Image Cover Sheet

CLASSIFICATION	SYSTEM NUMBER
UNCLASSIFIED	148249

TITLE

AN HIERARCHIC ALLIANCE OF FILTERS FOR FAULT TOLERANT NAVIGATION USING TWO INERTIAL SYSTEMS WITH AIDING SENSORS

System Number:

Patron Number:

Requester:

Notes:

DSIS Use only:

Deliver to: BA

An Hierarchic Alliance of Filters for Fault Tolerant Navigation Using Two Inertial Systems with Aiding Sensors

J. Chris McMillan, Jeff S. Bird
Communication and Navigation Section
Defence Research Establishment Ottawa
Ottawa, Ontario, Canada K1A 0Z4

Dale A.G. Arden
Computing Concepts Ltd.
10 Phylis St.
Nepean, Ontario, Canada K2J 1V2

1. SUMMARY

A Dual Inertial Integrated Navigation System (DIINS) is being developed for the Canadian Navy to improve the navigational accuracy and reliability on ships which have two inertial navigators plus other aiding navigation systems and sensors such as GPS, Loran-C, Omega, Doppler Speed Log(s) and so on. The sensor integration architecture being proposed to optimally combine all navigation sensors on such a vessel is an "hierarchic alliance" of Kalman filters, which is designed to allow sensitive error compensation as well as complete fault detection isolation and reconfiguration (FDIR). This architecture is ideally suited to central processing, can take advantage of parallel processing, and provides significant advantages over both the conventional unifilter approach and the "federated" (or cascaded) filter approach.

This hierarchic alliance consists of a specific set of optimal filters running in parallel, with each filter processing measurements from a different subset of the navigation sensors. These filters can be partially ordered so that primary and secondary filters can be defined. The primary filter(s) provide the optimal navigation solution, while the secondary filters provide uncorrupted backup in the event of a sensor fault.

The primary motivating factor for this architecture is to provide optimal integration under all conditions, and in particular after the occurrence of subtle sensor faults which could not be immediately detected and which could therefore corrupt the primary filter(s). This alliance of filters can provide an uncorrupted optimal solution, since it can be configured so that there will always be a secondary filter which, at the time of failure, was running independently of the faulty sensor. This removes the usual need to "back out" of a failure which was not immediately detected and thus substantially simplifies reconfiguration in response to such a failure.

Another motivating factor for this architecture is that the partial independence of the parallel filters also facilitates the detection and isolation of sensor faults. This can be accomplished by multiple levels of statistical hypothesis testing on a set of parallel Kalman filters. Fault detection techniques used include the usual sensor data reasonableness and filter residual tests as well as a chi-square hypothesis testing technique applied to the state vectors, and inter-filter voting applied to the residuals test results.

While this approach is computationally intensive, modern software techniques, and soon to be available processing power, are expected to make the real time implementation of this hierarchic alliance of filters quite practical. The system envisioned in this paper is being designed and built at the Defence Research Establishment Ottawa for the Canadian Navy and is expected to see initial real-time sea trials in 1993.

2. INTRODUCTION

During the 1980's the Defence Research Establishment Ottawa (DREO) developed a Kalman filter based Marine Integrated Navigation System (MINS), as described in McMillan (1990), that optimally integrated the navigation sensors found aboard Canadian naval vessels at the time. These included GPS, Transit, Loran-C, Omega, speed log, gyrocompass, and operator entered sextant measurements and position fixes. MINS employs a

single 17-state, dynamically re configurable, complementary Kalman filter (UD formulation) with the "system" states representing the dead reckoning errors. MINS units are being installed on all major Canadian naval vessels with the exception of the new patrol frigates which are still under construction. Unlike existing ships, these new frigates do not have dead reckoning systems, but are being equipped with twin inertial navigation systems (installed fore and aft) rather than a gyrocompass. Future submarines are expected to be similarly equipped, and present the added problem of submerged navigation, which is largely unaided. Consequently, a next generation of the MINS system will be required to take full advantage of the redundancy available from the INS's on the new vessels.

The new integrated system currently being designed and simulated at DREO is called DIINS (Dual Inertial Integrated Navigation System) and is a substantial revision of MINS. It will optimally integrate at least the following sensors:

- INS_1,
- INS_2,
- GPS,
- Loran-C,
- Omega,
- speed log,
- Hyperfix,
- operator entered sextant fixes.

The presence of two inertial systems has required a completely new set of error models in the Kalman filters and their redundancy allows a much higher degree of fault tolerance to be built into the system. The primary objective of DIINS is to enhance the accuracy and reliability of the navigation system, and the FDIR architecture and algorithms are the key to accomplishing this. The algorithms are designed to detect and isolate (identify) sensor failures promptly and to automatically remove the failed sensor from the navigation solution without introducing a significant discontinuity in the output. Since full failure mode operation requires that DIINS be designed to integrate any viable subset of the marine sensors mentioned above, this flexibility provides the added benefit that DIINS could be used on virtually any marine platform.

The primary reason for having two relatively expensive INS's on each ship is to provide reliability in the event of a failure. Much effort has therefore been directed towards the prompt detection and isolation of subtle faults in an INS, such as an accelerometer bias shift (perhaps due to temperature control failure) or an out of spec. gyro bias shift. It will be shown that this can be accomplished, even in the submarine scenario where the only aiding sensor is the speed log.

Without the sensor integration provided by DIINS, a slow failure of an INS could not be quickly detected. It would not be clear that an INS has failed until the discrepancy becomes quite large. Early failure isolation (determining which INS has failed), would be even more difficult, especially in the case of a submerged submarine, where accurate aiding sensor position measurements are not available.

