

Sensor Fault Detection and Identification in a Mobile Robot*

Stergios I. Roumeliotis, Gaurav S. Sukhatme[†] and George A. Bekey

stergios|gaurav|bekey@robotics.usc.edu

Department of Computer Science
Institute for Robotics and Intelligent Systems
University of Southern California
Los Angeles, CA 90089-0781

Abstract

Multiple model adaptive estimation (MMAE) is used to detect and identify sensor failures in a mobile robot. Each estimator is a Kalman filter with a specific embedded failure model. The filter bank also contains one filter which has the nominal model embedded within it. The filter residuals are postprocessed to produce a probabilistic interpretation of the operation of the system. The output of the system at any given time is the confidence in the correctness of the various embedded models. As an additional feature the standard assumption that the measurements are available at a constant, common frequency, is relaxed. Measurements are assumed to be asynchronous and of varying frequency. The particularly difficult case of 'soft' sensor failure is also handled successfully. A system architecture is presented for the general problem of failure detection and identification in mobile robots. As an example, the MMAE algorithm is demonstrated on a Pioneer I robot in the case of three different sensor failures.

1 Introduction

Failure¹ detection and identification (FDI) are important problems in the development of reliable, robust mobile robots. In this paper we present the application of a technique called Multiple Model Adaptive Estimation (MMAE) to the case of the FDI problem in mobile robots. Specifically, we apply the technique to three independent sensor failures. Two of the failures are 'hard' in the sense that the sensor is assumed to be stuck at a certain value and the measurements available from it are subsequently ignored. The other failure considered

in this work is 'soft' in the sense that a sensor degrades in quality but is not completely useless. To recognize that a soft failure has occurred using the conditional probability computation used in the 'hard' cases is an unsolved problem [11]. We estimate on-line the new noise profile in the case of such a failure. It may be noted that once a 'soft' failure is identified correctly, the sensor may still be used, albeit with decreased confidence.

Kalman filtering [5, 7] is a well known technique for state and parameter estimation. It is a recursive estimation procedure using sequential measurement data sets. Prior knowledge of the state is improved at each step by taking the prior state estimates and new data for the subsequent state estimation.

Using a bank of Kalman filters was pioneered by Magill [6] who used a parallel structure of estimators in order to estimate a sampled stochastic process. Subsequently Athans et al. [1] used a bank of Kalman filters that provided state estimates to an equal number of LQG compensators to provide control over different operating regimes of an aircraft. Each estimator relied on a set of system equations linearized about a different operating point. Later Maybeck et al. [9] used the same technique (with an adaptive control strategy) to control F-15 aircraft. Further, in work by Maybeck et al. [11, 8] the multiple model adaptive estimation (MMAE) technique was used to reliably detect and identify sensor and actuator failures for aircraft. It is a frequent assumption in the literature that the sensor measurements are available at a common, constant frequency. In the experiments reported here that restriction is lifted using the sequential processing technique; proofs and details are available in [15].

In recent years Kalman filter based localization has become common practice [2, 13] in the robotics literature. Since the MMAE technique relies upon a bank of Kalman filters it seems natural to apply it to fault

*This work is supported in part by JPL, Caltech under contract #959816 and DARPA under contracts #F04701-97-C-0021 and #F04701-97-C-0021

[†]contact author for correspondence

¹In this paper fault and failure are used synonymously



Figure 1: The Pioneer I

detection and identification in mobile robot systems. The important aspect of the method is to use analytical redundancy in the form of several system models (as opposed to say hardware redundancy which replicates hardware to identify a failure). A Kalman filter based framework provides a measure of the disparity (typically called a residual) between the measured sensor values and the values predicted by the model embedded within the filter. The residual is used in the filter to update the estimate and is an excellent indicator of failure. We demonstrate this fact using three sensor failures as examples.

Earlier work in FDI (though not applied to mobile robots) is due to [4] and [16]. In [3] a network of adaptive virtual sensors is used to maintain reliable performance of a walking robot with many sensors, actuators and computers. In [12] the authors investigate fault-tolerant techniques using redundant sets of control strategies. Previous work using MMAE applied to the case of mechanical failures (flat tires) on board mobile robots is due to Roumeliotis et al. [14].

In section 2 the robot model is presented; in section 3 the general fault tolerant architecture is described including the details of the MMAE approach and in section 4 the experimental results are given. Section 5 concludes with a summary and a discussion of ongoing and future research.

