NAWCWD TP 8832

Safety Considerations With Kalman Filters

by
Anthony S. Cantone
Kenneth R. Chirkis
Systems Safety Engineering Branch
Systems Engineering Department

AUGUST 2019

NAVAL AIR WARFARE CENTER WEAPONS DIVISION
CHINA LAKE, CA 93555-6100

Naval Air Warfare Center Weapons Division

FOREWORD

Navigation systems typically contain a Kalman filter, whose function is to develop a current optimal estimate of the state of a system based on the previous state, a system model, and measurements, while incorporating measurement noise statistics and instrumentation error characteristics. Modern navigation systems use extended Kalman filters, a real-time non-linear application of Kalman filters. However, for simplicity, a description of the processes typical of generic linear Kalman filters is presented along with hazards in the event of a Kalman filter malfunction. Hazards are considered from the viewpoint of an air vehicle application; although, this analysis can also be applied to arguably less hazardous applications such as an autonomous vehicle. A list of hazards to consider is presented for use by safety engineers.

This report was reviewed for technical accuracy by Glen McCue.

Approved by
H. D. KOOIMA, *Head*
Systems Engineering Department
9 August 2019

Under authority of
W. S. DILLON
RDML, U.S. Navy
Commander

Released for publication by
J. L. JOHNSON
Director for Research and Engineering

NAWCWD Technical Publication 8832

Published by .... Technical Communication Office
Collation .... Cover, 17 leaves
First printing .... 3 paper, 5 electronic media

REPORT DOCUMENTATION PAGE			*Form Approved* *OMB No. 0704-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
09-08-2019	Final	1 May 2018 – 31 January 2019

4. TITLE AND SUBTITLE Safety Considerations With Kalman Filters (U)	**5a. CONTRACT NUMBER** N/A **5b. GRANT NUMBER** N/A **5c. PROGRAM ELEMENT NUMBER** N/A
6. AUTHOR(S) Anthony S. Cantone and Kenneth R. Chirkis	**5d. PROJECT NUMBER** N/A **5e. TASK NUMBER** N/A **5f. WORK UNIT NUMBER** N/A
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Air Warfare Center Weapons Division 1 Administration Circle China Lake, California 93555-6100	**8. PERFORMING ORGANIZATION REPORT NUMBER** NAWCWD TP 8832
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) None.	**10. SPONSOR/MONITOR'S ACRONYM(S)** N/A **11. SPONSOR/MONITOR'S REPORT NUMBER(S)** N/A

12. DISTRIBUTION/AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

None.

14. ABSTRACT

(U) Navigation systems typically contain a Kalman filter, whose function is to develop a current optimal estimate of the state of a system based on the previous state, a system model, and measurements, while incorporating measurement noise statistics and instrumentation error characteristics. Modern navigation systems use extended Kalman filters, a real-time non-linear application of Kalman filters. However, for simplicity, a description of the processes typical of generic linear Kalman filters is presented along with hazards in the event of a Kalman filter malfunction. Hazards are considered from the viewpoint of an air vehicle application, although this analysis can also be applied to less hazardous applications such as an autonomous vehicle. A list of hazards to consider is presented for use by safety engineers.

15. SUBJECT TERMS

Estimation, Filter, Kalman, Propagation, Safety, Update

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Anthony Cantone
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAR	34	**19b. TELEPHONE NUMBER** *(include area code)* (760) 939-1820

CONTENTS

1.0 Introduction 5

2.0 What Is Filtering? 6

3.0 Design of a Typical Linear Kalman Filter 6

4.0 Kalman Filter Symbols 7

5.0 Kalman Filter Processes 8
- 5.1 State Transition Matrix 8
- 5.2 Propagation of the Error State Vector and the System Covariance Matrix 8
- 5.3 Update of the Error State Vector and the System Covariance Matrix 9

6.0 Kalman Filter Navigation Errors 10
- 6.1 Ill-Conditioning 10
- 6.2 Initialization of the System Covariance Matrix 11
- 6.3 Initialization of the Process Noise Covariance Matrix 11
- 6.4 Asymmetry of the System Covariance Matrix 11
- 6.5 Reasonableness Tests May Be Too Stringent or Too Lenient 12
- 6.6 Loss of Orthonormalization 13
- 6.7 Incomplete Model 14
- 6.8 Inadequate Alignment 14
- 6.9 Inappropriate Sampling Rate 14
- 6.10 Absence of In-Air Restart of the Kalman Filter 15

7.0 Aberrant Behavior 15

8.0 What Are the Hazards? 16

9.0 What Are Some of the Mitigations? 17

10.0 Behavior – Causes Relationship 17

11.0 Some Examples of Aberrant Behavior Related to Filter Defects 18
- 11.1 Experiment 1 – Stationary Vehicle, Measurements With Gaussian Noise 19
  - 11.1.1 Analysis 21
- 11.2 Experiment 2 – Erroneous Model With High Confidence 21
  - 11.2.1 Analysis 22

11.3 Experiment 3 – Erroneous Mode With Some Lack of Confidence 23
11.3.1 Analysis 23
11.4 Experiment 4 – Erroneous Model With Serious Lack of Confidence 24
11.4.1 Analysis 24
11.5 Experiment 5 – Erroneous Model With No Confidence in Model 25
11.5.1 Analysis 25
11.6 Experiment 6 – Measurements With Non-Gaussian Noise 26
11.6.1 Analysis 27
11.7 Experiment 7 – Rejection of Measurements 29
11.7.1 Analysis 29
11.8 Experiment 8 – Rejection of Noisy Measurements Due to Low Data Confidence Level 30
11.8.1 Analysis 31

12.0 Summary 32

References 33

Acronyms 34

Figures:

11-1. Normal Probability Plot, Experiment 1 20
11-2. Result of Experiment 1 21
11-3. Result of Experiment 2 23
11-4. Result of Experiment 3 24
11-5. Result of Experiment 4 25
11-6. Result of Experiment 5 26
11-7. Normal Probability Plot, Experiment 6 27
11-8. Result of Experiment 6 28
11-9. Result of Experiment 7 30
11-10. Result of Experiment 8 32

Tables:

6-1. $\chi 2$ Table (Reference 5), 3σ, 2 DOF 13
10-1. Observed Aberrant Behavior and Their Possible Causes 18
11-1. Measurements File for Experiment 1 19
11-2. Measurements File for Experiment 2 22
11-3. Measurements File for Experiment 6 26
11-4. $\chi 2$ Table, 2σ, 1 DOF 29
11-5. Measurements File for Experiment 8 31
11-6. $\chi 2$ Table, 1σ With 1 DOF 31

ACKNOWLEDGMENTS

The authors would like to thank Glen McCue for his technical assistance and review, and Erica Arroyo for her commentary. The authors would also like to thank Amy Wyatt for her excellent editorial contributions.

This page intentionally left blank.

1.0 INTRODUCTION

A typical hazard associated with a failed Kalman filter is the inability to furnish navigation aids. For an air vehicle, this could result in the inconvenience of having to navigate visually, to navigate using noisy sensors, or losing the ability to perform critical maneuvers such as landing or navigating through heavy air traffic.

The purpose of this report is to give the safety engineer a means to understand the structure of a typical linear Kalman filter. Even though filters found in current navigation systems are typically “extended,” which means “non-linear,” studying a linear filter will give the engineer a basic understanding of the Kalman filter so that s/he can associate aberrant behavior in the operation of an air vehicle navigation system with defects in the Kalman filter.

After a short chapter “What is Filtering?” in which filtering is defined, this report discusses Kalman filter design in the chapter “Design of a Typical Linear Kalman Filter.” A chapter with the title “Kalman Filter Symbols” presents the notation used to represent the parts of the filter, along with a description. “Kalman Filter Processes” describes the two processes that Kalman filters typically exhibit. There are several possible Kalman filter errors that are described in the chapter “Kalman Filter Navigation Errors.” This is the heart of the report and is designed to give the safety engineer the knowledge s/he needs to judge what the hazards associated with a particular navigator design might be. The chapter “Aberrant Behavior” describes some of the behaviors that a defective Kalman filter can exhibit. The “What Are the Hazards?” chapter describes the observable hazardous behavior that is possible if the design is faulty in the way described in the “Kalman Filter Navigation Errors” chapter. A mitigation is a feature that helps to alleviate the bad effects of a hazard. The chapter “What Are Some of the Mitigations?” is included to indicate possible mitigations. A list of behaviors and their associated causes appears in the “Behavior – Causes Relationship” chapter. Finally, examples of a single state filter with defects discussed in this report are displayed in the chapter “Some Examples of Aberrant Behavior Related to Filter Defects.”

Error state variables rather than state variables are employed in the formulation of the Kalman filter state vector.

2.0 WHAT IS FILTERING?

Filtering, as used in this report, refers to a process of estimating the state of a system at the current time based on past history and measurements (Reference 1, pp. 2-3). Since measurements are usually noisy, a means for estimating the effect on the state of the system due to the measurement noise is appropriate. A navigation system typically has as sources of errors process noise (due to inaccuracies in the accelerometers and gyroscopes, round off and chopping errors), instrument misalignment (due to assembly errors), measurement noise and biases, and other error sources. A Kalman filter is a mechanism for determining how much the measurements should influence the estimate of the state of a system, given the previous estimate, a model, the covariance of the state errors and measurement noise statistics.