This paper outlines an optimal filtering architectural concept, and how it can be applied to this particular set of sensors. It also briefly describes the basic Kalman filter design, and various techniques currently being considered and implemented for fault detection, isolation and reconfiguration (FDIR) in this context. These include:

- measurement reasonableness tests,
- Kalman filter residual tests,
- multiple filter residual test voting, and
- a χ^2 (chi-square) filter self test technique that is applied to the state estimate and involves propagating a "shadow" filter (with no measurements) for each regular Kalman filter.

Some simulation results are also presented to illustrate the power of the FDIR techniques to detect and isolate subtle inertial sensor faults, even in the submarine case where only speed log aiding is available.

Since DIINS is still at an early stage of development, the designs outlined in this paper are only preliminary. Full simulations have yet to be completed to verify the performance of the fault detection, isolation and reconfiguration algorithms under a full set of expected failure conditions.

3. HIERARCHIC ALLIANCE OF FILTERS

3.1. General Concepts

To discuss the proposed architectural concept, some notation will be helpful. Consider the usual case where one continuous system (inertial or dead reckoning) is to be integrated with n aiding sensors (GPS, Loran-C etc.). The usual unifilter shall be referred to as F^o. This is envisioned as a complementary or error state filter, with a state vector partitioned as follows:

$$\begin{bmatrix} X_0 \\ X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix} = \begin{bmatrix} \text{inertial system / sensor error states} \\ \text{first aiding sensor error states} \\ \text{second aiding sensor error states} \\ \vdots \\ n^{\text{th}} \text{ aiding sensor error states} \end{bmatrix} \quad (1)$$

The measurement vector can also be decomposed:

$$\begin{bmatrix} Z_1 \\ Z_2 \\ \vdots \\ Z_n \end{bmatrix} = \begin{bmatrix} \text{first aiding sensor measurement} \\ \text{second aiding sensor measurement} \\ \vdots \\ n^{th} \text{ aiding sensor measurement} \end{bmatrix} \quad (2)$$

Let F^i (for $i=1,2,..n$) refer to the filter in which the i^{th} aiding sensor states and measurements have been removed. Let F^{ij} refer to the filter in which the i^{th} and j^{th} aiding sensor states and measurements have been removed. Similarly define F^{ijk} and so on.

The F^i can be considered the first generation filters, the F^{ij} second generation, and so on. The set of all such filters form a tree, as shown in Figure 1. Each filter in this tree can easily be spawned from the filter immediately above it (its parent filter), by simply eliminating the appropriate states and measurements. This reconfiguration can quite easily be done dynamically (in real time) without introducing any discontinuity, since the remaining states will be unchanged, as will their covariance matrix.

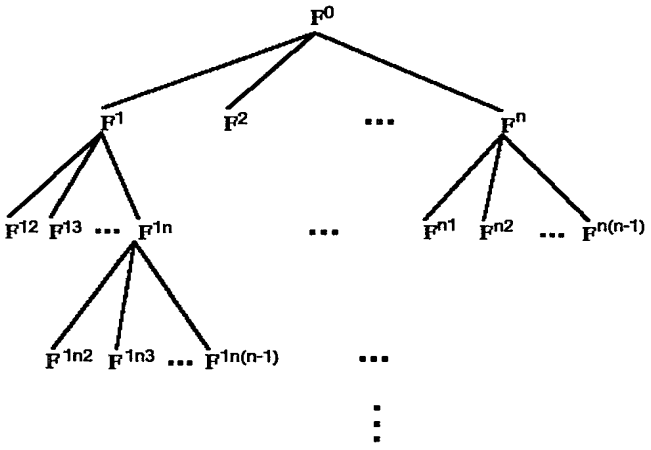

Figure 1. Hierarchic Alliance Structure

Any set of filters sharing a common parent filter, together with that parent shall be referred to as a family of filters. In the proposed approach, only one such family of filters is required at any given time. This provides the optimal filter (the parent) and a set of backup filters. The fact that each backup filter is independent of one of the sensors assists in all three aspects of fault detection, isolation and reconfiguration.

This architecture provides many advantages over both the conventional unifilter approach and the "federated" (or cascaded) filter approach. The different filters share sensor measurements only, rather than state vector information and therefore can be designed and run completely independently. They are therefore not subject to the filter design constraints which lead to sub optimality of the federated filters.

As discussed below, the partial independence of the parallel filters substantially improves FDIR, with failure detection being more sensitive, isolation more decisive, and reconfiguration much simpler.

With one INS and n aiding sensors it would be necessary to run at most $n+1$ filters at a time. With practical considerations, this number may be reduced. Although this may be considered burdensome, with modern software techniques and recently available processing power, the practical implementation of this hierarchic alliance of filters is now possible.

3.2. Reconfiguration

This architecture provides substantial capability for reconfiguration in the event of sensor failure, which is not available to a unifilter or to federated filters. The greatest difficulty with reconfiguration under both the unifilter and the federated filter schemes, is that the primary filter may have been corrupted by the sensor failure. This requires either a complete filter re initialization, with the resulting total loss of state estimate information, or some complex "backing out" scheme which attempts to use the state estimate from some pre-selected point in the past, hopefully before the failure occurred. The hierarchic alliance approach resolves this problem by simply having "hot" backup filters running at all times.