2 Robot Model

The Pioneer I used for experiments is a three wheeled robot shown in Figure 1. The front two wheels are actuated independently thereby enabling differential steering. The rear wheel is a passive caster. The kinematics of the Pioneer I are given in Equations 1-2. Each wheel of the Pioneer has a position encoder. A gyro was used to measure the yaw rate of the robot.

$$v_L = r_L \dot{\theta}_L \quad v_R = r_R \dot{\theta}_R \quad (1)$$

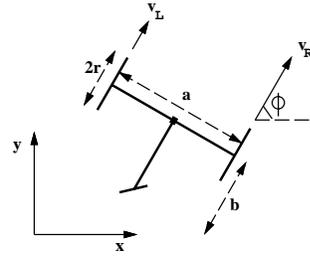


Figure 2: The Robot Kinematics

$$\dot{\phi} = \frac{v_R - v_L}{a} \quad v_{tot} = \frac{v_R + v_L}{2} \quad (2)$$

where a is the axle length, r_L and r_R are the radii of the left and right wheels respectively. The yaw rate of the robot in the x - y plane is denoted by $\dot{\phi}$ and the rotational speeds of the left and right wheels are denoted by $\omega_L = \dot{\theta}_L$ and $\omega_R = \dot{\theta}_R$. The linear speeds of the left and right wheels are denoted by v_L and v_R . These kinematic quantities are shown in Figure 2.

The basic idea behind the approach used in this work is to do fault detection by processing the residual signature of the Kalman filter and fault identification by having a particular filter respond to its matching failure with minimal measurement residual (and therefore maximal confidence). In the experiments reported here the measurement vector is composed of the two rotational speeds of the left and right wheels and the yaw rate of the chassis as measured by a gyro. The estimated measurement vector is denoted by $\hat{\mathbf{z}}$, the actual measurements from the sensors are denoted by \mathbf{z} and the residual vector is denoted by \mathbf{r} . We have

$$\mathbf{z} = [\omega_L \ \omega_R \ \dot{\phi}]^T \quad \hat{\mathbf{z}} = [\hat{\omega}_L \ \hat{\omega}_R \ \hat{\dot{\phi}}]^T \quad \mathbf{r} = \mathbf{z} - \hat{\mathbf{z}} \quad (3)$$

The various filters developed in this work use the kinematics in Equations 1-2 and the measurements shown in Equation 3.

3 Fault Tolerant Architecture

The proposed architecture is depicted pictorially in Figure 3. The thrust of the current work is in the fault detection and identification modules (FDI) as seen in the bold part of the figure. The control module is part of the future work. The first stage of the proposed approach is the detection of a fault. In the results presented here thresholding the residual is sufficient to detect a fault. However identification requires residual

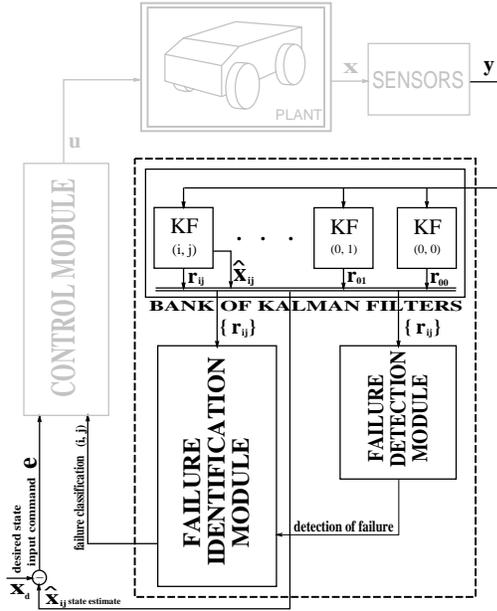


Figure 3: Failure Detection and Identification

postprocessing using a hypothesis conditional probability computation.

The bank of Kalman filters shown in Figure 3 is labeled using pairs of numbers (i, j) . If the first number in the ordered pair is the same for two filters then the same set of kinematic system equations (with different parameters) is used in both filters. Four filters were used in the experiments reported here; all share the same robot kinematics. If the first entry is unequal the two filters have different kinematic equations of the system. This allows a general application of the method covering the cases when the failure causes a change in the kinematic equations of the robot. Each filter produces a residual \mathbf{r}_{ij} . The nominal model results in a residual of \mathbf{r}_{00} . All the residual vectors are fed into the FD (failure detection) and FI (failure identification) modules. The output of the FD module is a signal notifying a failure. When this is true the FI module becomes active. The FD module stays active even after a failure has been detected in case the failure spontaneously disappears or a new failure appears. A small portion of this scheme (a bank of four filters) is implemented here. In the scheme proposed here, the key insight is to use a bank of estimators. The fault detection/identification depends on the selection of the state estimator with the minimal residual. This estimator is the one which assumes currently sound knowledge of the system description, i.e. it has incorporated