Filtering should not be confused with smoothing or predicting. While filtering estimates the system state at the time of the last measurements, smoothing estimates the system state given measurements both in the past and in the future time of the estimate, and predicting uses past history to estimate the system state past the end of the measurements available.

Filtering has another meaning that of extracting a signal from noise. In Signal Processing, a filter is a process that reduces the amplitude of unwanted frequencies from a signal. This should not be confused with the word filtering as used in this report, which is a process of estimating.

Every Kalman filter is designed to be recursive, which means that all the past history is incorporated in the latest state and covariance estimates. Therefore, in a recursive filter, it is unnecessary to store the past states or measurements.

3.0 DESIGN OF A TYPICAL LINEAR KALMAN FILTER

A Kalman filter that is linear in both the state propagation and measurement update can be described typically as two processes working in parallel: a propagation or extrapolation process and an update process. For Kalman filters found in navigation systems, the propagation process is based on a model that predicts navigation quantities such as position and velocity at the current time by using previous position, velocity, gyroscope, and accelerometer information. The update process corrects this prediction using measurement data. Instrumentation and measurement errors are estimated and considered in the propagation and update process.

The filter assumes Gaussian noise in the accelerometer, gyroscope, and measurement sensors.

4.0 KALMAN FILTER SYMBOLS

The following list contains the symbols used to represent the parts of the filter.

X Error State Vector: Contains the difference between propagated state and the updated state due to measurements. Initially set to zero.

F System Dynamics Matrix: Dynamics of the system.

Φ State Transition Matrix: $\Phi = e^{F\Delta t}$.

H Measurements Matrix: (aka Measurement Sensitivity Matrix) Transformation between state space and measurement space.

R Measurements Noise Matrix: Sensor noise characteristics. If measurements are independent, then off-diagonal terms are 0.

Z Observations Vector: Vector of measurements.

r Measurements Residual Vector: $r = Z - HX^{-}$ (X^{-} is the *a priori* state vector).

P System Covariance Matrix: Contains variances of the errors of each of the states on the main diagonal and covariance between the states at the off-diagonal positions.

K Kalman Gain Matrix: Answers the question of how much of the measurements to incorporate into the propagated state estimate.

Q Process Noise Covariance Matrix: Contains accelerometer, gyroscope errors, round-off, chopping, computational errors in the states. Often, this matrix can be used to *tune* the filter.

C_E^L Direction Cosine Matrix (DCM) from earth-centered inertial (ECI) frame of reference to local level frame of reference.

C_B^L DCM from body frame of reference to local level frame of reference.

χ^2 Statistic associated with the measurement residual.

5.0 KALMAN FILTER PROCESSES

This section defines propagation and update of both the error state vector and the system covariance matrix. This section also demonstrates calculation of the Kalman gain matrix during the update process.

5.1 STATE TRANSITION MATRIX

The state transition matrix $\Phi\ (t_2, t_1)$ transforms or propagates a state vector solution $X(t_1)$ at time t_1 to the state vector solution $X(t_2)$ at time t_2. The state transition matrix is derived from the model as it appears in the system dynamics matrix F.

5.2 PROPAGATION OF THE ERROR STATE VECTOR AND THE SYSTEM COVARIANCE MATRIX

In a Kalman filter, the values of the error state vector X and system covariance matrix P are extrapolated, or propagated, in time by use of the state transition matrix Φ. If there is no measurement update, the state vector X and the system covariance matrix P are still propagated in time.

The equations of propagation are as follows:

Propagation of error state vector: $X_n^- = \Phi_{n-1}^n X_{n-1}^+$

Propagation of system covariance matrix: $P_n^- = \Phi_{n-1}^n P_{n-1}^+ (\Phi_{n-1}^n)^T + Q_n$ (T= transpose)

In the above propagation of error equation, X_n^- is the error state vector after the last propagation cycle and before the next update, called the *a priori* state vector; X_{n-1}^+ is the state vector before the next propagation and after the last update, called the a posteriori state vector; the subscript n-1 corresponds to time t_{n-1} before the state update; the subscript n corresponds to time t_n after the state update (Reference 2, p. 143, Table 8.1).

The system covariance matrix P_n^- reflects the opinion of the filter about the accuracy of the state variables. The initialization of this matrix is the estimated variances along the main diagonal. If an initialization error of this matrix is made, the values in this matrix should adjust quickly to more accurate values as each measurement is made. If the initial values of some state vector elements are unknown, the variances of the unknown state vector elements should be set to large values. The superscripts and subscripts carry the same implication as with the propagation of error equation.

The process noise covariance matrix Q_n reflects the confidence of the designer that the model is accurate (the transition equations are correct and there are no missing states). Smaller values along the diagonal in this matrix reflect a higher confidence. (Reference 3, p. 574).

5.3 UPDATE OF THE ERROR STATE VECTOR AND THE SYSTEM COVARIANCE MATRIX

The update process begins with the collection of a measurement. This measurement is collected in measurement space, which is a coordinate system that may be different from state space, the coordinate system in which the state vector is defined. For example, a radar system locates a vehicle in a polar coordinate frame of reference using range and direction. The estimated state vector (in a Cartesian coordinate system) will be transformed to polar coordinates using an appropriate H matrix, so that the measurements may be compared to the estimated states.

A comparison between the measurements and the estimate of the current state is performed. The purpose is an estimate of the reasonableness of the measurements. This guards against incorporation of spurious measurements due to instrument failure but may also neglect to incorporate measurements associated with high dynamics maneuvers. See the chapter “Reasonableness Tests may be too Stringent or too Lenient” for details on how this is done. If the decision is that the measurements are not reasonable, they are ignored, and the update cycle is skipped. If the measurements are reasonable, the Kalman gain is then calculated.

The Kalman gain is a matrix that computes the weight which is applied to the measurements. Equation 5-1 shows computation of the Kalman gain matrix:

$$K_n = P_n^- H_n^T (H_n \, P_n^- H_n^T + R)^{-1} \tag{5-1}$$

Once the Kalman gain matrix has been computed, the decision having been made that the measurements are to be incorporated, the system covariance matrix and error state vector are updated. That is, the states and covariance are updated using information provided by the measurements. The measurement incorporation is weighted by the Kalman gain matrix. If only low weight is given to the measurements, this reflects the filter’s belief that the measurements are not very reliable, and the filter tends to trust its propagation more than the measurements. High weight given to the measurements reflects the filter trusting the measurements as much as or more than the propagation.

In the state and covariance measurement update, the superscript “–” is used before a measurement update takes place, and “+” is used after the measurement update.

The update of the error state vector is shown in Equation 5-2 (Reference 1, p. 110):

$$X_n^+ = X_n^- + K_n[Z_n - H_n X_n^-] \tag{5-2}$$

where X_n^- is the error state vector after the last propagation cycle, H_n transforms the state space to measurement space, and the quantity in the brackets is the difference between the last propagation and the corresponding measurements, called the measurement residual vector.

In an analogous way, the update of the system covariance matrix is performed as in Equation 5-3, where I is a unit matrix (Reference 1, p. 110):

$$P_n^+ = (I - K_n H_n)P_n^- \tag{5-3}$$

If the elements of the Kalman gain matrix are effectively 0, the system covariance will not be changed. The larger the elements of the Kalman gain matrix, the more change will be incorporated into the covariance.

6.0 KALMAN FILTER NAVIGATION ERRORS

6.1 ILL-CONDITIONING

One of the symptoms of an ill-conditioned problem is the behavior of the solution with respect to variations of the input data. In an ill-conditioned problem, wild and chaotic behavior of the solution is possible with small changes in the input data. Such changes can occur from random variations, such as noisy data. Ill-conditioning is not only a result of noisy data but also of faulty algorithms. For example, there are ways that the computation of the Kalman gain can be affected by modeling errors in the Φ, Q, H, and R matrix parameters, large ranges of values of these matrix parameters, ill-conditioning of intermediate results, computer round-off errors, and large number of states (Reference 4, p. 195).

Ill-conditioning can play a part in determining observability of a dynamic system. One of the characteristics of an unobservable system is the inability to invert the observability matrix because it is singular. However, arbitrarily small changes in the elements of a singular matrix can make it nonsingular. These small changes can occur because of round-off or truncation errors. Then, what one has is a dynamic system that is "almost observable." Various schemes, such as using the reciprocal of the condition number of the Gramian matrix as a quantitative measure of un-observability, can be employed (Reference 4, p. 43).

6.2 INITIALIZATION OF THE SYSTEM COVARIANCE MATRIX

The effect of initialization errors is not a crucial factor in the performance of the filter. The initial value of the variances along the main diagonal of the System Covariance Matrix (*P*) reflects the knowledge (or lack thereof) of the initial values in the state vector. For example, if wind buffeting of a vehicle is estimated to cause a variation of 1 degree (1σ) in the tilt errors of the instrument platform, the variance of the tilt error state should be set to 1. After a few filter cycles, the system covariance matrix should converge to more reasonable values.

6.3 INITIALIZATION OF THE PROCESS NOISE COVARIANCE MATRIX

The initialization of the process noise covariance (*Q*) matrix is crucial for a satisfactory behavior of the filter. If the accelerometers and gyroscopes are used in a self-alignment process, statistics from the data sheets of these instruments can be included on the main diagonal. Otherwise, the values in the Q matrix reflect the noise associated with round-off or truncation errors that appear in the state estimates.