For any given set of aiding sensors, there will be a maximal filter in the complete tree (see Figure 1) which integrates all available (not failed) sensors. To provide an uncorrupted filter for single failures, it is only necessary to run the family containing the maximal filter as parent. For example if F^i is the maximal filter, and a soft failure (one who's detection is delayed) occurs in sensor j, then F^i will have been corrupted, but F^{ij} will not have been. Therefore F^{ij} can be used as the new maximal filter, with new subfilters F^{ijk} being spawned from F^{ij}. This is the basic hierarchic alliance reconfiguration concept.

It should be mentioned that if the above failure in sensor j was a hard failure, then F^i will not have been corrupted, in which case F^i would be used to spawn F^{ijk} directly, since F^i should have better state estimates than F^{ij}. Therefore if sensor j is not expected to ever have soft failures then the corresponding filter F^{ij} is not necessary for reconfiguration purposes. However, these secondary filters have an equally important role in fault isolation, as described below.

It should also be mentioned that inherent in this approach is the assumption that a second soft failure

will not occur in the time interval between the occurrence and detection of the first. It is therefore important to be able to detect these soft failures as quickly as possible in order to minimize this possibility.

Extending this technique to handle multiple simultaneous failures is possible and in fact is conceptually quite simple. To handle m simultaneous failures, it would be necessary to include all different filters for m generations below the maximal filter in the tree of Figure 1. For reasons of practicality however this has not been pursued further.

3.3. Failure Isolation

This architecture provides significant fault isolation capability which is not available to a unifilter or to cascaded filters.

When a sensor failure occurs, one consequence is that the error model assumptions underlying the Kalman filter are no longer valid. This can lead to unanticipated effects. Therefore some of the most sensitive fault detection schemes indicate only that there is a statistically significant problem within a Kalman filter but give no indication of where the fault lies.

For example the chi-square test indicates only that there is a statistically significant discrepancy between a Kalman filter's design model (represented by covariance information) and the measurements being processed by that filter (represented by state estimate information). Assuming that the fault is not with the filter design, it can then only be deduced that the sensor data is behaving in unexpected way, most likely due to a sensor failure. Some other means must then be used to determine which sensor is at fault. In an hierarchic alliance family of filters however, one and only one of the filters should be unaffected by the failure of a single sensor, namely the filter not using the failed sensor. It should therefore be possible to apply simple deductive logic to isolate the failed sensor. This is described in more detail in section 5.3 below, in the context of DIINS.

Furthermore, some standard FDI techniques, such as residual testing, are often used to isolate failures within a single filter even though they are actually ambiguous in this context. As will be explained in section 5.2 below, this ambiguity can be resolved by using residual test results from the multiple filters proposed.

3.4. Failure Detection

Failure detection and isolation are not entirely separable, since techniques which can detect but not isolate a failure are of limited use. The enhanced fault isolation capability of this architecture, as described above, therefore improves fault detection. It does this by allowing the use of more sensitive detection schemes such as the chi-square test, which would not otherwise allow isolation.

Ultimately sensor failures are detected by comparisons made between different sensor measurements (using *a priori* information regarding their different error behaviours and perhaps also regarding absolute physical limits on platform dynamics). A full family of filters, as described above, provides a richer set of sensor comparisons than is available to a unifilter or to a federation of filters. This should also improve the likelihood of failure detection.

With several different filters providing somewhat independent test results, voting can be used. This provides a level of verification not available to a unifilter, or even to a federation of filters. This can be used in two ways, or some combination of both:

- individual failure detection trip levels can be made lower (more sensitive) while maintaining the same level of confidence, or
- trip levels can be left the same, so that detection confidence would increase (false alarms reduced).

4. DIINS FILTERS

The presence of the second INS in DIINS slightly complicates this picture, since there will initially be two trees such as shown in Figure 1, plus a special purpose INS1/INS2 filter. Since both trees are equivalent, and are derived from the same unifilter, a brief description of this unifilter will be given.

4.1. The DIINS Unifilter

The usual multisensor filtering approach is taken. For navigation purposes this is normally a complementary, or error state, filter. This type of filter primarily estimates the errors of a continuous system such as an inertial (or dead reckoning) system, and coincidentally estimates the aiding sensor errors that are observable. The navigation output then comes from the primary sensor, corrected by the appropriate state estimates.

To minimize the effect of non linearity in the measurement equation, an extended filter formulation is used, whereby the filter measurements are formed by taking the difference between aiding measurements (such as GPS pseudo ranges etc.) and a prediction of what the aiding measurement should be at the filter's estimated position (or velocity etc.).

As long as an INS is functional, it is considered the primary sensor for position, velocity and attitude for navigation and related command and control

functions. In this case the INS error estimates are used to correct the output of the INS for these purposes.

The details of the error modeling developed for the unifilter are well beyond the scope of this paper, however it is useful to display the state vector and measurement vector. The unifilter state vector for DIINS is shown in Table 1.

Table 1. DIINS Unifilter State Vector

State Number	Description
1-2	INS horizontal position error
3-4	INS horizontal velocity error
5-7	INS platform misalignment
8-9	horizontal accelerometer biases
10-12	gyro biases
13-18	GPS Pseudo range errors
19	GPS clock bias
20	GPS clock drift
21-22	Speed Log errors
23-24	Ocean Current
25-26	Loran-C hyperbolic errors
27-33	Omega phase errors

This has been partitioned into the inertial system states (9), the inertial sensor error states (6), the GPS states (8), the Speed Log error states (4), the Loran-C error states (2) and the Omega error states (7). Of course there are two such filters corresponding to the two inertial systems.

The measurement vector can be similarly partitioned, as shown in Table 2. It is expected that more sensors may be added to this list as development proceeds.