the failure effect in its structure. The FDI module thus provides a high quality estimate to the control module and allows for graceful degradation of the system. An increase in robustness is expected since the proposed scheme uses additional knowledge about expected failures. Thus it is capable of dealing with scenarios that in other schemes might be considered catastrophic. The third stage in the process shown in Figure 3 is labeled control. This stage deals with modifying the control structure of the robot (fault accommodation) so that it can continue functioning after the fault has occurred.

Failure identification is accomplished using hypothesis testing on the residuals. In this case we have four residuals (one from the nominal model and three from the failed cases). Accordingly, we form four hypotheses. The first hypothesis H_1 is that the system is operating nominally. The second hypothesis H_2 is that the gyro which measures yaw rate has degraded such that its noise profile has changed. The third hypothesis H_3 is a hard failure of the gyro - this is modeled as a stuck-at fault. The fourth hypothesis H_4 is that the left wheel encoder of the robot has suffered a hard failure (again modeled by a stuck-at fault). Hard failures are modeled using different parameters in the sensor matrix H . Soft failures are modeled using different parameters in the sensor noise covariance matrix R .

Following Maybeck [8] we compute the conditional density function of the measurement \mathbf{z} at time t_i given that failure k has occurred ($1 \leq k \leq 4$) and the measurement history $\mathbf{Z}_{i-1} = [\mathbf{z}^T(t_1) \dots \mathbf{z}^T(t_{i-1})]^T$ as

$$f_{\mathbf{z}(t_i)|\mathbf{a}_k, \mathbf{Z}_{i-1}}(\mathbf{z}(t_i)|\mathbf{a}_k, \mathbf{Z}_{i-1}) = \beta e^{-\frac{1}{2} \mathbf{r}_k^T(t_i) \mathbf{S}_k^{-1} \mathbf{r}_k(t_i)} \quad (4)$$

where \mathbf{a}_k is a vector of parameters specific to the k^{th} failure and the coefficient β is given by

$$\beta = \frac{1}{(2\pi)^{m/2} |\mathbf{S}_k|^{1/2}} \quad (5)$$

m is the dimension of the measurement vector, $\mathbf{z}(t_i)$ and \mathbf{S}_k is the covariance of the residual at time k . Maybeck [8] has shown that one way to achieve fast response to failures is to set $\beta = 1$. We successfully tested this approach (termed β elimination) for the different cases of sensor failure considered here.

The next step is to determine the conditional probabilities of the hypotheses H_1 through H_4 . These probabilities are denoted by $p_k(t_i)$, ($1 \leq k \leq 4$). The k^{th} probability is computed as the conditional probability that fault k has occurred given the measurement vector sequence \mathbf{Z} using the following iterative formula.

$$p_k(t_i) = \frac{f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}_{i-1}}(\mathbf{z}(t_i)|\mathbf{a}_k, \mathbf{Z}_{i-1})p_k(t_{i-1})}{\sum_{j=1}^4 f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}_{i-1}}(\mathbf{z}(t_i)|\mathbf{a}_j, \mathbf{Z}_{i-1})p_j(t_{i-1})} \quad (6)$$

It may be noted that $\sum_{k=1}^4 p_k(t_i) = 1$; each $p_k(t_i)$ has a maximum value 1 and an artificially imposed minimum of 0.01. This minimum value ensures that no failure hypothesis is too weak to influence subsequent future conditional probability evaluations quickly enough. The above expression for updating the conditional probabilities makes the assumption that all measurements are available at once. This is however not the case. As an example H_4 is the hypothesis that the left wheel encoder has suffered a hard failure. If at a given time instance t_i data is not available from this encoder then we update the conditional probabilities $p_k(t_i)$, $k = 1, 2, 3$ as follows:

$$p_k(t_i) = (1 - p_4(t_{i-1})) \frac{f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}_{i-1}}(\mathbf{z}(t_i)|\mathbf{a}_k, \mathbf{Z}_{i-1})p_k(t_{i-1})}{\sum_{j=1}^4 f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}_{i-1}}(\mathbf{z}(t_i)|\mathbf{a}_j, \mathbf{Z}_{i-1})p_j(t_{i-1})} \quad (7)$$