Adjusting these values may help the filter behave in a way to facilitate the estimation process. In other words, the Q matrix can be used to *tune* the filter so that its performance shows the appropriate confidence in the model. If too much confidence (small Q) is shown in the model, the resulting Kalman gain will allow too little a percentage of the measurements and too much of the propagated states to influence the next state vector values. Too little confidence (large Q) will on the other hand weigh on the side of incorporating too low a percentage of the propagated states and too high a percentage of the noisy measurements in the next state vector values. This occurs because the values in the Q matrix add directly to the system covariance matrix.

However, be aware that although a low value in the Q matrix can result in the filter diverging from reality by ignoring the measurements in favor of a possibly faulty model, larger values in the Q matrix can result in the filter overly relying on the measurements and ignoring the model propagation.

6.4 ASYMMETRY OF THE SYSTEM COVARIANCE MATRIX

If the Kalman filter is implemented using the entire system covariance matrix rather than an upper triangular mechanization, computer round-off and/or truncation error can cause asymmetry in the system covariance matrix. Over a long period of time, this asymmetry can result in zero or negative variances along the main diagonal of the system covariance matrix. Since the variances are theoretically positive definite, this will eventually cause significant errors in the calculation of the states. In order to forestall this possibility, it is recommended that the system covariance matrix be symmetrized after each propagation. The simplest method for making a matrix symmetric is accomplished by copying the upper triangular to the lower triangular portion of the matrix. Another possibility is to average each term with its image across the main diagonal.

6.5 REASONABLENESS TESTS MAY BE TOO STRINGENT OR TOO LENIENT

The χ^2statistic associated with the measurement residual *r*, shown in Equation 6-1:

$$r = Z - HX^-, \tag{6-1}$$

where X^- is the *a priori* state vector, can be used to protect the filter from spurious measurements. It can be designed to protect against large deviations from the expected measurements, for example, on the order of five or more times standard deviation. This will ensure that spurious measurements, due to instrument failure or noise spikes, which are most likely to appear as extremely large deviations from the expected measurements, will be rejected, while fairly unlikely measurements will still be accepted. The risk is that measurements during times of large maneuvers may be rejected.

The measurement residual (Equation 6-1) is assumed to be normally distributed with mean zero. If *r* is the measurement residual vector with *d* degrees of freedom (number of measurements), and *S* is its covariance matrix in Equation 6-2,

$$S = HP^-H^T + R \tag{6-2}$$

then the quadratic form in Equation 6-3

$$\chi^2 = r^T S^{-1} r \tag{6-3}$$

is known to be χ^2 distributed with *d* degrees of freedom (DOF) (Reference 5). This scalar can be compared to a critical chi-squared value that corresponds to a preselected significance level for the number of DOF, i.e., the dimension of the measurement residual vector.

The usual approach is to choose a significance level and form the null hypothesis "The current measurement is valid."

For example, assume a filter with two independent measurements and that all measurements generating a residual farther than 3σ from the expected measurements are to be rejected. Since 99.73% of the area under a Gaussian distribution lies in the 3σ range, 99.73% is the confidence level; the significance level 0.27% is the percent of measurements rejected.

From a chi-squared table, it is seen that this corresponds to a χ^2 statistic of approximately 12 for 2 DOF, so the χ^2 statistic is established at 12. If the computed χ^2 statistic exceeds the established limit, the null hypothesis is rejected, and no update is performed. In Table 6-1, the probability level of 0.0027 has an associated χ^2 statistic somewhere around 12 for 2 DOF (see in Table 6-1).

TABLE 6-1. χ^2 Table (Reference 5), 3σ, 2 DOF.

DOF	Significance Level					
	0.5	0.10	0.05	0.02	0.01	0.001
1	0.455	2.706	3.841	5.412	6.635	10.827
2	1.386	4.605	5.991	7.824	9.210	13.815
3	2.366	6.251	7.815	9.837	11.345	16.268

If the significance level is too small, the rejection rate will be too low. The result is that the Kalman filter will tend to accept all measurements, and the filter will be exposed to spurious measurements if a sensor fails. If the significance level is too large, the rejection rate will be too high. The result is that the filter will tend to reject all measurements as though it has become smart and has more confidence in its extrapolations than in the measurements, and it will continue uncorrected.

6.6 LOSS OF ORTHONORMALIZATION

Coordinate transformations use DCMs. The accelerometers and gyroscopes give outputs relative to an inertial frame of reference. In order for the dynamics equations to be solved, a transformation from the inertial frame to the earth-centered earth-fixed (ECEF) frame is needed, followed by a transformation from the ECEF to the local level frame in a tangent plane at the location of the aircraft. Finally, a transformation from the local level to the body frame of the aircraft is needed. The inertial measurement unit (IMU) is fixed to the body frame, but there may be alignment errors, called tilts, which can be modeled as error quantities in the error state vector.

The body frame to local level frame DCM should be an orthonormal matrix. This DCM can be determined from the vehicle Euler angles by considering the yaw, roll, and pitch angles of the vehicle. Round-off errors contribute to a deviation from the orthonormal condition during propagation and update. Therefore, it would be prudent to install a method to preserve the orthonormalization of this matrix. A normal matrix is a square matrix that can be diagonalized with real eigenvalues on the diagonal. Orthogonal matrices have the useful property that their inverses are identical to their transposes. This makes it possible to calculate the inverse of an orthogonal matrix without risking the numerical chaos that ill-conditioning would bring.

Suppose C is a 3x3 DCM. A formal analysis (Reference 6) shows that post-multiplication by $(C^T C)^{-\frac{1}{2}}$ will orthogonalize the matrix C. A useful modification can be developed as follows.

Assuming that C is almost orthonormal, $C^T C$ can be written as $C^T C = I + \Delta C$ where I is the 3x3 identity matrix, and ΔC is a 3x3 matrix whose terms are small when compared to C. Then, if M is the modified (i.e., orthonormal) C matrix,

$$M = C(I + \Delta C)^{-\frac{1}{2}} \tag{6-4}$$

which becomes, to a first order approximation,

$$M = \frac{1}{2} C(3I - C^T C) \tag{6-5}$$

The body to local level DCM should be orthogonalized periodically. The period is a tunable parameter. Note that orthogonal matrices are normal.

6.7 INCOMPLETE MODEL

Unknown and unmodeled forces will cause the behavior of the system to diverge from the modeled behavior. This will eventually result in the rejection of the measurements because the measurement residual r will become very large, and the solution will not be corrected. An example is the failure to model drag on an earth satellite due to ballooning of the atmosphere because of uneven heating of the earth's surface.

6.8 INADEQUATE ALIGNMENT

The reason for the alignment process is to determine where level is and the direction of north is. Level and north are stored in a DCM that determines the relationship between the IMU and the body axes of the vehicle. The longer the alignment process lasts, the more accurate the level and north will be. Navigation with an inaccurate north direction accumulates error the farther one travels in the north-south direction.

6.9 INAPPROPRIATE SAMPLING RATE

Accelerometer and gyroscope integration are used during alignment and in some models to form the transition matrix. Integration of accelerometer and gyro outputs generally is noisy, resulting in a potential for poor propagation. The problem is worsened when the readings are either held constant or averaged in the interval between samples.

It is reasonable that the integration interval of the accelerometer and gyro outputs should be small to facilitate accurate integration. Trapezoidal integration or the use of Simpson’s Rule for integration of the accelerometer and gyro outputs also helps towards accurate integration when data throughput limits the size of the integration interval. Otherwise, integration errors may result in erroneous propagations.

6.10 ABSENCE OF IN-AIR RESTART OF THE KALMAN FILTER

If the Kalman filter ceases working, a possible mitigation is to restart the filter while the vehicle is in the air. The process is actually an in-air realignment of the navigation system. In-air alignment has the advantage (over on-ground alignment) of being able to use the last known heading and leveling information if attempted soon enough after the Kalman filter ceases working.

7.0 ABERRANT BEHAVIOR

The Kalman filter can suffer a fatal exception and cease working (i.e., crash). If this occurs, look for loss of symmetry in the system covariance matrix caused perhaps by round-off or truncation error, particularly in single precision applications. Fatal exceptions are typically caused by a lack of exception handling for errors such as square roots of negative numbers, stack overflow, null pointers, and other errors associated with risky constructs in various high-level computer languages.

The Kalman filter can reject all measurements and continue to propagate. If this occurs, look for a significance level that is too large, or the process noise covariance matrix values are too small.

The Kalman filter can accept all measurements and operate as a recursive least squares estimator. If this occurs, look for a significance level that is too small, or the process noise covariance matrix values are too large.

The Kalman filter can operate smoothly but diverge from the real solution. If this occurs, look for a missing state, i.e., a force or bias that is not modeled, or inadequate alignment, or faulty accelerometer/gyroscope integration.

If the navigation system accumulates error in its solution rapidly, look for inadequate alignment. The accumulation of error may actually take place slowly if the alignment is only moderately faulty, for example, as caused by not enough time in the alignment mode.

If the solution deviates from reality during periods of high maneuvering, one can expect to find that the significance level used in the reasonable test is too large, resulting in too many measurements being rejected.

The values in the process noise covariance matrix can reflect confidence, or lack thereof, in the accuracy of the model. Generally, large values in the Q matrix reflect a lack of confidence in the model, as evidenced by the solution tending to follow the measurement at the expense of the extrapolation via the model. The opposite is true for Q values being too small.