Table 2. DIINS Unifilter Measurement Vector

Measurement Number	Description
1-2	Speed Log
3-8	GPS pseudo ranges
9-14	GPS pseudo range rates
15-16	Loran-C time delays
17-23	Omega phase differences
24-25	position fix

4.2. DIINS Filter Configuration

Since the two INS's provide more than one primary sensor, in the context of complementary filtering, it was only natural to consider multiple filters, even though it is still possible to construct a complementary unifilter. Although such a unifilter would provide optimal integration under normal conditions, a gradual sensor failure could corrupt the entire state vector before it was detected, causing serious problems. Individual filters that would combine one INS and one additional sensor, such as the federated filter of Carson (1987), were considered. This consists of a filter for each INS - single aiding sensor combination and one "federal" filter to integrate the subfilters. However, the interdependency and sub optimality of the filters, and the reliance on an uncorrupted federal filter precludes the availability of an optimal backup filter, and thus does not provide completely graceful degradation.

It was therefore decided to investigate an architecture such as described in section 3 above. This expands on ideas drawn from Widnall (1987).

Another basic characteristic of the DIINS system is a special filter called the INS1-INS2 filter that is very useful in helping to detect and isolate INS failures. Since this filter has very low measurement noise, it will be very sensitive to INS failures. However it is not directly used in the navigation solution since it estimates only the relative errors between the two INS's.

If the aiding sensors are referred to as S1, S2, S3 and S4 (representing for example GPS, Loran-C, Speed Log and Omega), then the initial set of filters comprising the hierarchic alliance are listed in Table 3, with the alliance structure illustrated in Figure 2. After reconfiguration this set will of course change, and if one of the INS's fails, then there will only be one tree. If both INS's fail, then a non-complementary filter would be required to integrate the remaining sensors. This level of reconfiguration is beyond the scope of this paper.

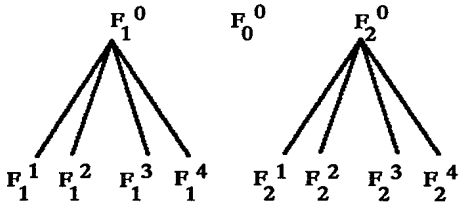

Figure 2. DIINS Initial Set of Filters.

Table 3: Initial Set of Filters for DIINS

Filter	SENSORS	
F_0^0	INS1, INS2	
F_1^0	INS1, S1, S2, S3, S4	(INS1 unifilter)
F_2^0	INS2, S1, S2, S3, S4	(INS2 unifilter)
F_1^4	INS1, S1, S2, S3	(i.e. F_1^0 w/o S4)
F_2^4	INS2, S1, S2, S3	(i.e. F_2^0 w/o S4)
F_1^3	INS1, S1, S2, S4	(i.e. F_1^0 w/o S3)
F_2^3	INS2, S1, S2, S4	(i.e. F_2^0 w/o S3)
F_1^2	INS1, S1, S3, S4	(i.e. F_1^0 w/o S2)
F_2^2	INS2, S1, S3, S4	(i.e., F_2^0 w/o S2)
F_1^1	INS1, S2, S3, S4	(i.e. F_1^0 w/o S1)
F_2^1	INS2, S2, S3, S4	(i.e. F_2^0 w/o S1)

5. DIINS FAILURE DETECTION

Of course, sometimes fault detection is trivial to the navigation officer monitoring the sensors. The DIINS system must be aware of such obvious faults even while it is monitoring the integration filters for subtle faults that the operator has yet to see. This suggests multiple levels of fault detection to be built into the system that can roughly be ordered as follows:

1. Operator disable - Often the operator knows that a sensor is faulty and should not be integrated.

2. Built-in-Test - The sensors' built-in-test indicators are used at a very high level to isolate a faulty sensor.

3. Reasonable Data - The raw sensor data is subjected to upper and lower bounds to ensure that it is physically realistic.

4. Consistent Data - These are tests on the physical reasonableness of the *change* in sensor output from one measurement to the next.

5. Reasonable States - The state estimates are tested against physically meaningful upper and lower bounds.

6. Residual Tests - The integration filter compares the magnitude of the filter residuals associated with the incoming sensor data with what it expects them to be based on the error models in the filters.

7. Inter-Filter Residual Test Voting - Simple deductive reasoning is applied to the results of the individual filter residual tests to isolate failed sensors.

8. Filter Self Test - Each filter can test its own state vector against the state vector of a "shadow" filter (an identical one propagated with no measurements).

The first 5 of these tests provide a very efficient and easily implemented prescreening of sensor measurements, before the more sensitive but more complex tests are performed. These simple tests are very effective in protecting filters from spurious measurements and rapid sensor failures.

The sixth and seventh (residual testing) provide a more sensitive, statistically based test which is effective in detecting and isolating discontinuity-type failures (hard failures) and is described in the following section. The eighth test (chi-square) is another statistically based test, which is most sensitive in detecting and isolating slow accumulative-type failures (soft failures). This is also described in more detail below.

5.1. Residual Testing

Assuming that a sensor measurement has passed the preliminary FDI tests (the first 5 tests listed above), it must still pass the residual test before being incorporated into the filter. A brief discussion of filter mechanization is needed to explain this test.