$$p_4(t_i) = p_4(t_{i-1}) \quad (8)$$

The idea being that since no information is available from the encoder and only a gyro measurement appears, the *a posteriori* probability of the encoder failing is considered to be the same as the *a priori* probability. The other terms are weighted suitably to ensure $\sum_{k=1}^4 p_k(t_i) = 1$ at time t_i . At this point it is worth mentioning that due to the β elimination the *a priori* probabilities do not sum to one. We relax this restriction because we are interested in their relative values only (as opposed to their absolute values). This is evident in the calculation of the *a posteriori* probabilities. The most probable hypothesis is forced to a value close to one while the rest of the probabilities are forced to stay close to zero. Due to Equations 6, 7 and 8, the *a posteriori* probabilities sum to one. This is valid in the cases where a single failure has occurred.

4 Experimental Results

In the experiments reported here, three faults were considered and four filters were used in the filter bank to process the data stream from the robot's sensors as it transitioned from its nominal behavior to the faulty behavior. In Figure 4 each subfigure shows the case of a different failure. The first one (at the top of the page) is the case of the soft gyro failure. In this case the gyro noise characteristics are assumed to degrade by increasing the standard deviation of the noise by a

factor of 2 at time $t = 14$ seconds. The nominal conditional probability (solid line) responds by decreasing sharply and the soft failure conditional probability rises sharply to 1. After a soft failure occurs the new level of noise is calculated on-line following [10].

The middle subfigure in Figure 4 shows the case of a hard gyro failure. It is assumed that the gyro is stuck at a fixed value at time $t = 14$ seconds. As before the nominal conditional probability responds by degrading sharply (solid line). However there is a small interval of time where both the soft and hard gyro failures respond positively. This is to be expected since the system cannot immediately decide whether the failure is due to a change in noise characteristics or due to a constant bias. However more samples (and hence a larger time interval) allow the system to calculate its conditional probabilities appropriately and ultimately the hard gyro failure conditional probability saturates to 1 as expected and the soft gyro failure conditional probability decays back to near 0.

The bottom subfigure in Figure 4 shows the case of a hard encoder failure on the left wheel of the robot. The encoder on the left wheel is stuck at a fixed value at time $t = 10$ seconds. The response of the nominal model conditional probability decays from 1 to nearly 0 rapidly and the response of the hard encoder failure filter rises to 1. It may be noted that in each of the three case no false positive identification occurred. In each case the system was able to detect and correctly identify the failure in a timely way.

It will be noted that the soft gyro failure assumed a specific noisy degradation in the gyro data by increasing the standard deviation of the gyro measurement noise by 3. In practice however, one is unlikely to know *a priori* the exact noise profile after the degradation. In Figure 5 we show the case of a soft gyro failure with a different (smaller) noise degradation (by a factor of 1.5 in the noise standard deviation). As the figure shows the system correctly detects and identifies the failure, though with a slower response time. This provides preliminary evidence that the FDI system is robust and will function in the unmodeled soft failure regime as well.

5 Conclusions and Future Work

In this paper we have presented a Multiple Model Adaptive Estimation (MMAE) based technique for sensor failure detection and identification on board a mobile robot. Experimental evidence is presented to show that the technique works well for several different failures. The implementation described here is able to use measurements from several sensors asynchronously and

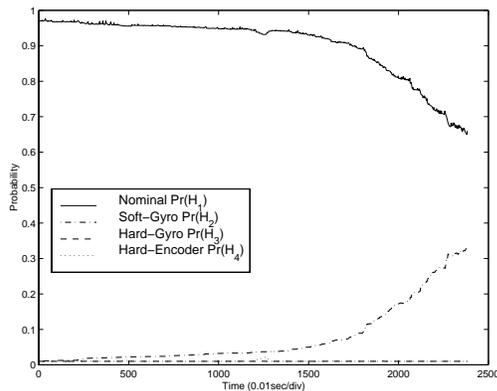


Figure 5: An example of an unmodeled soft gyro failure. The response is accurate but slow.

at non-constant frequencies. Soft gyro failure (modeled by an increase in noise standard deviation) is correctly identified by the system using online estimation. In all cases, correct and timely failure identification preserves the validity of the state estimate along with an updated associated confidence level. In the present treatment the appearance of a bias in the gyro measurement is treated as a hard failure. In the future we plan to incorporate existing models of gyro drift into the system to convert this case into a soft failure case so that the information from a biased gyro need not be discarded as it presently would be. Future research is also aimed at a better understanding of soft failure modes as well as extensions to other sensor failure applications. This work is part of a larger research effort which seeks to detect and identify common failures in mobile robots as well as reconfigure the system once identification is done. To that end future work will also include research into actuator and mechanical failure diagnosis and recovery.