If there is an error in the recursive design resulting in the old measurements not being discarded, that could result in memory overflows.

Ill-conditioning problems can result in wild variations of the state vector.

If DCM orthonormalization is lost, this may be evidenced by chaotic orientation of body axes. This is not a direct result of a Kalman filter fault, but can occur in connection with navigation operation and can be difficult to diagnose.

Some navigation systems do not have the capability for in-air alignment in the event that the navigation system ceases operation. Then, once Global Positioning System (GPS) aiding is lost, it cannot be restored while in flight.

8.0 WHAT ARE THE HAZARDS?

The observable hazardous behavior that is possible if the filter design is faulty in the way described in Section 6.0 may include the following:

1. Position, time, and velocity references supplied by the navigation system are missing. (Navigator ground and inflight routing not needed with GPS.)
2. The position and timing accuracy reverts to thousands of feet instead of the centimeter accuracy sent from the navigation hardware and software to the mission computers.
3. Time accuracy synchronization is only to the second instead of to the millionth of a second.
4. Landing capability lost in areas where no landing systems exist, and there is limited visibility because of weather conditions, darkness, or vehicle design.
5. Flights to any point in the Continental United States (CONUS) would become problematic. Routing is longer, necessitating waypoint planning, using excessive fuel and time.

6. In military applications, target error increases without GPS.
7. In some unmanned applications, loss of GPS can result in the inability to locate landing strips, and so result in damage to or destruction of the vehicle if this hazard is not mitigated.

9.0 WHAT ARE SOME OF THE MITIGATIONS?

A mitigation is a feature that helps to alleviate the bad effects of a hazard. Possible mitigations include the following:

1. In-air restart of the Kalman filter (in-air alignment)
2. Instrument landing system (ILS)
3. Pilot visuals
4. On top position fixes
5. Lengthening filter alignments
6. Simpson's rule during samplings

10.0 BEHAVIOR – CAUSES RELATIONSHIP

Aberrant behavior of a navigation system is most likely found in the testing environment. However, if the behavior is not detected during tests because of inadequate testing, a hazard may exist that, without mitigation, could put personnel and/or an expensive piece of equipment at risk. Table 10-1 ties observed aberrant behavior to possible flaws in the design of the Kalman filter. This will give safety engineers a tool that can be used to identify possible causes of aberrant behavior.

TABLE 10-1. Observed Aberrant Behavior and Their Possible Causes.

Observed Behavior	Possible Causes
Fatal exception	Loss of symmetry in the system covariance eventually results in negative variances
All measurements rejected	Too large a significance level in the reasonableness test Q values too small
All measurements accepted	Too small a significance level Q values too large
Smooth operation but diverges from the real solution	Missing state (faulty model) Inadequate alignment Faulty accelerometer/gyroscope integration
The navigation system accumulates error in its solution rapidly	Inadequate alignment
Solution deviates from reality during periods of high maneuvering	Too many measurements rejected due to too large a significance level in the reasonable test
Solution tends to follow measurements, ignoring the extrapolation	Lack of confidence in the model, indicated by Q being too large
Solution tends to follow the extrapolation at the expense of the measurements	Q values too small
Stack/heap overflow	Old measurements not discarded; recursive design has an error
Wild variations of state vector	Ill-Conditioning
Chaotic orientation of body axes	DCM orthonormalization lost
Inability to restore GPS	Absence of in-air realignment

11.0 SOME EXAMPLES OF ABERRANT BEHAVIOR RELATED TO FILTER DEFECTS

Some of the following experiments are a modification of examples found in Reference 7. There are others that were formulated by the authors.

The experiments are described in the following sections.

11.1 EXPERIMENT 1 – STATIONARY VEHICLE, MEASUREMENTS WITH GAUSSIAN NOISE

An estimate of the location of a vehicle believed to be stationary is desired. A Kalman filter with a single state (position) is used to derive this estimate. Noisy sensor measurements of this position are used to attempt to establish the location of the vehicle. The issues are as follows:

1. The position is completely unknown. Relative to a local coordinate system, the location is assumed at 0.0. To reflect the lack of confidence in this assumption, the variance (P^{-}) of the position is set at 1,000 ft^2. There is an external measurement every 1 second.

2. One of the critical assumptions of a Kalman filter is that the noise characteristics of the measurements are Gaussian. Considering that the number of sample points is small, a normal probability plot can be used to verify this assumption, using the Normal Probability Plot Maker (Reference 8). The measurements are normalized N(0,1) (Table 11-1).

TABLE 11-1. Measurements File for Experiment 1.

t, sec	Z, ft
1	0.9
2	0.8
3	1.1
4	1.0
5	0.95
6	1.05
7	1.2
8	0.9
9	0.85
10	1.15

If Figure 11-1 even remotely resembles a straight line, it can reasonably be concluded that the measurements are a sample from a data set with Gaussian noise.

The filter is not expected to perform correctly if the measurement noise does not reflect accurately the measurement statistics. The initial measurement variance (R) is set at 0.1 ft^2.

Another approach at this analysis is to compute the variance of the measurements, but remember that these measurements are merely a sample of a universe of measurements with the stated variance, so the variance of the actual measurement set may vary. This is one of the problems in working with such a small sample, that the variance of the sample may be so different from that of the universe from which the sample is derived. For those interested in sample theory, consult the Central Limit Theorem in any reference on statistics.

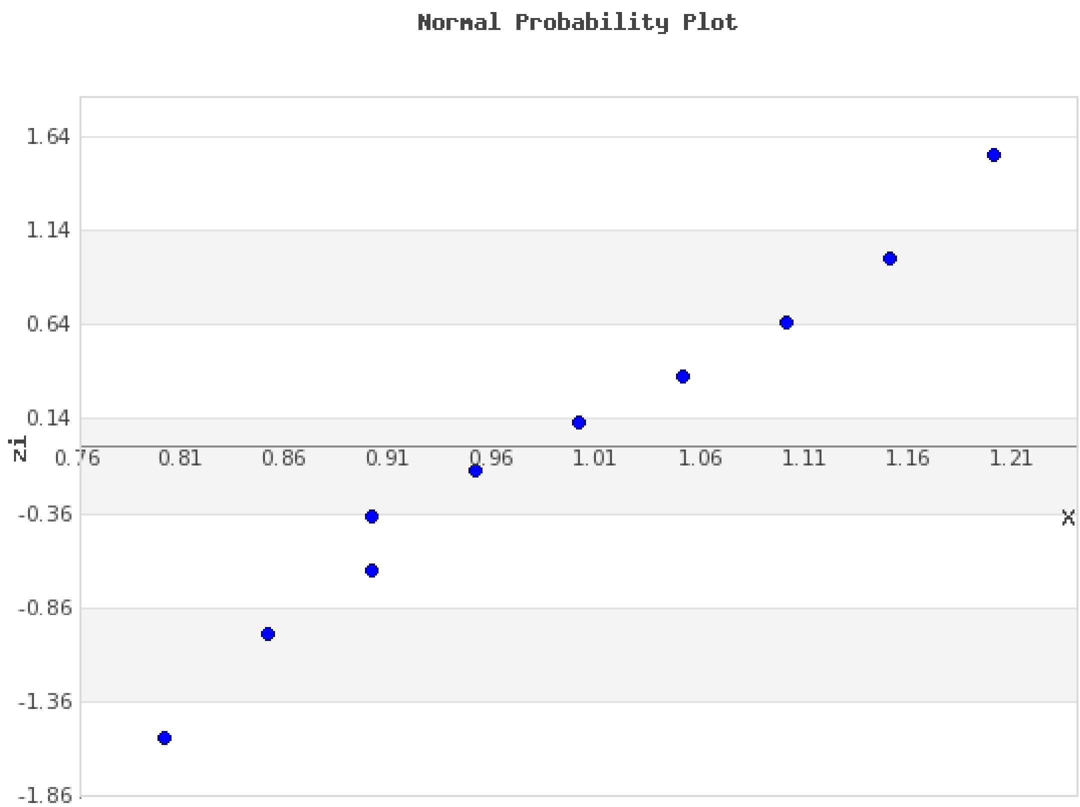

FIGURE 11-1. Normal Probability Plot, Experiment 1.

3. A model for the propagation of the state and system covariance is created. Confidence in the correctness of this model is high. To reflect this high confidence, the process noise covariance (Q) is set to a very small value 0.0001. The filter will not perform satisfactorily if the process noise covariance does not reflect the expected error model. Q can be used to tune the filter. Q is a constant scalar in this application.

Distribution Statement A.

11.1.1 Analysis

Even with a terrible original position estimate and noisy measurements, one can see in Figure 11-2 that the estimate converges quickly to the truth. Note that one of the characteristics of a good Kalman filter is that the estimate noise is much less than the measurement noise.