The filters are mechanized with Bierman's UDU^T form with scalar measurement updating and diagonalized noise covariance matrices. If $z(k)$ denotes the measurement vector for a filter at time k then:

$$z(k) = Hx(k) + w(k) \quad (3)$$

where H, $x(k)$, and $w(k)$ are the corresponding measurement geometry matrix, state vector and measurement noise vector, respectively. In a complementary error state filter, such as used in DIINS, these measurements z are generally *misclosures*, or differences between some quantity y_{meas} measured by an aiding sensor and a prediction of what that quantity should be, based on the position, velocity etc. of the primary sensor (the INS) y_{INS}.

$$z(k) = y_{meas}(k) - y_{INS}(k) \quad (4)$$

The component i of the residual vector ν, is then defined as

$$\nu_i(k) = z_i(k) - E\{z_i(k)|\hat{x}(k|k-1)\} \quad (5)$$

where $\hat{x}(k|k-1)$ is the estimate of the state vector of the filter at time k given all data up to time $k-1$. If the covariances of the state vector error and the measurement noise vector are denoted by P and R, respectively, then the variance of the residual is

$$E\{\nu(k)\nu^T(k)\} = HP(k|k-1)H^T + R(k) \quad (6)$$

where we have assumed a linear measurement equation (Eq. 3) and no correlation between the state vector error and the measurement noise.

If neither INS has failed, the DIINS residual testing algorithm uses F_0^0 and the most accurate pair of filters remaining (F_1^0 and F_2^0 in the no fail case). Note that each sensor is used in 2 of the 3 filters. Sensor failure detection and isolation using residual testing is implemented as follows:

1. At each update time, form the residual for each measurement in each filter, Eq. 5.

2. Form the estimate of the variance of each residual from the H, P, and R matrices from each filter according to Eq. 6.

3. If the square of the residual exceeds some preset number, A^2, times its filter computed expected variance:

$$\nu_i^2(k) > A^2[HP(k|k-1)H^T + R(k)]_{ii} \quad (7)$$

then that particular measurement is considered bad and is not incorporated in the Kalman filter update (typically A is set to 3 to provide a "3-sigma" test).

This is the usual application of residual testing within a single complementary filter, and constitute FDI test level 6 of the previous section. From Eq. 4 and Eq. 5 it can be seen that two assumptions are implicit in this method:

A. the residual ν in Eq. 5 is large because of the misclosure z rather than the state estimate $\hat{x}(k|k-1)$, and

B. the misclosure z in Eq. 4 is large because of the aiding sensor measurement y_{meas} rather than the primary sensor prediction y_{INS}.

Within a single complementary filter these are reasonable hypotheses because there is no useful response under the alternative hypothesis (a filter failure or a primary sensor failure). With the proposed architecture however, there is both the means to determine where the fault lies and the means to respond appropriately if the filter or the primary sensor (the INS) has failed. This is especially true when two primary sensors (INS's) are available, as is the case with DIINS.

5.2. Inter-Filter Residual Test Voting

The usual residual testing, described by the three steps above, can be extended as follows, to properly deal with failure of the primary sensor:

4. If a filter residual fails the 3-sigma test for at least M of the last N updates, then that residual is declared "suspect" in that filter.

5. If a residual is declared "suspect", then both sensors associated with that residual (the aiding sensor and the primary sensor) are declared suspect.

6. If a sensor i (aiding or INS) is "suspect" in all top level filters which use it, and its residual-associated sensor (INS or aiding) is not suspect in at least one top level filter which uses it, then the sensor i is declared *failed*.

Steps 5 and 6 above describe roughly how to apply deductive logic to these residual test results in order to isolate the failed sensor. Step 6 eliminates assumption B above by ascribing failure to the appropriate sensor. Assumption A is arguably less important, however it too can be somewhat eliminated by adding to step 6 a requirement intended to verify the validity of the state vector estimate. This is more difficult to couch in general terms, but in the case of DIINS it can be assumed that if the residual associated with the most accurate aiding sensor (and the INS) is not suspect, then the state estimate is not suspect.

In the DIINS baseline system of two INS's and four secondary sensors, the residual tests would be conducted in all 11 filters and measurements would only be included in those filters if they passed the residual tests. The top level filters, as shown in Figure 2 above, are the two unifilters, F_1^0 and F_2^0, along with the INS1-INS2 filter, F_0^0. Only these would participate in the sensor fault isolation of Steps 5 and 6 above. Table 4 below shows the logic table implementing these steps. This table is not complete, however it does includes all decisive cases.

Table 4: Residual Test FDI Voting

F_1^0 Residuals (INS1/S1/S2/S3/S4)				F_0^0 INS1 INS2	F_2^0 Residuals (INS2/S1/S2/S3/S4)				Failed Sensor
S4	S3	S2	S1		S1	S2	S3	S4	
.	.	.	X	X	0	.	.	.	INS1
.	.	.	0	X	X	.	.	.	INS2
.	.	0	X	0	X	0	.	.	S1
.	.	X	0	0	0	X	.	.	S2
.	X	.	0	0	0	.	X	.	S3
X	.	.	0	0	0	.	.	X	S4

0 - Residual is not suspect
X - Residual is suspect: M out of N residual rejections
. - Irrelevant

An "X" in the table means that a residual associated with that aiding sensor has failed the "3-sigma" test for M out of the last N samples, so that the residual is "suspect." For the aiding sensor to be declared failed, a residual formed from that sensor in both unifilters must be suspect and there must be evidence that neither the primary sensor (INS) nor the state vector estimate are at fault.

Acceptable residuals in F_0^0 provide evidence that the primary sensor (INS) is not at fault, while acceptable residuals for the most accurate remaining sensor in the unifilters provide evidence that the unifilter state vectors are not at fault.

