far in characterizing the measurements. In particular, the above two algorithms may be combined by modelling the dead reckoning as a biased random variable. We will also study the temperature effects which cause drift in the radio beacon system. This drift is very slow and occurs over hours; it did not affect the data collected here because the runs were so short in time.

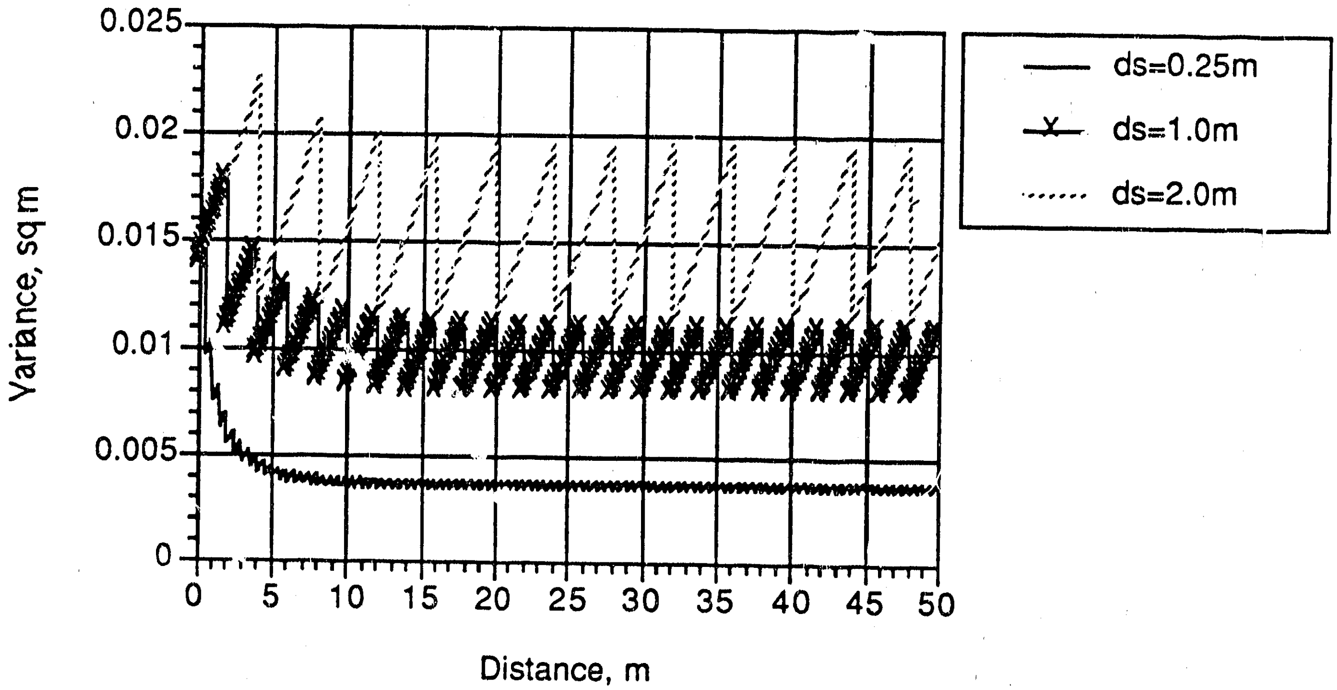

Figure 1. Theoretical Variance Propagation

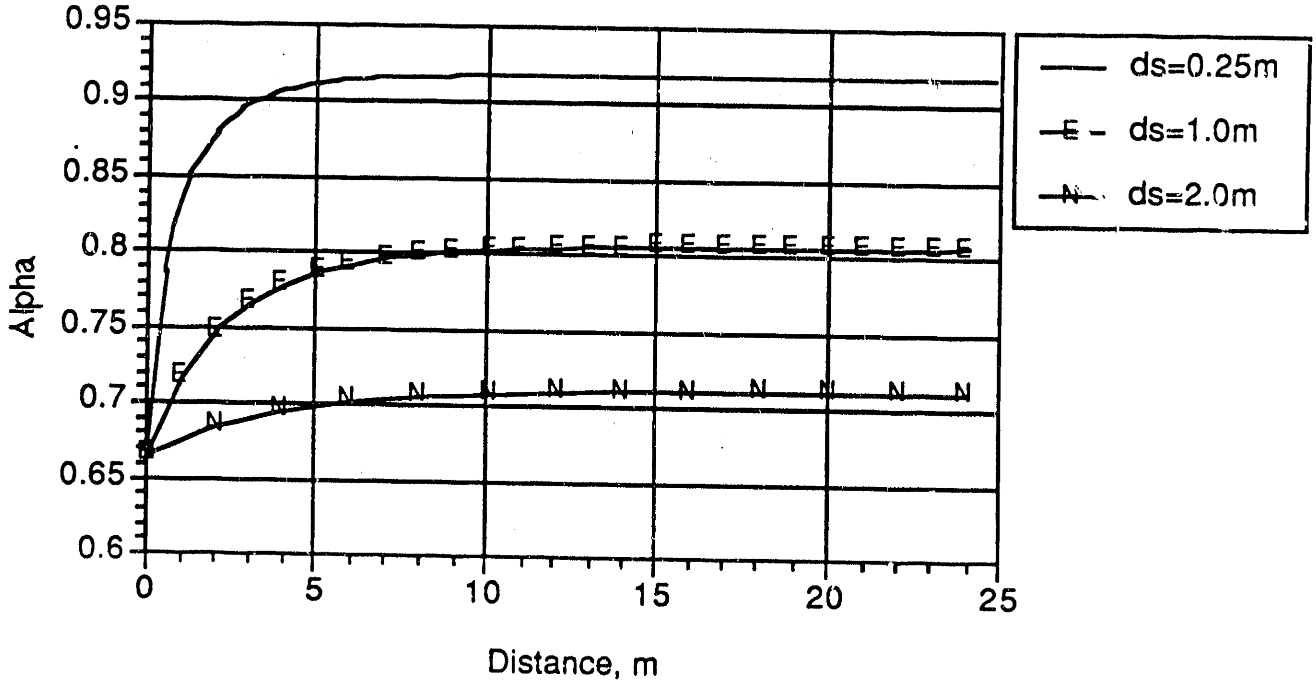

Figure 2. Theoretical Alpha Propagation; Q=5%

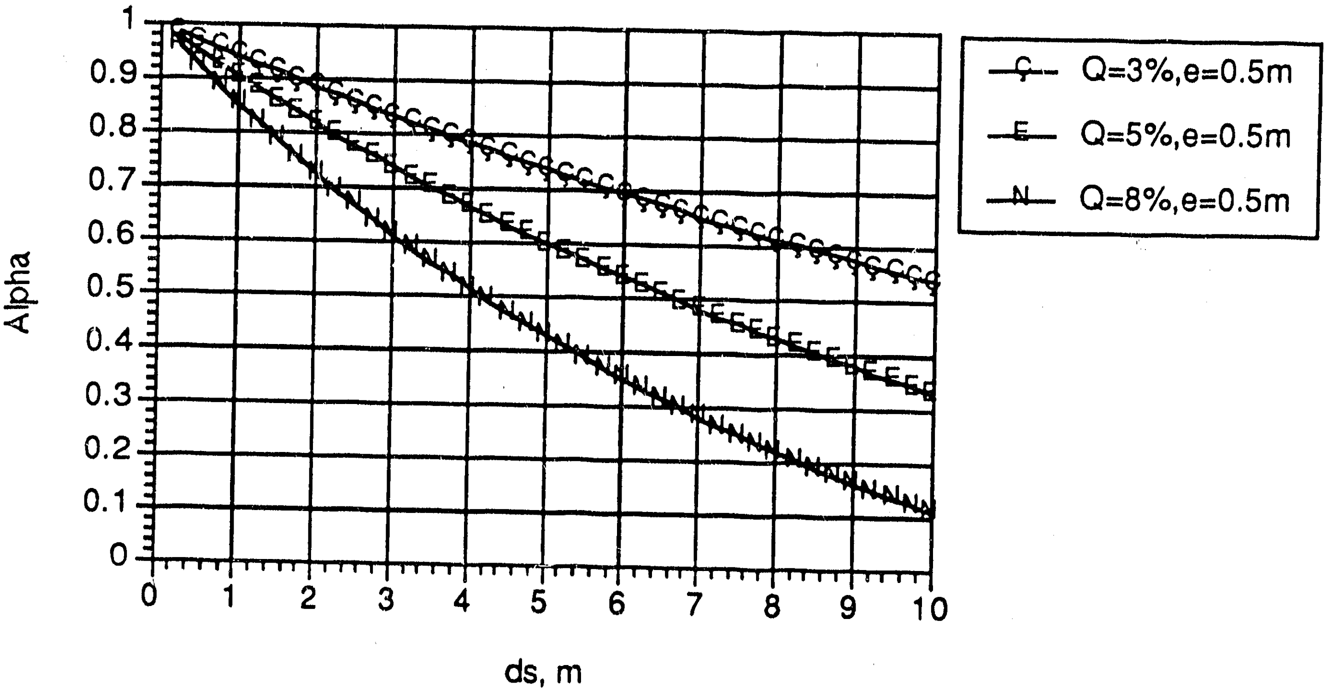

Figure 3. Steady-State Alpha vs Ds