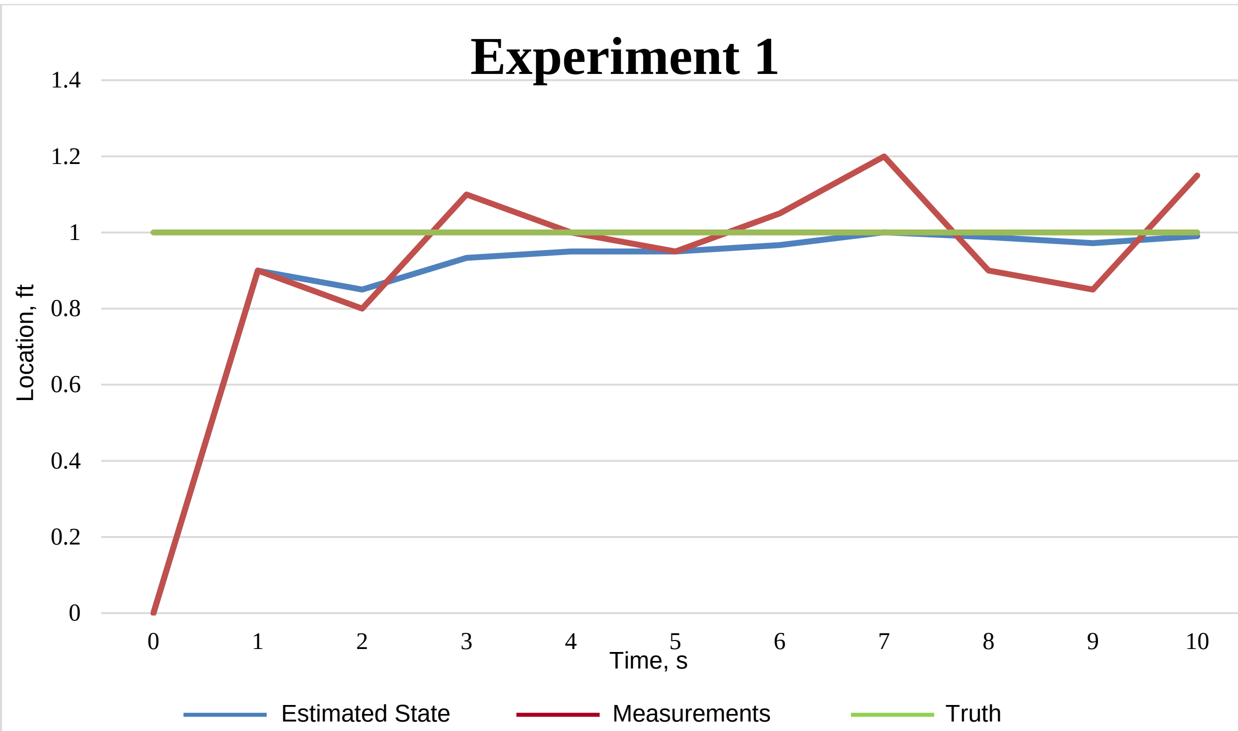

FIGURE 11-2. Result of Experiment 1.

11.2 EXPERIMENT 2 – ERRONEOUS MODEL WITH HIGH CONFIDENCE

In this navigation experiment, the vehicle is moving. We are trying to establish the position using noisy measurements (Table 11-2). This navigation system is subject to the following two major error sources:

1. Initial position error. The initial values of the system covariance reflect the confidence in the accuracy of the initial values of the position. The system covariance also reflects round-off and other process errors. As in Experiment 1, the initial position is unknown, estimated at 0, and the variance of the position is set at a large value, 1000 ft^2.

2. Model error. The model of Experiment 1 is retained. Since the filter is attempting to track a moving vehicle, the underlying model, which assumes a stationary vehicle, is

erroneous. The process noise covariance reflects the confidence in the accuracy of the model. The filter will not perform correctly if the process noise covariance does not reflect the expected error model. Q can be used to tune the filter. Q is constant with a small value in this application (0.001). This reflects a moderate confidence in the erroneous model.

TABLE 11-2. Measurements File for Experiment 2.

t, sec	Z, ft
1	0.11
2	0.29
3	0.32
4	0.50
5	0.58
6	0.54
7	0.75
8	0.79
9	0.90
10	1.09
11	0.98
12	1.35
13	1.42
14	1.58
15	1.45
16	1.50
17	1.59
18	1.83
19	1.84
20	2.15

11.2.1 Analysis

The solution is seen to proceed smoothly, except that it bears no resemblance to the truth (Figure 11-3). The solution drifts slowly away from the truth.

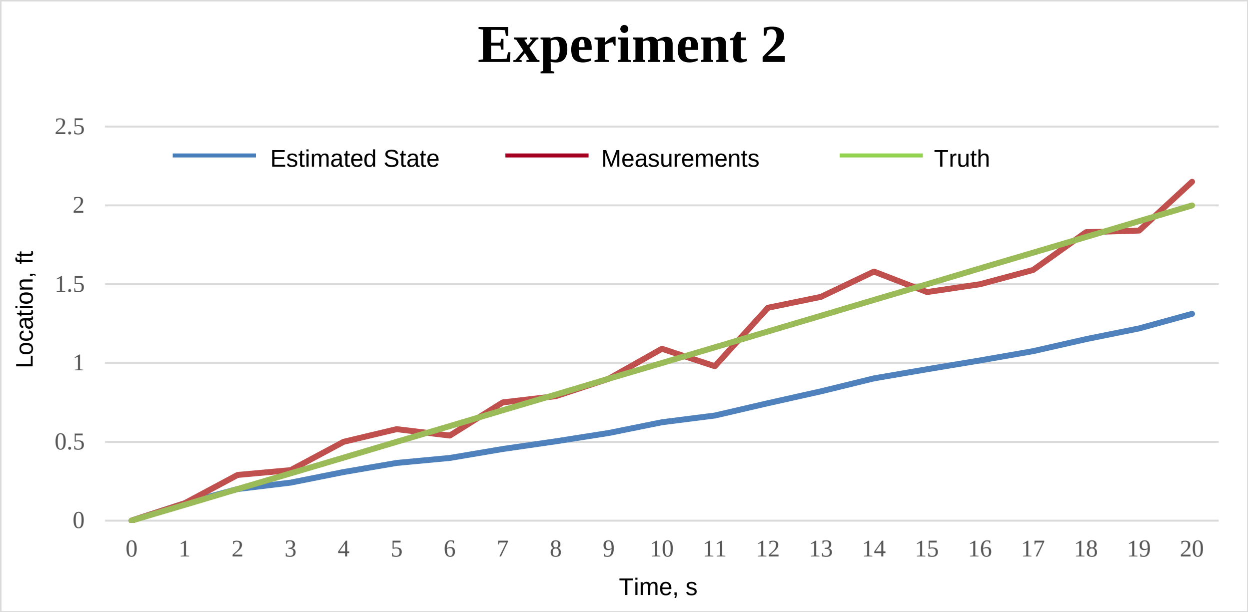

FIGURE 11-3. Result of Experiment 2.

11.3 EXPERIMENT 3 – ERRONEOUS MODE WITH SOME LACK OF CONFIDENCE

If one is not sure what the correct model is, one can enlarge the value of the process noise covariance, which will reflect some lack of confidence in the accuracy of the model. Let us try 0.01, keeping everything else the same (Figure 11-4).

11.3.1 Analysis

Now the solution follows the truth better, except that the noise of the estimated position is still larger than the measurement noise. Not good.

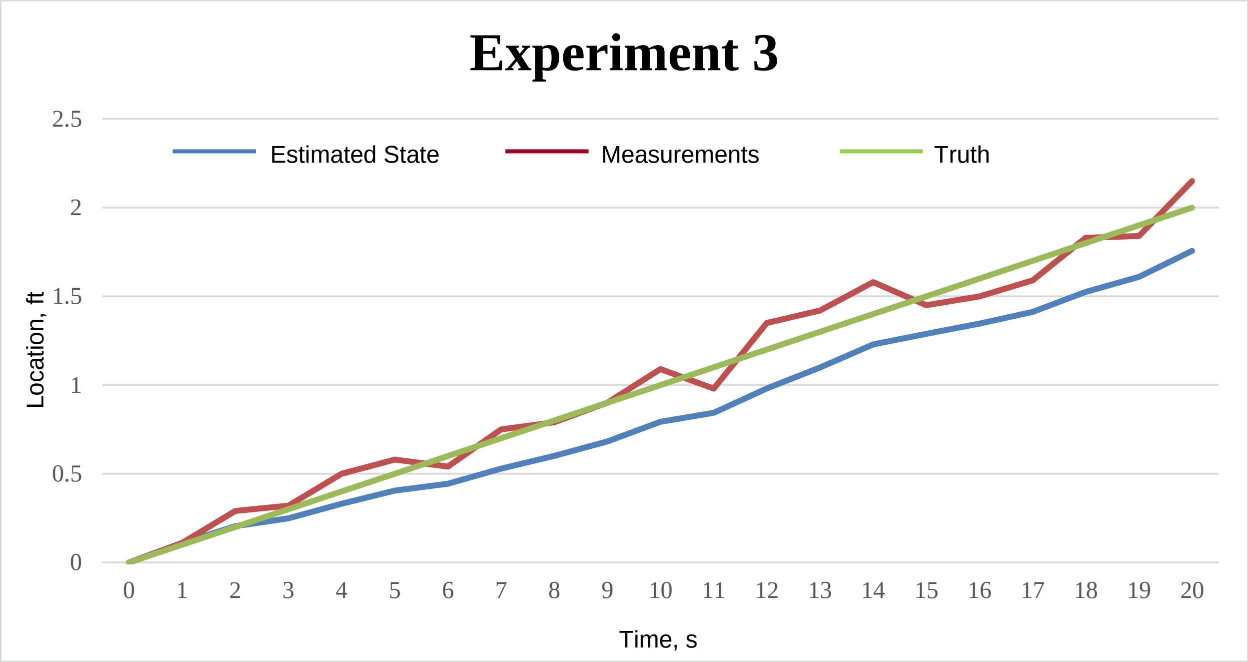

FIGURE 11-4. Result of Experiment 3.