5.3. Filter Self Test - The Chi-Square Test

Not all failures will be observable through filter residual monitoring. Any slowly accumulating error, such as an unmodeled INS errors, for example, would tend to be absorbed in the state estimates of other INS errors and would show little effect in the residuals. Thus a test on the state vector of a filter as a whole, as outlined by Brumback and Srinath (1987) and based on earlier work (see Kerr (1987)), is being investigated for use in DIINS. The idea is to have two solutions of each Kalman filter run in parallel, *one with measurements and the other without*. (We call the filter without measurements a "shadow" filter and look for a difference between a filter and its "shadow.") This shadow filter provides for a statistically significant reasonableness test that is applied to each filter's state estimate.

Consider one filter, to simplify the explanation. Define $\delta\hat{x}$ as the difference in the two state vectors:

$$\begin{aligned}\delta\hat{x} &\equiv \hat{x}_m - \hat{x}_e \\ &= \left(x - \tilde{x}_m\right) - \left(x - \tilde{x}_e\right) \\ &= \tilde{x}_e - \tilde{x}_m \\ &= \delta\tilde{x}\end{aligned} \tag{8}$$

where $\hat{x}_m$ is the state vector estimated using measurements and $\hat{x}_e$ is the state vector estimated without measurements (thus if $\hat{x}_e$ is initially zero then it will always remain zero). The covariance of the difference in the state estimates can be formed from the covariances of their errors:

$$\begin{aligned}P_{\delta\hat{x}} &= E\left[\delta\hat{x}\delta\hat{x}^T\right] = E\left[\delta\tilde{x}\delta\tilde{x}^T\right] \\ &= E\left[\tilde{x}_e\tilde{x}_e^T\right] - E\left[\tilde{x}_e\tilde{x}_m^T\right] - E\left[\tilde{x}_m\tilde{x}_e^T\right] + E\left[\tilde{x}_m\tilde{x}_m^T\right] \\ &= P_e - P_{em} - P_{me} + P_m\end{aligned} \tag{9}$$

It is fairly easy to show that under conditions of optimal filter gains, identical state models and identical initial conditions for both solutions that

$$P_{em} = P_{me} = P_m \tag{10}$$

so that

$$P_{\delta\hat{x}} = P_e - P_m \tag{11}$$

Since $\delta\hat{x}$ is Gaussian (it being the difference of two Gaussian vectors) of zero mean and covariance given by Eq. 10, its distribution is completely defined. The test for a failure consists of computing the scalar test statistic

$$k^2 = \delta\hat{x}^T\left[P_{\delta\hat{x}}\right]^{-1}\delta\hat{x} \tag{12}$$

The test statistic k^2 has a chi-square distribution with n degrees of freedom, n being the dimension of the state vector under test. A failure is declared with a confidence level of $(1-\alpha)$ when the test statistic exceeds the appropriate threshold:

$$k^2 > \chi^2_{n\alpha} \tag{13}$$

where $\chi^2_{n\alpha}$ is determined from the tables of the chi-square distribution. Note that α is the probability of false alarm; declaring a failure when there is none.

Note that this test can be just as easily applied to any

subset of the vector $\delta\hat{x}$, buy simply extracting its covariance matrix from $P_{\delta\hat{x}}$. The dimension n of the subvector used will of course affect the failure threshold. One current topic of investigation is to determine the most effective set of states for the purpose of detecting gyro and accelerometer failures.

Failure of this filter self-test, nominally called the chi-square test, indicates that the Kalman filter solution has deviated from its predetermined model (under the assumption that the filter has been properly modeled and tuned). The alternative solution from the shadow filter only propagates the initial conditions using the given model and cannot be corrupted by bad data. The failure of the test is attributed to measurement errors which in turn point to undetected sensor failures. There is not enough information in the test to conclude which sensor in the filter is not performing correctly. Sensor failure isolation requires multiple filters using different combinations of the available sensors.

Consider the full set of 11 filters listed in Table 3 and illustrated in Figure 2. If a chi-square test is done on each one of these filters (against their respective shadow filters) individual sensor failure isolation is possible. In Table 5 below, an "X" under a filter indicates that the filter has failed the test, a "0" indicates the test passed, and "." indicates a don't care. A sensor is isolated only in the scenarios listed. Other cases are ambiguous and no action is taken.

Table 5: Chi-Square Failure Isolation

Results of chi-square tests:											Isolate
F_1^1	F_1^2	F_1^3	F_1^4	F_1^0	F_0^0	F_2^0	F_2^4	F_2^3	F_2^2	F_2^1	Failure
.	X	X	X	X	X	0	0	0	0	0	INS1
0	0	0	0	0	X	X	X	X	X	.	INS2
0	X	X	X	X	0	X	X	X	X	0	S1
X	0	X	X	X	0	X	X	X	0	X	S2
X	X	0	X	X	0	X	X	0	X	X	S3
X	X	X	0	X	0	X	0	X	X	X	S4

0 - Filter passes chi-square test
X - Filter fails chi-square test
. - Irrelevant

The essence of Table 5 is that a sensor failure can be isolated by the chi-square test *only if* all filters that use that sensor fail and all filters that do not use that sensor pass. The only exception to this statement might be when isolating an INS failure (the first two cases in Table 5). It is probably better *not* to wait for filters F_1^1 and F_2^1 to fail. These filters are deprived of the most accurate aiding sensor, S1, and so they may not be sufficiently sensitive to promptly detect a soft INS failure.