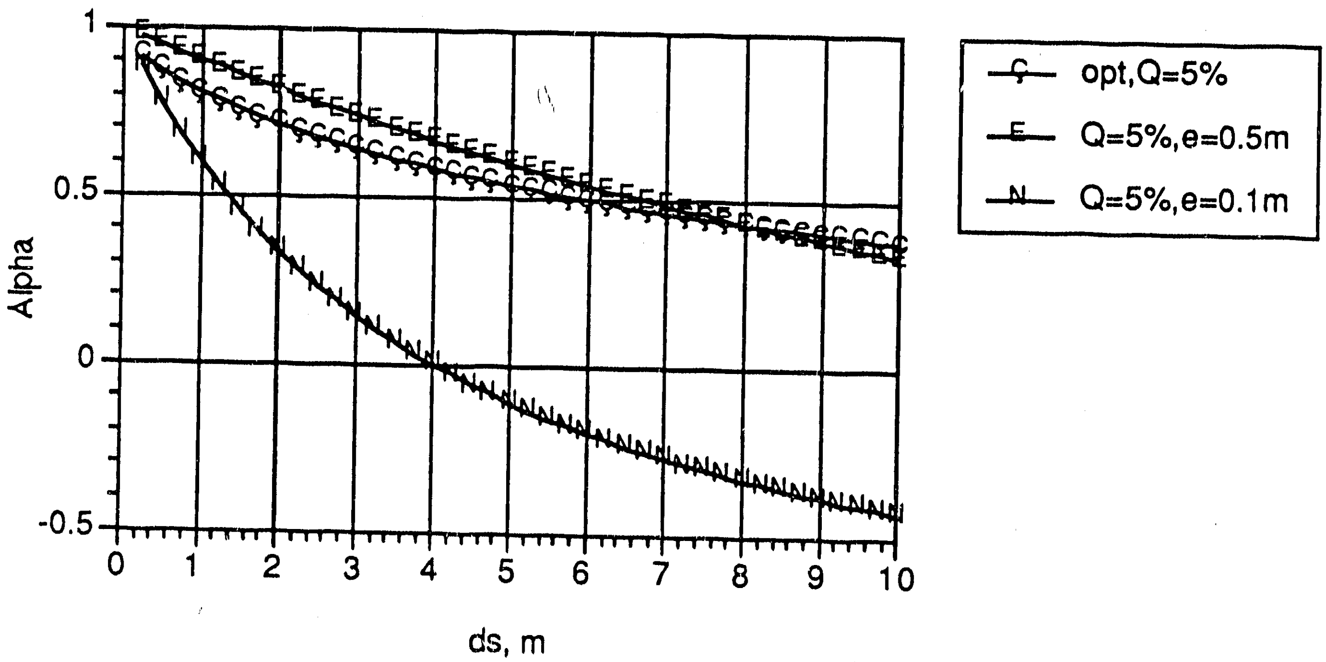

Figure 4. SS Alpha vs ds; Optimal vs Avg Err Comparison

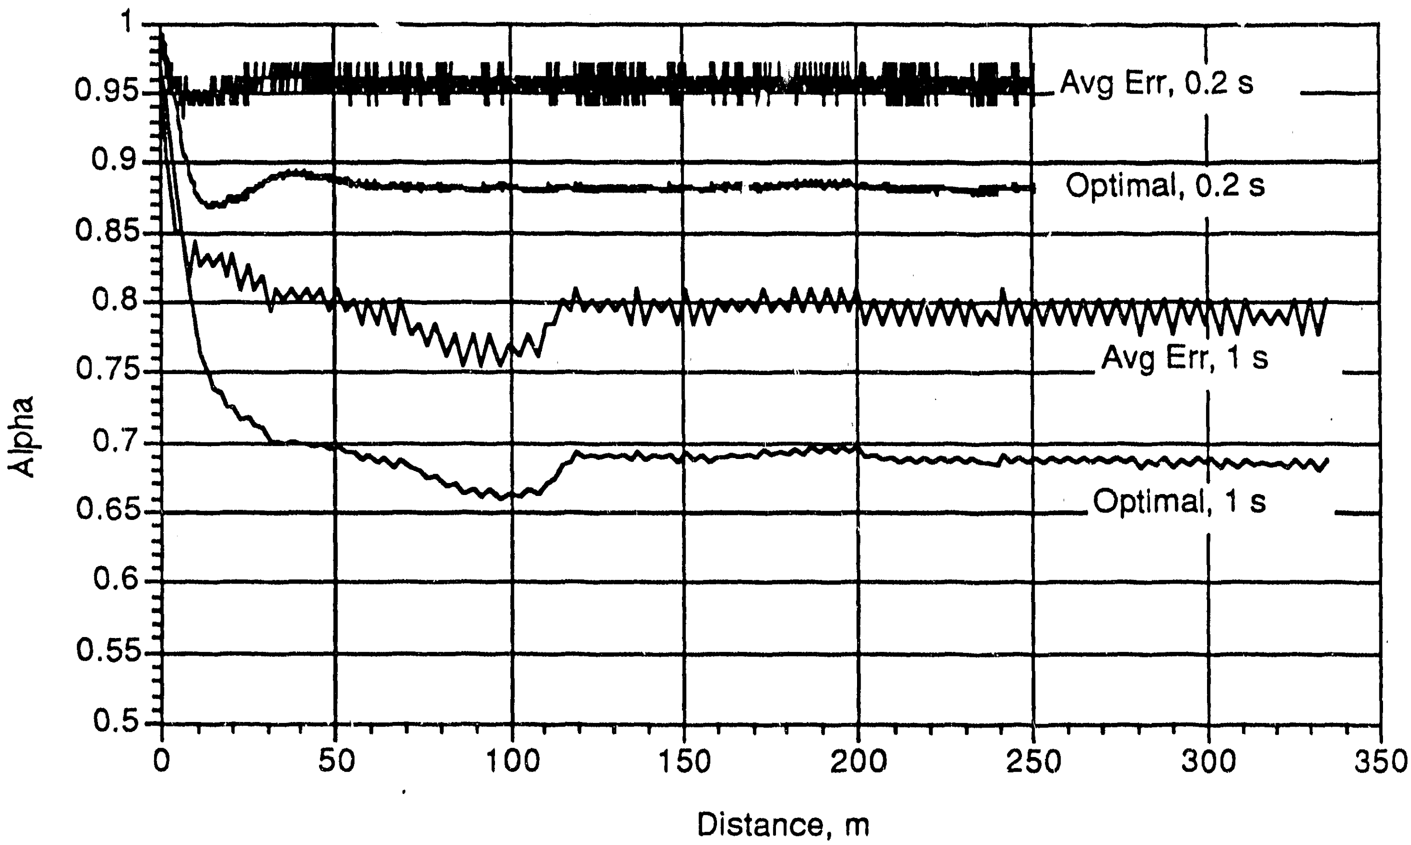

Figure 5. Alpha Propagation Measurements

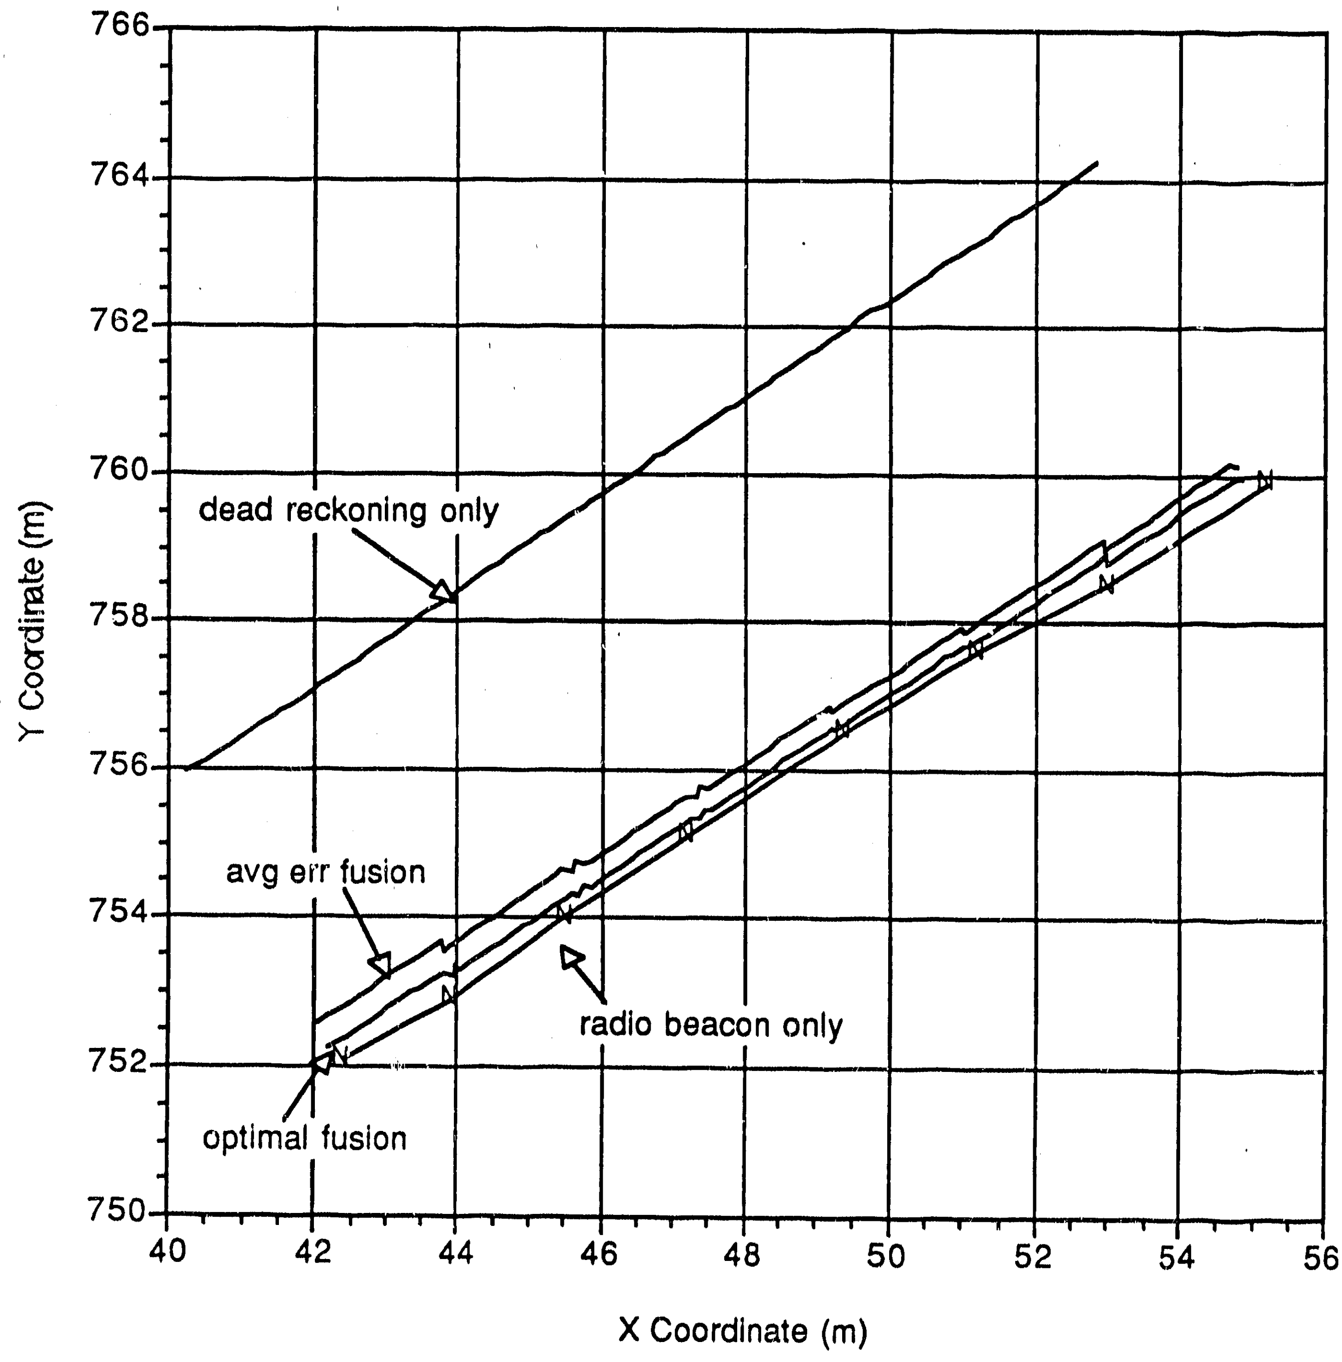

Figure 6. XY Trace Measurements, 1s Update

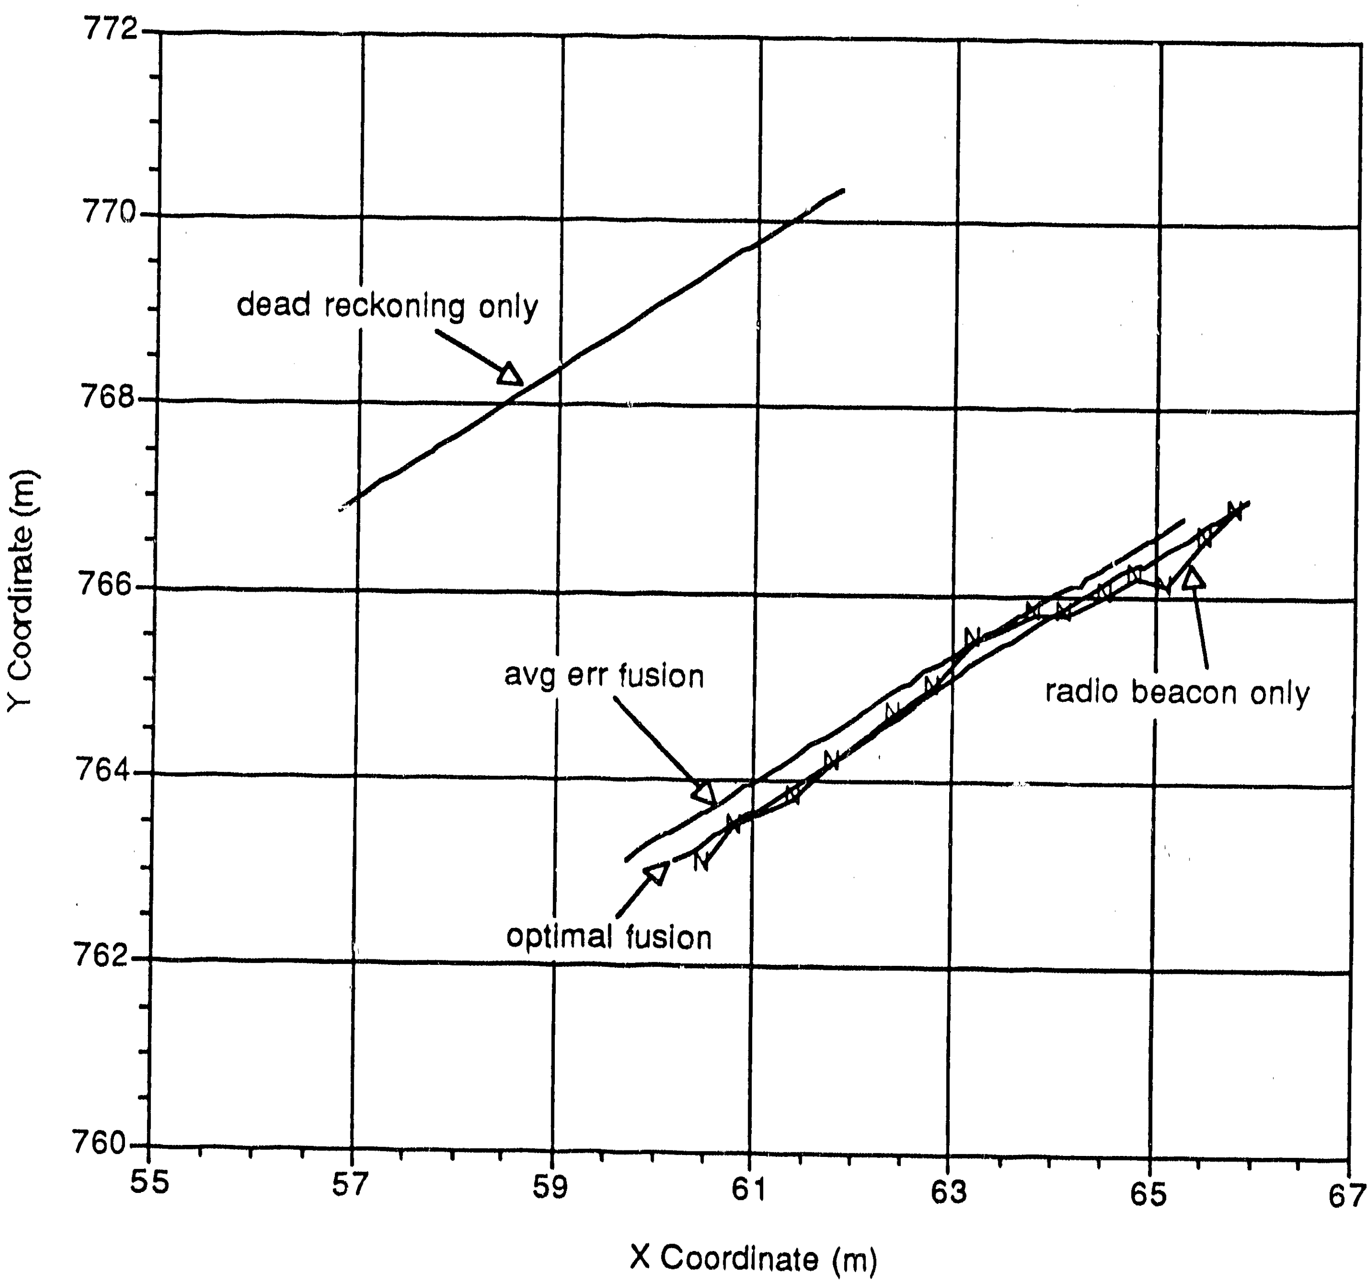

Figure 7. XY Trace Measurements, 0.2s Update

APPENDIX A

The derivation of the steady-state error for fusion of a constant bias, Q, dead reckoning measurement with an unbiased radio beacon measurement is done for motion along a straight line. The distance between fusion updates is Δs and is a constant.