Acknowledgments

The authors thank B. Werger, N. Mohan and A. Mohan for help with data collection. The authors also thank B. Balam, S. Hayati, G. Rodriguez, P. Schenker, R. Volpe and C. Weisbin for discussions.

References

- [1] M. Athans, D. Castanon, K-P Dunn, C. S. Greene, W. H. Lee, N. R. Sandell Jr., and A. S. Willsky. The stochastic control of the f-8c aircraft using a multiple model adaptive control (mmac) method- part i: Equilibrium flight. *IEEE Transactions on Automatic Control*, AC-22(5):768–780, October 1977.
- [2] B. Barshan and H. F. Durrant-Whyte. Inertial navigation systems for mobile robots. *IEEE Transactions on Robotics and Automation*, 11:328–342, January 1995.
- [3] C. Ferrell. Failure recognition and fault tolerance of an autonomous robot. *Adaptive Behavior*, 2(4):375–398, 1994.
- [4] D. T. Horak. Failure detection in dynamic systems with modeling errors. *AIAA Journal of Guidance, Control and Dynamics*, 11(6):508–516, Nov-Dec 1988.
- [5] R. E. Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 86:35–45, 1960.
- [6] D. T. Magill. Optimal adaptive estimation of sampled stochastic processes. *IEEE Transactions on Automatic Control*, AC-10(4):434–439, 1965.
- [7] P. S. Maybeck. *Stochastic Models, Estimation and Control*, volume 1. New York: Academic Press, 1979.
- [8] P. S. Maybeck and P. D. Hanlon. Performance enhancement of a multiple model adaptive estimator. *IEEE Transactions on Aerospace and Electronic Systems*, 31(4):1240–1253, October 1995.
- [9] P. S. Maybeck and D. L. Pogoda. Multiple model adaptive controller for the stol f-15 with sensor/actuator failures. In *Proceedings of the 20th Conference on Decision and Control*, pages 1566–1572, December 1989.
- [10] R. K. Mehra. On the identification of variances and adaptive kalman filtering. *Transactions on Automatic Control*, AC-15(2):175–184, April 1970.
- [11] T. E. Menke and P. S. Maybeck. Sensor/actuator failure detection in the vista f-16 by multiple model adaptive estimation. *IEEE Transactions on Aerospace and Electronic Systems*, 31(4):1218–1229, October 1995.
- [12] D. Payton, D. Keirse, D. Kimple, J. Krozel, and K. Rosenblatt. Do whatever works: A robust approach to fault-tolerant autonomous control. *Journal of Applied Intelligence*, 2:225–250, 1992.
- [13] S. I. Roumeliotis and G. A. Bekey. An extended kalman filter for frequent local and infrequent global sensor data fusion. In *SPIE International Symposium on Intelligent Systems and Advanced Manufacturing*, October 1997.
- [14] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey. Fault detection and identification in a mobile robot using multiple-model estimation. In *Proc. 1998 IEEE International Conference on Robotics and Automation*, pages 2223–2228, May 1998.
- [15] H. W. Sorenson. *Advances in Control Systems*, volume 3, chapter 5, pages 256–260. Academic Press, 1966.
- [16] M. L. Visinsky, I. D. Walker, and J. R. Cavallaro. New dynamic model-based fault detection thresholds for robot manipulators. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 1388–1395, 1994.