5.4. False Alarms

Since a failure detection results in reconfiguration and long term loss of a sensor, the probability of false alarm must be kept very low; much lower than would be acceptable in a one-time measurement rejection test. Fortunately the voting schemes described in the previous two sections makes this possible without using very high trip levels on the individual tests. For example the INS failure detection by the chi-square test as shown in Table 5 requires the confirmation of five failed chi-square tests and five passed. If these ten individual tests were all independent, with individual false alarm probabilities of α, then the combined probability of a false alarm in either one of the INS's would be:

$$\Pr(\text{fa}) = 2\alpha^5(1-\alpha)^5 \qquad (14)$$

Thus a 95% chi-square threshold ($\alpha = .05$) on the individual tests would yield a very small combined false alarm probability of only $5\text{x}10^{-7}$. The individual thresholds could be increased to 90% or even 75% and still have a combined false alarm probability of .00001 or .0005 respectively.

A case of special interest is the submarine scenario, where only one aiding sensor (the Speed Log) may be available. In this case only two failed and one passed chi-square test are available to detect an INS failure. Thus individual test levels of 95% and 90% yield combined false alarm probabilities of .0048 and .018 respectively.

6. SIMULATION RESULTS: CHI-SQUARE FDI

This section presents some early simulation results performed to investigate the effectiveness of these chi-square tests. Different levels of accelerometer and gyro failures were simulated at different points in time, and different subsets of the state vector were used in the chi-square test. These were evidently among the factors determining effectiveness.

As would be expected, the detectability of a failure depends strongly on the magnitude of the failure simulated, since this determines the size of $\hat{x}_m$ and hence of $\delta\hat{x}$ and k(via Eqs. 8 and 12). For the purpose of this report, moderately accurate INS's are simulated (1 nmi/hour) and the inertial sensor failures are about eight times larger than specified as normal.

The temporal sensitivity arises from two independent factors:

- the presence of maneuvers, and
- the time since initialization of the shadow filter.

The first factor influences the observability of the inertial sensor errors, and hence the degree to which the corresponding states grow. The second factor is probably a result of the growth in the covariance of the shadow filter state estimates, particularly for the position states.

Two different subsets of the state vector were tested; one with seven elements and one with five. As will be shown below, the difference in effectiveness was not very significant.

Figures 3 to 8 illustrate these effects by showing six different cases. The same 5.5 hour trajectory was used in each case. Each figure shows the test statistic k^2 (from Eq. 12) for the two top level navigation filters, each of which is using the same aiding sensors but a different INS. For each figure one INS has a sensor failure (either accelerometer or gyro) while the other INS has no failure. In all cases the failure was detected, however it is the speed and decisiveness of detection that is of interest.

In Figures 3 to 6 the failed sensor was the x-axis accelerometer. In Figures 7 and 8 the failed sensor is the x-axis gyro. The simulated failures were of the "ramp to constant" type, whereby an additional unmodeled error (accelerometer bias or gyro drift rate) is introduced at a zero level at the "failure time" and increases linearly over an interval (500 seconds) to a maximum value, thereafter remaining fixed. The maximum values used for these failures were $1000\,\mu g$ for the accelerometer and 0.1 deg/hr for the gyro.

Figures 3 and 4 can be considered the baseline results. These both illustrate the prompt detection of an accelerometer failure which occurs immediately before a ship maneuver, using a seven degree of freedom (dof) chi-square test. For Figure 3 only Speed Log aiding was used, while for Figure 4 both GPS and Speed Log aiding was used. From this we can see that the Speed Log alone was quite adequate.

Figure 5 illustrates the same situation as in Figure 3 (accelerometer failure before a maneuver, Speed Log aiding) except that a five dof chi-square test was used. This seems to be a slightly less decisive test in this case.

Figure 6 is also the same situation as Figure 3 (accelerometer failure, Speed Log aiding, seven dof) except that the failure takes place 50 minutes before the maneuver. This significantly delays the failure detection.

Figures 7 and 8 illustrate gyro failure detection using the seven dof chi-square test.

In all of these figures, 3 to 8, the difference between the test statistic k^2 from the failed and unfailed filters is quite dramatic. As discussed in 5.4 lower thresholds could probably be used to improve response time.

7. CONCLUSIONS

The techniques outlined in this paper (multi-filter residual testing and chi-square testing) show considerable promise in being capable of promptly and reliably detecting subtle inertial sensor faults, even with Speed Log aiding only. However further investigation is required to determine the best choice of filters and the most effective FDIR techniques. In particular, the choice of state vector elements for use in the chi-square testing, the choice of confidence levels for use in the individual tests, the choice of individual tests to combine for final failure detection, and so on. Longer simulations are also needed to determine whether or not the method continues to loose sensitivity as the shadow filter covariance grows.

8. REFERENCES

Bird, J.S., McMillan, J.C. and Arden, D.A.G. (1992). "A Highly Fault Tolerant, Dual Inertial Integrated Navigation System (DIINS)," Proceedings of the National Technical Meeting of the Institute of Navigation, pp 235-244.

Brumback, B.D., and M.D. Srinath (1987). "A Fault-Tolerant Multisensor Navigation System Design," *IEEE Trans. Aerospace and Electronic Systems, AES*-23, No. 6, pp. 738-756.

Carlson, N.A. (1987). "Federated Square Root Filter for Decentralized Parallel Processes," Proc. *National Aerospace and Electronics Conf.*, pp. 1448-1456. Also appears in *IEEE Trans. Aerospace and Electronic Systems, AES*-26, No. 3, 1990, pp. 517-525.

Kerr, T. (1987). "Decentralized Filtering and Redundancy Management for Multisensor Navigation," *IEEE Trans. Aerospace and Electronic Systems, AES*-23, No. 1, pp. 83-119.