11.4 EXPERIMENT 4 – ERRONEOUS MODEL WITH SERIOUS LACK OF CONFIDENCE

So let us enlarge the value of the process noise covariance by another factor of 10, to 0.1, and see what happens. We are telling the filter that we have a serious lack of confidence in the design of the model (Figure 11-5).

11.4.1 Analysis

The estimated solution seems to follow the data, not putting much trust in the propagation at all. Now the noise in the estimated solution is actually less than the measurement noise. Are we getting somewhere?

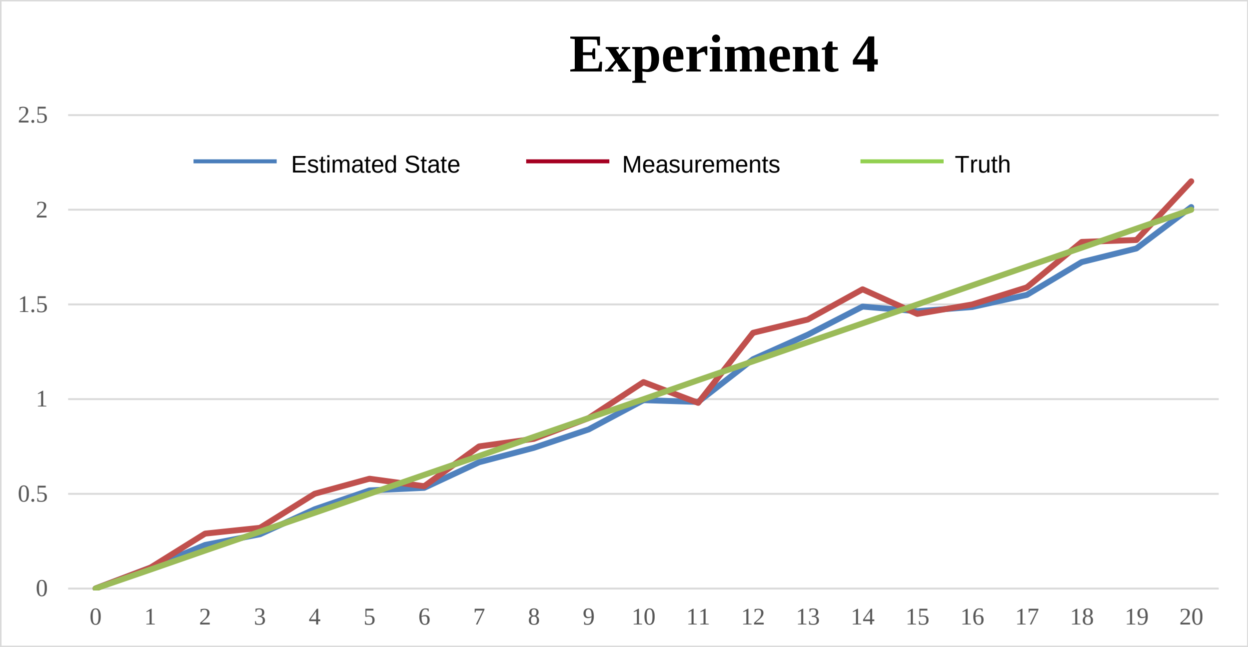

FIGURE 11-5. Result of Experiment 4.

11.5 EXPERIMENT 5 – ERRONEOUS MODEL WITH NO CONFIDENCE IN MODEL

In one last try, we enlarge the value of the process noise covariance to 1, announcing that we just do not trust the model. And sure enough, the estimated solution follows the data almost exactly, ignoring the propagation from the model altogether (Figure 11-6). So, why have a filter? Just trust the measurements!

11.5.1 Analysis

If your model is not correct, you will not get a good estimate. However, if you suspect this, you can increase *Q*, which feeds into the system covariance matrix and influences the Kalman gain, which in turn determines how much of the model state vs. the measurements to include in the updated value of the state. The risk is that the filter will tend to disregard the model propagation and believe the noisy measurements. You can see this when the estimate of the state begins to resemble the measurements.

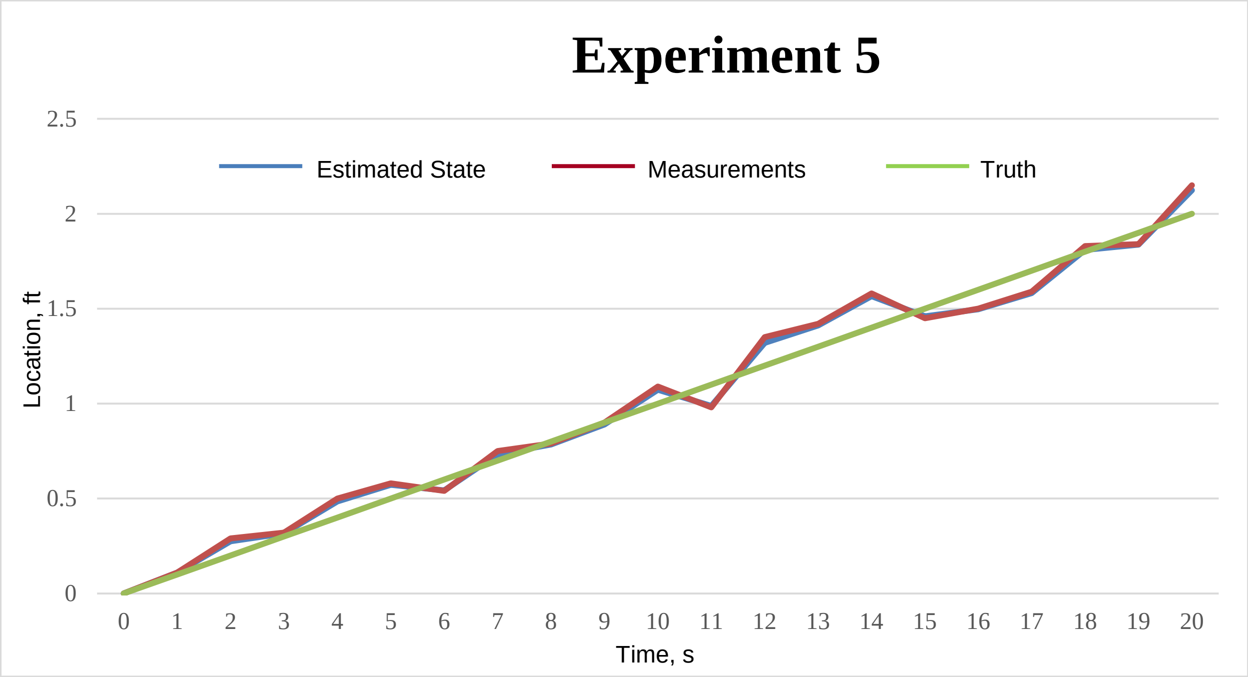

FIGURE 11-6. Result of Experiment 5.

11.6 EXPERIMENT 6 – MEASUREMENTS WITH NON-GAUSSIAN NOISE

What if the measurement noise is not Gaussian? Let us consider the following measurements (Table 11-3).

TABLE 11-3. Measurements File for Experiment 6.

t, sec	Z, ft
1	1.20
2	0.75
3	1.25
4	0.25
5	0.23
6	0.85
7	0.25
8	0.76
9	1.45
10	1.50
11	10.0

The model assumes a stationary vehicle at 1.0 foot. A normal probability plot of the measurements reveals that a straight line is hardly indicated by the standardized measurements. See Figure 11-7.

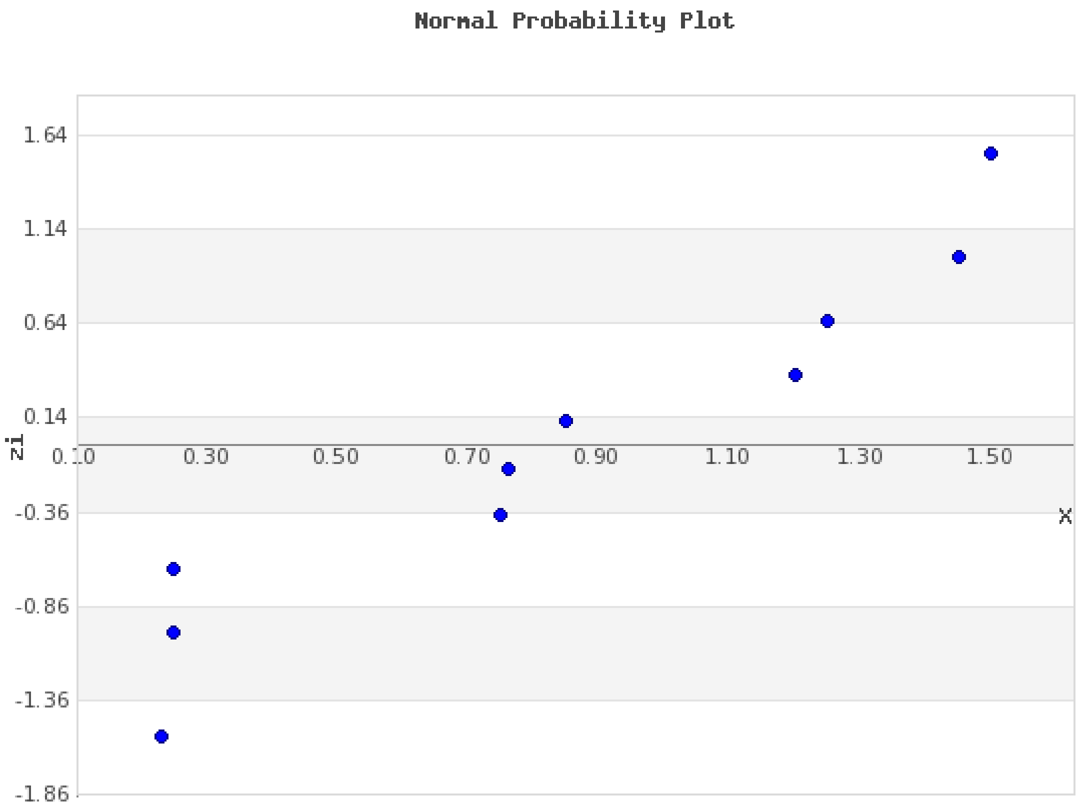

FIGURE 11-7. Normal Probability Plot, Experiment 6.