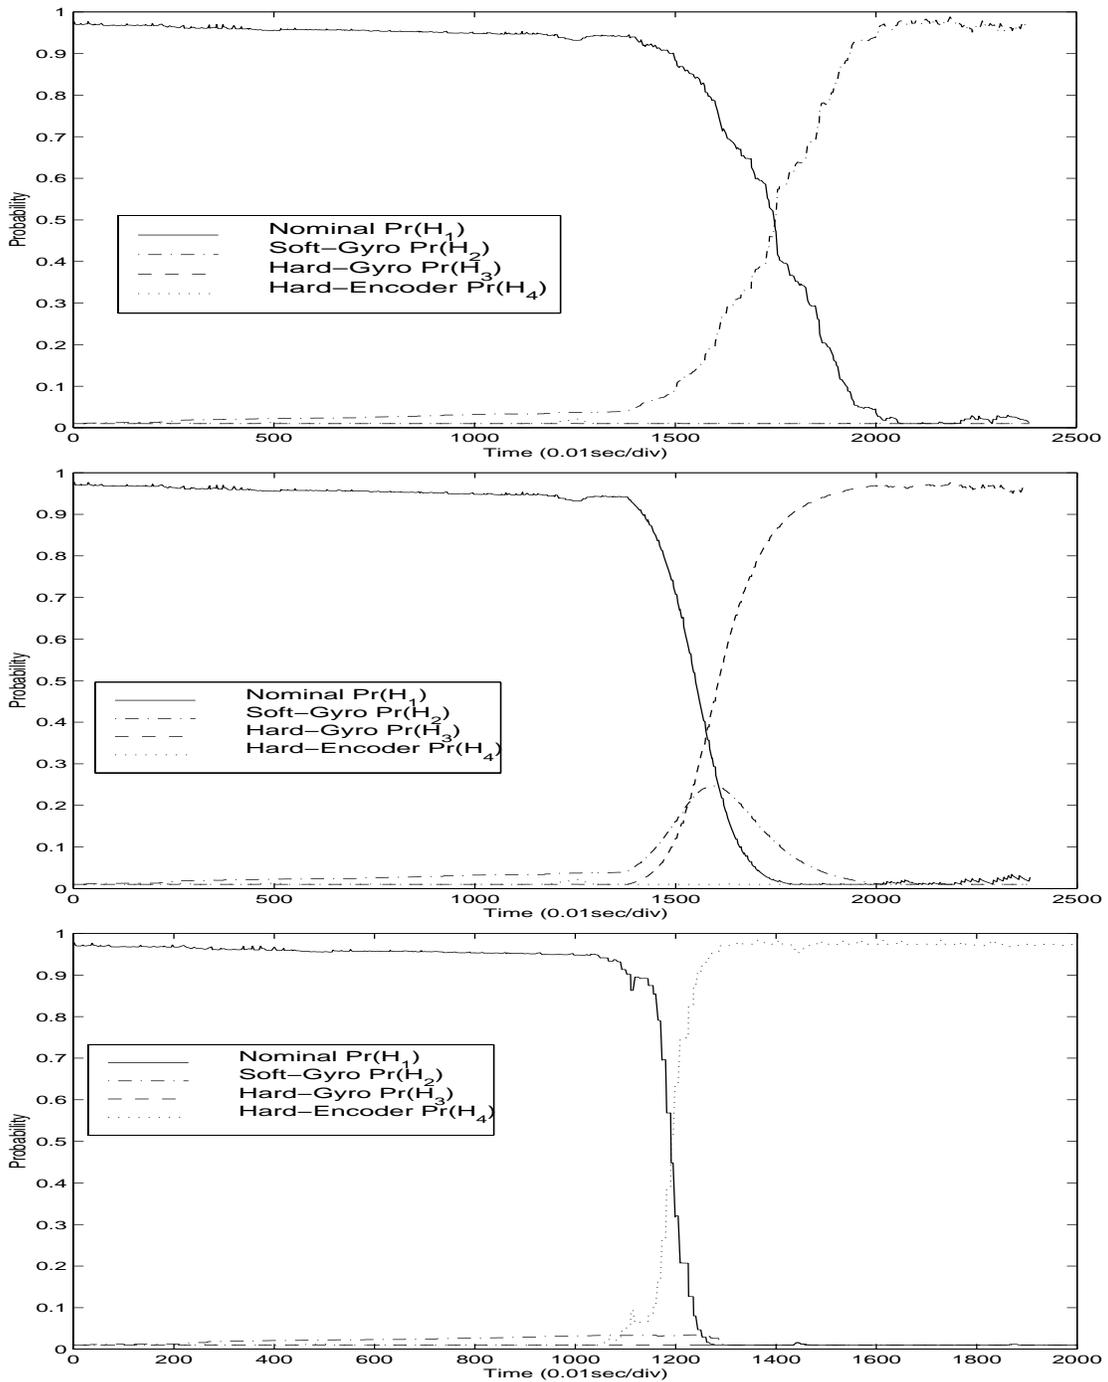


Figure 4: Each figure above shows four traces. In each subfigure the nominal filter conditional probability starts out near 1 and all the others are close to 0. In each case after the failure the nominal probability declines and the appropriate conditional probability rises to nearly 1. The figure at the top shows the case of a soft gyro failure, the middle figure shows the case of a hard gyro failure and the bottom figure shows the hard encoder failure case.