Magill, D.T. (1965). "Optimal Adaptive Estimation of Sampled Stochastic Processes," *IEEE Trans. Automatic Control, AC*-10, No. 4, pp. 434-439.

McMillan, J.C. (1990). "MINS-B II: A Marine Integrated Navigation System," in *Navigation Land, Sea Air & Space*, M. Kayton editor, IEEE Press, NewYork, pp. 161-171.

Widnall, W. (1987) private communication.

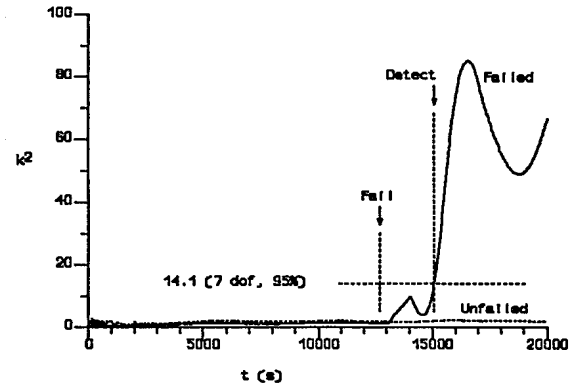

Figure 3. Accelerometer Failure, Speed Log Aiding, 7 dof Chi-Square Test

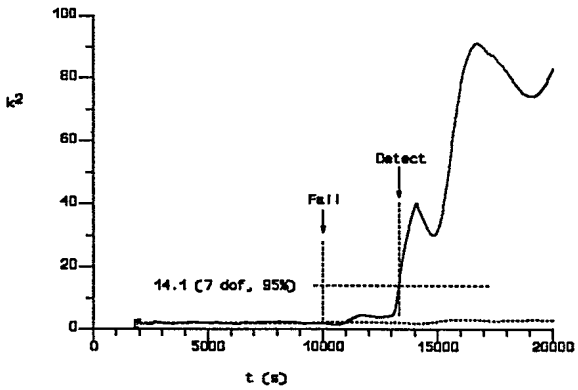

Figure 6. Accelerometer Failure, Speed Log Aiding, 7 dof Chi-Square Test, Delayed Maneuver

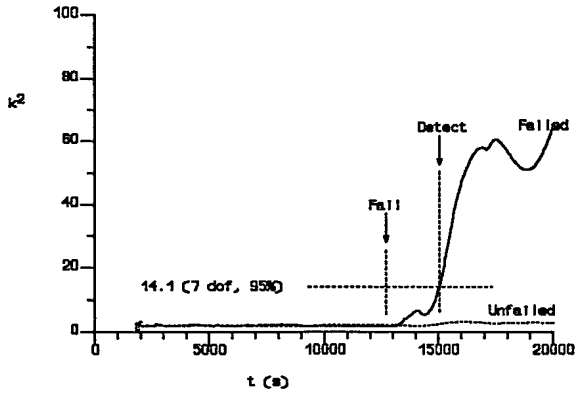

Figure 4. Acceleromter Failure, GPS & Log Aiding, 7 dof Chi-Square Test

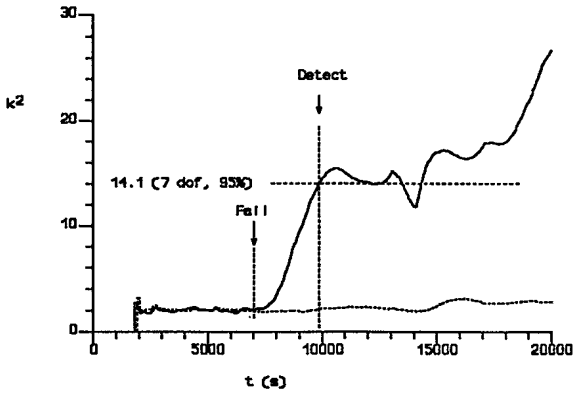

Figure 7. Gyro Failure, Speed Log Aiding, 7 dof Chi-Square Test

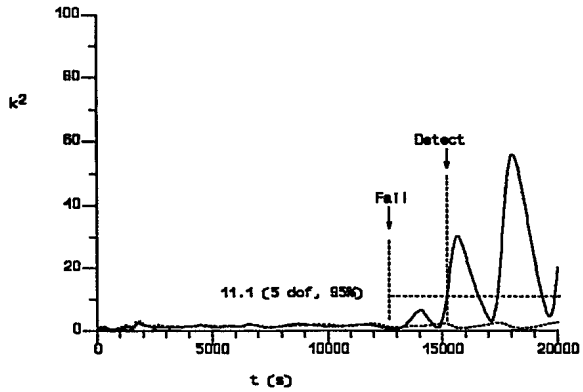

Figure 5. Accelerometer Failure, Speed Log Aiding, 5 dof Chi-Square Test

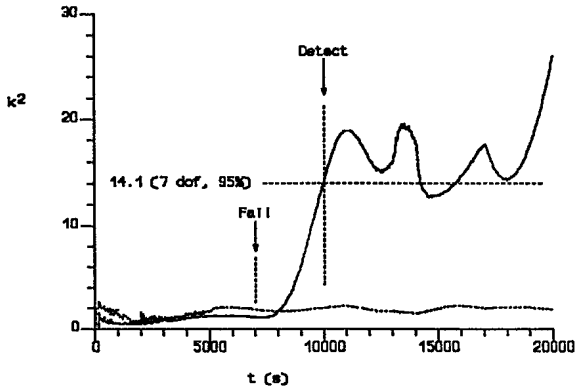

Figure 8. Gyro Failure, GPS & Log Aiding, 7 dof Chi-Square Test