11.6.1 Analysis

It is not clear that the estimated solution (blue line in Figure 11-8) will ever converge to the truth, but if the measurements continue mainly to be biased to one side of the truth, there will probably be no convergence.

The last measurement reflects a sensor gone bad. That measurement was rejected. One can see that the estimated solution is not affected by the spurious measurements. The rejection of measurements criterion is set up so that all measurements beyond 3σ are to be rejected, but except for the last measurement, all measurements were accepted. This is called the data confidence level, not to be confused with the confidence in the model.

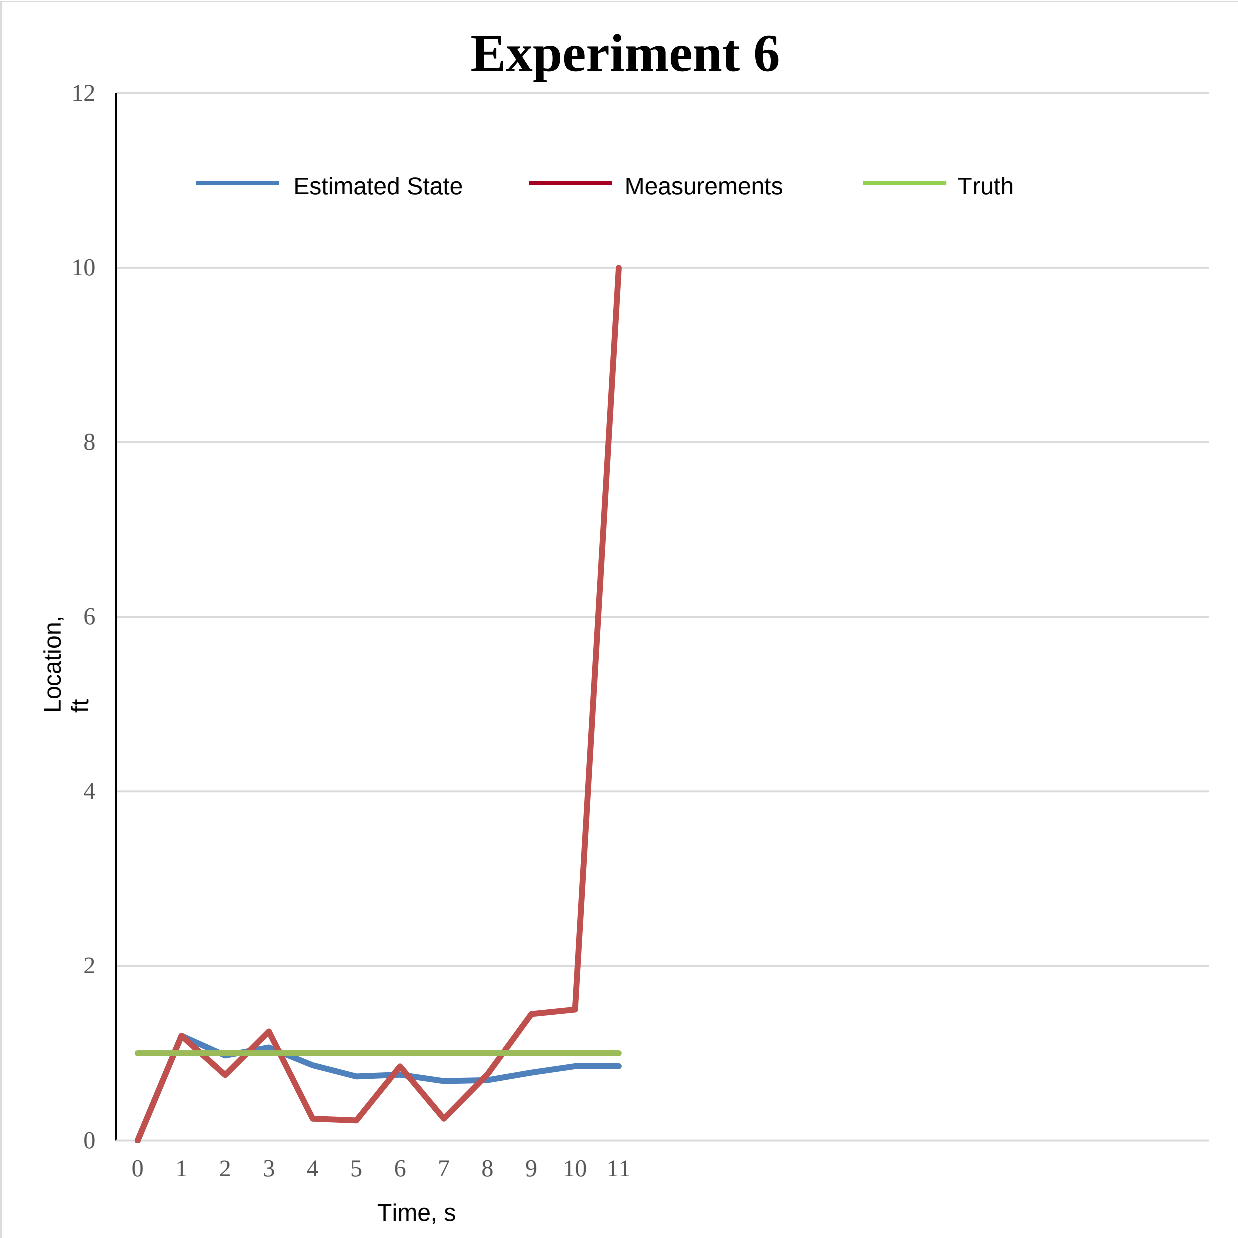

FIGURE 11-8. Result of Experiment 6.

11.7 EXPERIMENT 7 – REJECTION OF MEASUREMENTS

The model assumes that the vehicle at 1.0, stationary, and that the measurements are very noisy. The measurement distribution is not Gaussian, similar to Experiment 6 (except no spurious measurement at t = 11).

Q is small (Q = 0.0001), reflecting confidence in the model.

Initial position is completely unknown, 0 assumed with large variance (P = 1000 ft^2). All measurements beyond 2σ are to be rejected. Table 11-4 shows how to establish the χ^2 limit for 1 DOF at 2σ. The data confidence level is set at 95% while the significance level is at 5% (100% – data confidence level). The χ^2 limit is set therefore at 3.841.

TABLE 11-4. χ^2 Table, 2σ, 1 DOF.

DOF	Significance Level					
	0.5	0.10	0.05	0.02	0.01	0.001
1	0.455	2.706	3.841	5.412	6.635	10.827
2	1.386	4.605	5.991	7.824	9.210	13.815
3	2.366	6.251	7.815	9.837	11.345	16.268

11.7.1 Analysis

Measurements at t = 4, 5, and 7 seconds were rejected. The estimated solution appears much more reasonable relative to the truth (Figure 11-9). But what if these were valid measurements during maneuvers? 2σ may cause a rejection of too many valid measurements. This can be detected by observing that the solution deviates from reality during periods of high maneuvering.