There are two critical times to deal with: just before the fusion and just after the fusion. Dead reckoning propagates the measurement from just after the last fusion to just before the next fusion and the error grows linearly with the distance travelled between the two fusions.

Recall the fusion equation,

$$z = \alpha(DR) + (1-\alpha)(RB)$$

and note that the biasing error is only contained in the (DR) term and thus is propagated with the α factor.

If we let (e1+) indicate the initial error and (e2-) indicate the error just before the first fusion, then the following expression can be written:

$$(e2-) = (e1+) + Q\ \Delta s$$

where Δs is the distance travelled until the first fusion. After the first fusion the error is expressed as (e2+) and can be written as:

$$(e2+) = \alpha\ (e2-)$$

or

$$(e2+) = \alpha\ [\ (e1+) + Q\ \Delta s\ ]$$

Likewise, the error just before the next fusion is expressed as (e3-) and can be written as:

$$(e3-) = (e2+) + Q\ \Delta s$$

The error just after the next fusion is expressed as (e3+) and can be written as:

$$(e3+) = \alpha\ (e3-)$$

or

$$(e3+) = \alpha\ [\ (e2+) + Q\ \Delta s\ ]$$

or, substituting for (e2+),

$$(e3+) = \alpha\ [\ \alpha[(e1+) + Q\ \Delta s] + Q\ \Delta s\ ]$$

or, simplifying,

$$(e3+) = \alpha^2(e1+) + \alpha\ Q\Delta s + Q\Delta s$$

The general equation for the nth error term can be written by extrapolating the previous examples. The equation is given by:

$$(en+) = \alpha^{n-1}(e1+) + \alpha^{n-2}\,Q\Delta s + \alpha^{n-3}\,Q\Delta s + \ldots + Q\Delta s$$

Assuming that the position has been initialized exactly and thus the initial error, (e1+), is zero, this expression can be simplified to:

$$(en+) = Q\Delta s \sum_{k=2}^{k=n} \left[\alpha^{n-k} \right]$$

The steady-state after fusion error expression can be obtained by letting $n \to \infty$ and noting that the infinite series converges to $\alpha/(1-\alpha)$:

$$(e\infty+) = Q\ \Delta s\ \alpha/(1-\alpha) = e_{min}$$

The after fusion error is also the smallest error since it immediately begins to grow (linearly) with dead reckoning propagation. Since the maximum error occurs just before the fusion, the steady-state maximum error can be written as:

$$(e\infty-) = (e\infty+)/\alpha = Q\ \Delta s\ 1/(1-\alpha) = e_{max}$$

The average steady-state error is written as:

$$(e_{avg}) = (e_{min} + e_{max})/2$$

or

$$(e_{avg}) = Q\ \Delta s\ (1+\alpha)/(1-\alpha)$$

APPENDIX B

This appendix is a summary of the data collected to summarize the radio beacon measurement variance. Five different sets of data were measured at five distinct locations on the driving course at the Sandia National Laboratories' Robotics Vehicle Range in Albuquerque, NM.

The vehicle was parked at the location and 250 radio beacon readings were taken. Thus, there were 250 x,y coordinates measured at each location. The histograms of this data are listed in this appendix. The histogram "bin" width is not equal among the different histograms; bin width was computed using the maximum and minimum values of each histogram.

The y-values have the following variances (in m^2): 0.029, 0.031, 0.019, 0.011, 0.029. The average of these is 0.0238, which corresponds to a 6σ spread of about 0.93m.

The x-values have the following variances (in m^2): 0.019, 0.017, 0.020, 0.053, 0.014. The average of these is 0.0246, which corresponds to a 6σ spread of about 0.94m.

For most of the calculations, the variance used was slightly higher, 0.0278, which corresponded to an even 1.0m 6σ spread.

POINT A RADIO BEACON MEASUREMENTS

Y Distribution at Point A

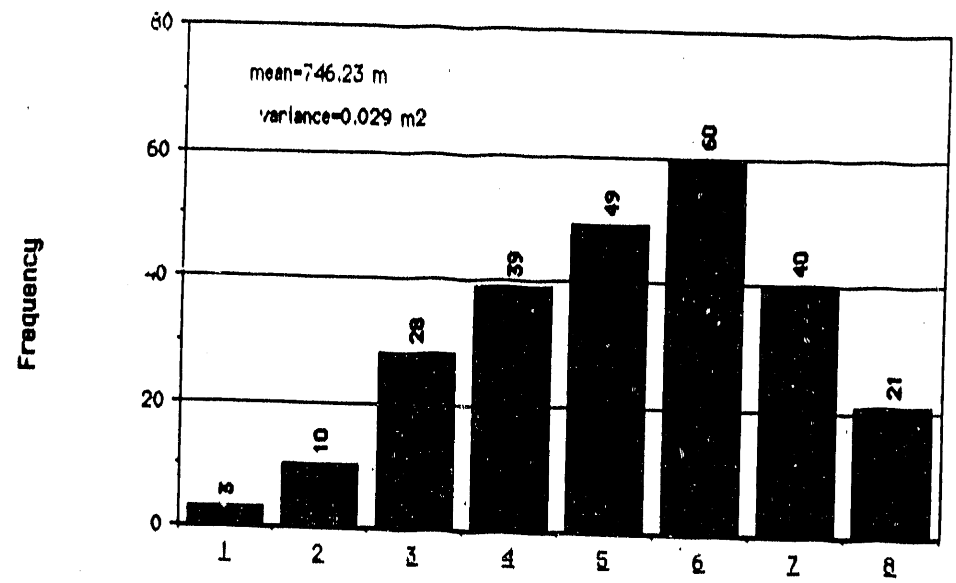

X Distribution at Point A

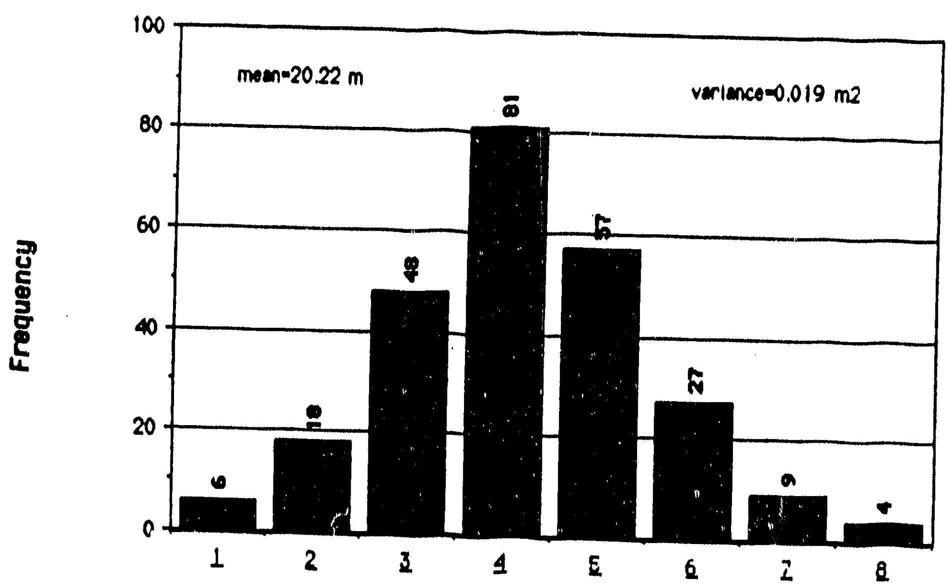

POINT B RADIO BEACON MEASUREMENTS

Y Distribution at Point B

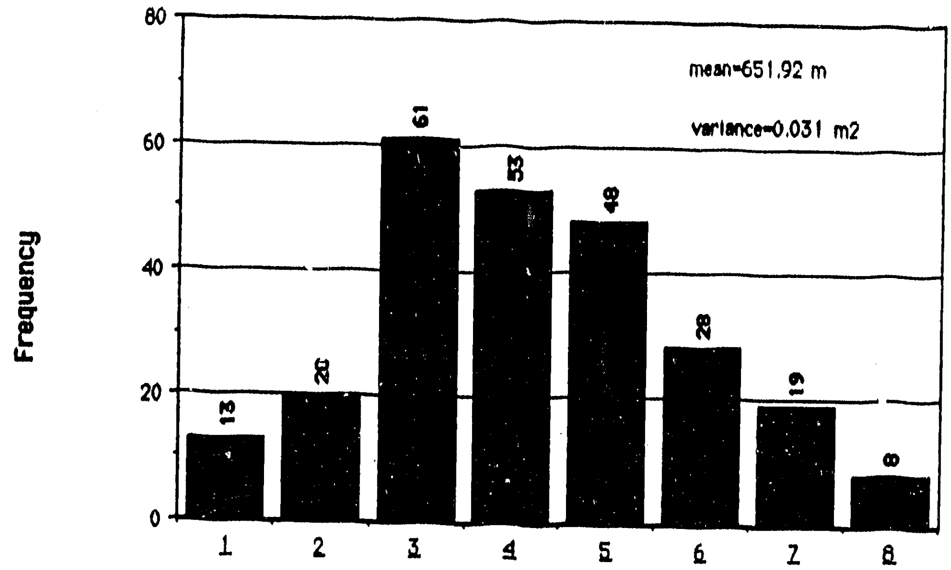

X Distribution at Point B

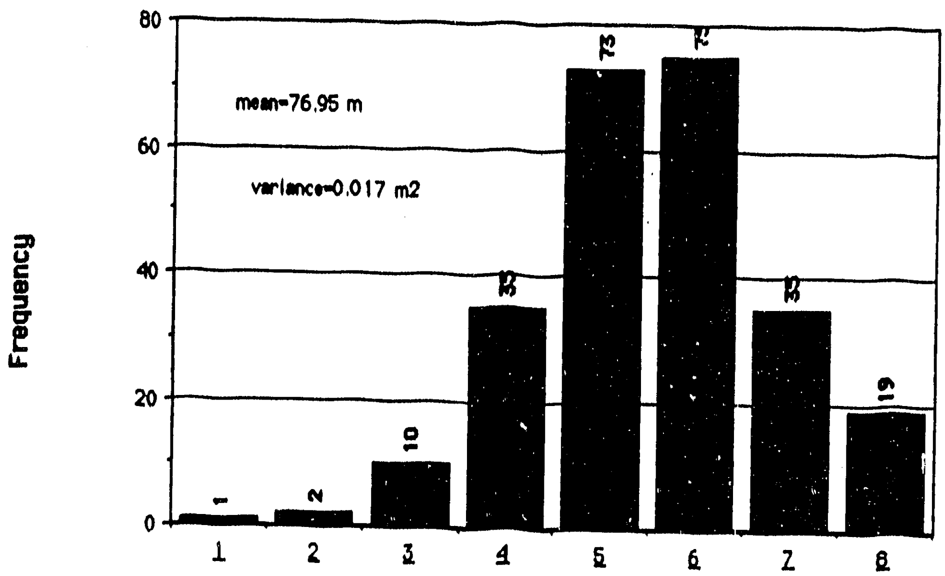

POINT C RADIO BEACON MEASUREMENTS

Y Distribution at Point C

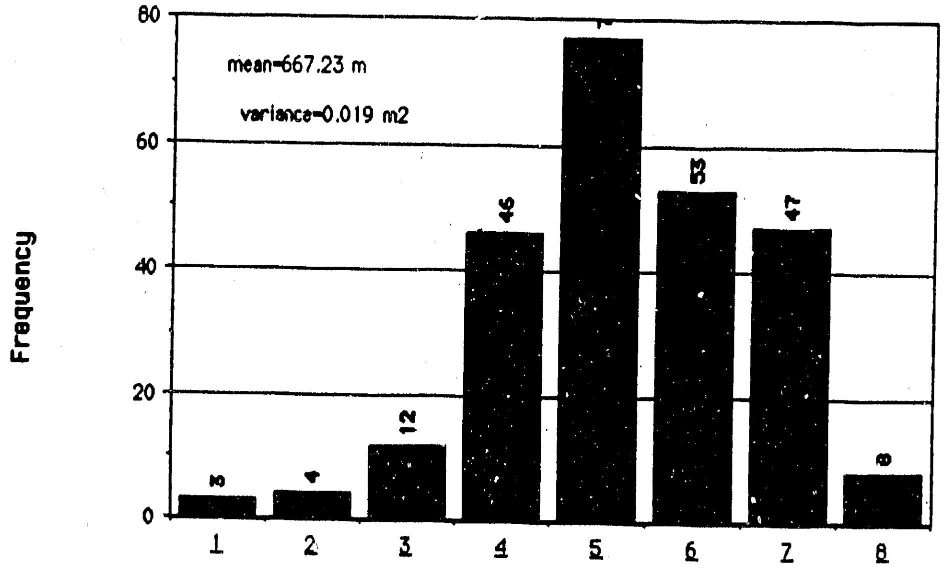

X Distribution at Point C

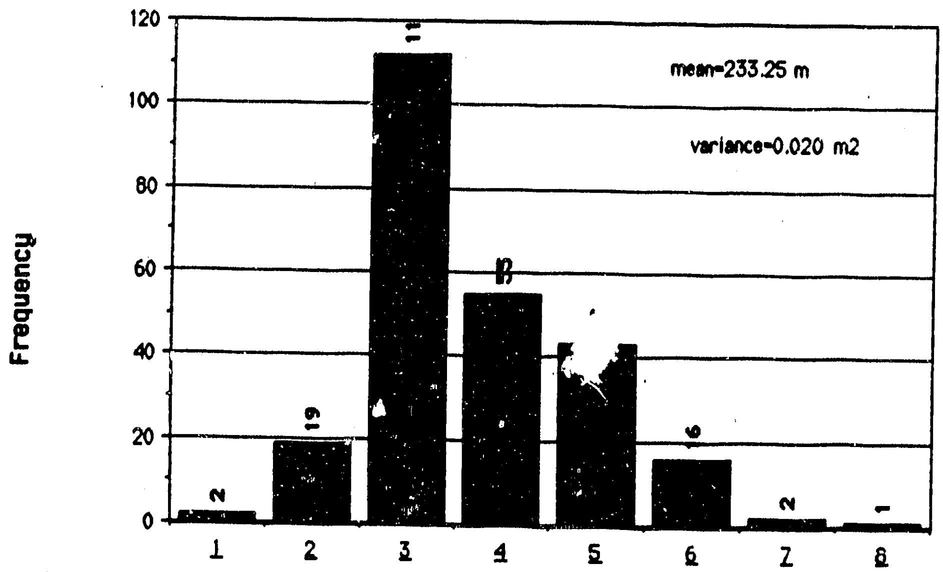

POINT D RADIO BEACON MEASUREMENTS

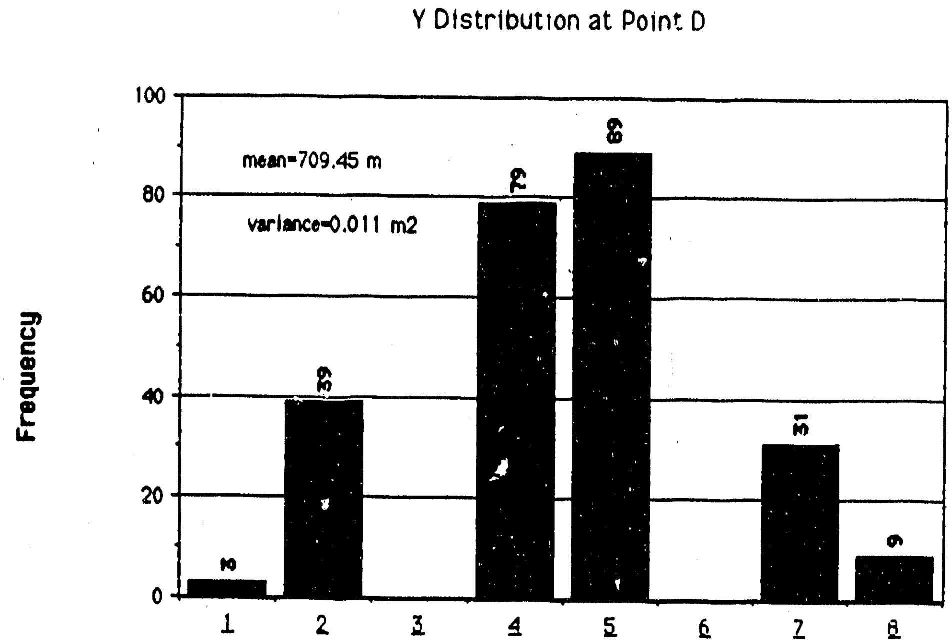

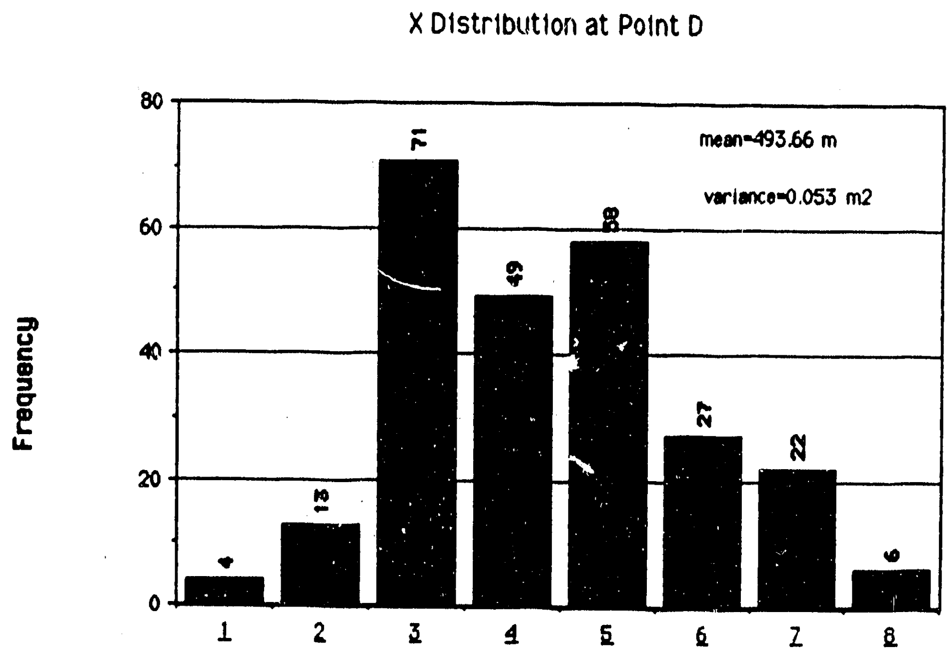

POINT E RADIO BEACON MEASUREMENTS

Y Distribution at Point E

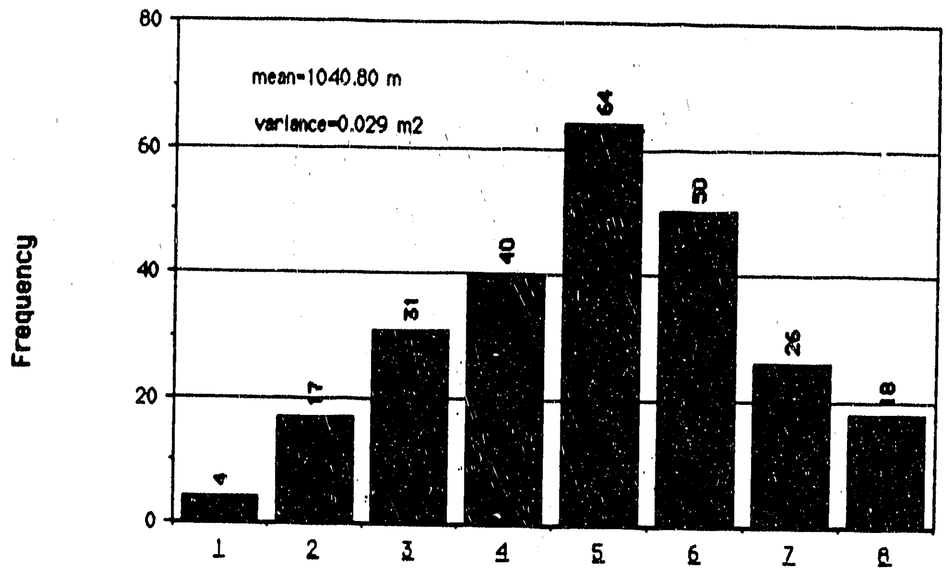

X Distribution at Point E

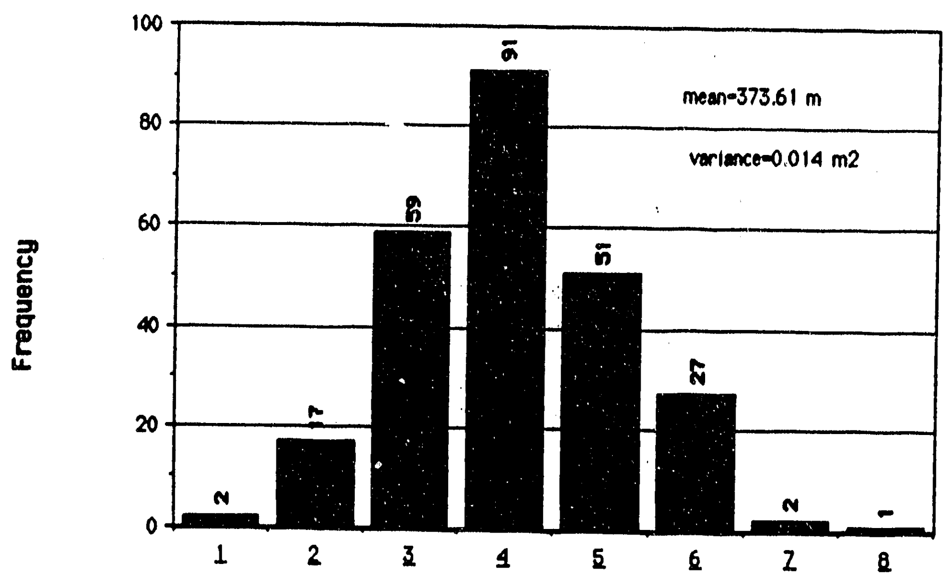

APPENDIX C

This appendix contains the derivation of the steady-state value of α as a function of the dead reckoning drift rate Q, and the distance travelled between radio beacon updates for the optimal fusion algorithm.

Recall that the value of α for the optimal algorithm is given by:

$$\alpha = \frac{P_{rb}}{P_{rb} + P_{dr}} .$$

where P_{rb} is the radio beacon variance and P_{dr} is the dead reckoning variance.
Also recall the equation for the variance of the fused estimate:

$$P = \alpha^2 * P_{dr} + (1-\alpha)^2 * P_{rb}$$

By combining the two above equations, the value of the fused estimate variance can be written as a function of the two measurement variances as follows:

$$P = \frac{P_{rb} \; P_{dr}}{P_{rb} + P_{dr}} .$$

Note that the fused estimate variance is smaller than both measurement variances, indicating that some additional unbiased information always improves the current estimate of position.

In these espressions, the radio beacon variance is a constant, but the dead reckoning variance increases as the position estimate is propagated with dead reckoning calculations, and then decreases abruptly when the position estimate is fused with radio beacon data. Thus, the value of α actually varies with dead reckoning variance, which in turn varies (by increasing) with distance travelled.

For a constant distance between updates, Δs, the steady-state condition occurs when the growth in variance during dead reckoning propagation (between radio beacon updates) equals the drop in variance when fusion with the radio beacon measurement is made. The growth in variance between radio beacon updates is related to the distance travelled between updates, Δs, and the dead reckoning drift rate as:

$$\Delta P = (Q\Delta s/6)^2$$

The drop in variance when the fusion with the radio beacon measurement is done is the variance before the fusion (P-) minus the variance after the fusion (P+). So the steady state condition occurs when:

$$\Delta P = (P-) - (P+)$$

where ΔP is the growth in variance during dead reckoning propagation (over Δs), (P-) is the estimate variance just before fusion with the radio beacon data, and (P+) is the estimate variance just after fusion with the radio beacon data.

(P+) and (P-) are related by the following expression:

$$(P+) = \frac{P_{rb}\ (P-)}{P_{rb} + (P-)} \ .$$

So the expression (P-) - (P+) can be written as:

$$(P-) - (P+) = \frac{(P-)^2}{P_{rb} + (P-)}$$

where (P-) is the steady-state value of the variance just before fusion. Also recall that the dead reckoning variance is equal to (P-) just before fusion and (P+) just after fusion.

Thus, the steady-state value of α is a function of (P-), or the steady-state dead reckoning variance just before fusion:

$$\alpha = \frac{P_{rb}}{P_{rb} + (P-)} \ .$$

Looking back at the steady-state condition,

$$\Delta P = (P-) - (P+)$$

(P-) can be related to Q, Δs, and P_{rb} by plugging in the equivalent expressions for ΔP and (P-) - (P+):

$$(Q\Delta s/6)^2 = \frac{(P-)^2}{P_{rb} + (P-)} \ .$$

Solving this quadratic for (P-) gives:

$$(P-) = (Q\Delta s/6)^2 \left[1 + \sqrt{1 + \frac{4\ P_{rb}}{(Q\Delta s/6)^2}} \right]$$

So the steady-state value of α is given by:

$$\alpha = \frac{P_{rb}}{P_{rb} + (P-)}$$

where (P-) is defined above.

REFERENCES

[1] Maybeck, Peter S., "Stochastic Models, Estimation, and Control, Volume 1", Academic Press Inc, 1979, pp 9-15.

[2] Banta, Larry E., "Self Tuning Navigation Algorithm", IEEE, 1988, pp 1313-1314.

END

DATE FILMED

12 / 13 / 90