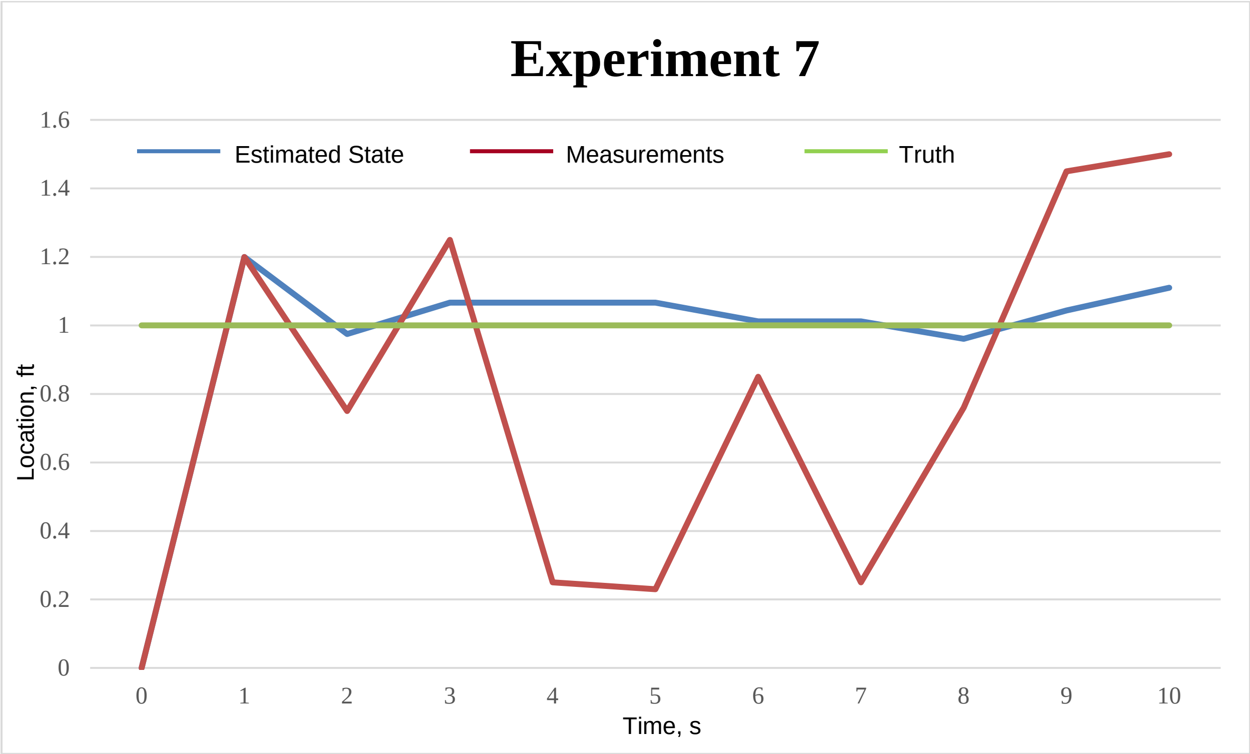

FIGURE 11-9. Result of Experiment 7.

11.8 EXPERIMENT 8 – REJECTION OF NOISY MEASUREMENTS DUE TO LOW DATA CONFIDENCE LEVEL

The model assumes that the vehicle at 1.0, stationary, and that the measurements are very noisy. See Table 11-5.

The measurement distribution is Gaussian.

Q is small (Q = 0.0001), reflecting confidence in the model.

Initial position is completely unknown, 0 assumed with large variance (P = 1000 ft^2). All measurements beyond 1σ are to be rejected. Table 11-6 shows how to establish the χ^2 limit for 1 DOF at 1σ. The confidence level is set at 68.27% while the significance level is at 31.73%. The χ^2 limit is set therefore at 1.0095.

TABLE 11-5. Measurements File
for Experiment 8.

t, sec	Z, ft
1	0.60
2	0.20
3	1.40
4	1.00
5	0.80
6	1.20
7	1.80
8	0.60
9	0.40
10	1.60

TABLE 11-6. χ^2 Table, 1σ With 1 DOF.

DOF		Significance Level					
	0.5	0.3173	0.10	0.05	0.02	0.01	0.001
1	0.455	1.0095	2.706	3.841	5.412	6.635	10.827
2	1.386	2.28	4.605	5.991	7.824	9.210	13.815
3	2.366	3.51	6.251	7.815	9.837	11.345	16.268

11.8.1 Analysis

The only measurements accepted are at times 1, 2, 8, and 9 seconds. However, the Kalman gain is fairly high at times 1 and 2, but low at 8 and 9 seconds, so the estimated state tends to follow the measurement at times 1 and 2, and the extrapolation (i.e., the model) in the remainder of the problem (Figure 11-10). Clearly the confidence level is set at much too low a level.

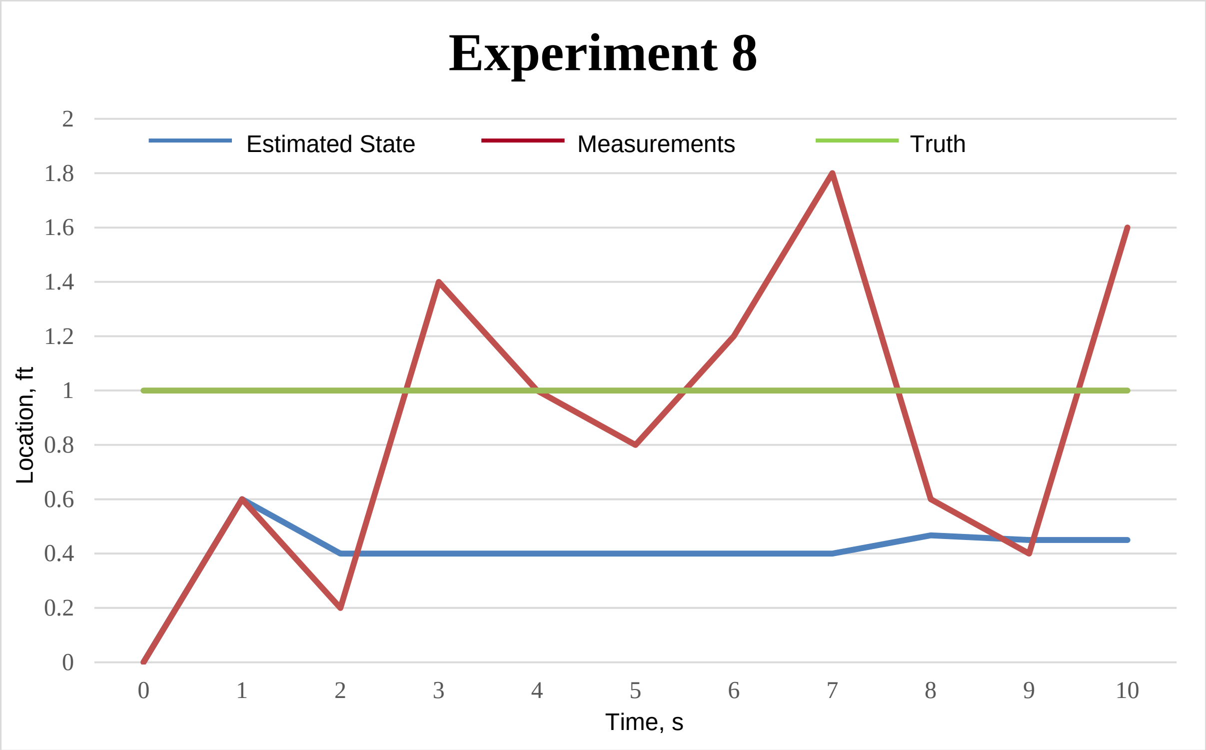

FIGURE 11-10. Result of Experiment 8.

12.0 SUMMARY

Kalman filtering has been used in a wide variety of aided inertial navigation systems. This report provides the system safety engineer a means to understand the structure of a typical Kalman filter found in the navigation system so that s/he can associate aberrant behavior in the operation of an air vehicle with defects in the Kalman filter. Kalman filtering and its overall role in the integrated system is still not well understood by many in the navigation safety community. This report is largely directed toward defining a tool for the safety engineer who is presented with the problem of including Kalman filtering into the design of a navigation system. This perspective is particularly useful in helping the engineer understand potential observed aberrant behaviors and their possible causal factors (see Table 10-1). Aberrant behavior is most likely observed in the software or hardware integration lab, but can be observed in the field if lab testing is not adequate.

REFERENCES

1. Arthur Gelb. “Optional Estimation. Introduction.” *Applied Optimal Estimation*, MIT Press, 1984, pp. 2-3.

2. Robert M Rogers. “Direction Cosine Matrices.” *Applied Mathematics in Integrated Navigation Systems*, 2nd ed., American Institute of Aeronautics and Astronautics, 2003, pp. 37-39. Education Series.

3. Nina P. G. Salau, et al. “A New Process Noise Covariance Matrix Tuning Algorithm for Kalman Based State Estimators.” *IFAC Proceedings Volumes*. Originally published in *ScienceDirect*, Vol. 42, No. 11, 2009, pp. 572-77.

4. Mohinder S. Grewal and Angus P. Andrews. *Kalman Filtering Theory and Practice*. Englewood Cliffs, New Jersey, Prentice-Hall, 1993.

5. StatLect. “Normal Distribution - Quadratic Forms.” *StatLect*, www.statlect.com/probability-distributions/normal-distribution-quadratic-forms. Accessed 4 December 2018.

6. Paul G. Savage. *Strapdown Inertial Navigation Lecture Notes, Feb. 1, 1990*. Plymouth, Minnesota, Strapdown Associates, Inc.

7. “Subject MI63: Kalman Filter Tank Filing.” *Academia*, September 2008, www.academia.edu/7547346/. Accessed 26 June 2018.

8. Mathcracker. “Normal Probability Plot Maker.” *Mathcracker*, 2018, mathcracker.com/ normal-probability-plot-maker.php#results. Accessed 26 June 2018.

ACRONYMS

CONUS	Continental United States
DCM	Direction Cosine Matrix
DOF	Degrees of Freedom
ECEF	Earth-Centered Earth-Fixed
ECI	Earth-Centered Inertial
GPS	Global Positioning System
ILS	Instrument Landing System
IMU	Inertial Measurement Unit
NAWCWD	Naval Air Warfare Center Weapons Division

INITIAL DISTRIBUTION

1 Defense Technical Information Center, Fort Belvoir, VA

ON-SITE DISTRIBUTION

2 Code 4F0000D (archive copies)
1 Code 4G0000D (file copy)
4 Code 416100D
  Cantone, A. (2)
  Chirkis, K. (2)