

Fault Detection and Identification in a Mobile Robot using Multiple-Model Estimation*

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

This paper introduces a method to detect and identify faults in wheeled mobile robots. The idea behind the method is to use adaptive estimation to predict (in parallel) the outcome of several faults. Models of the system behavior under each type of fault are embedded in the various parallel estimators (each of which is a Kalman Filter). Each filter is thus tuned to a particular fault. Using its embedded model each filter predicts values for the sensor readings. The residual (the difference between the predicted and actual sensor reading) is an indicator of how well the filter is performing. A fault detection and identification module is responsible for processing the residual to decide which fault has occurred. As an example the method is implemented successfully on a Pioneer I robot. The paper concludes with a discussion of future work.

1 Introduction

Fault¹ tolerant behavior in mobile robots is desirable for a variety of reasons including safety and economics. Fault tolerant behavior refers to the autonomous detection and identification of faults as well as the ability to continue functioning after a fault has occurred. This paper deals with the first two components of the problem namely detection and identification. The ultimate goal of this work is to develop a methodology that allows mobile robots to autonomously detect, identify and rectify faults due to sensor failure, actuator failure as well as mechanical failure.

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

[†]contact author for correspondence

¹In this paper fault and failure are used synonymously

Earlier work in the field (though not applied to mobile robots) is due to [6] and [15]. In [4] a network of adaptive virtual sensors is used to maintain reliable performance of a walking robot with many sensors, actuators and computers. The basic idea is that the virtual sensors reconfigure the way they use sensor information when a failure is detected. In [13] the authors investigate fault-tolerant techniques using redundant sets of control strategies. Recent work in generating fault residuals in robotics includes [14], a grid cell consistency measure based approach and [3], an observer-based approach to fault isolation in robot manipulators.

Kalman filtering [7], [9], [5] is a well known technique for state and parameter estimation. Kalman filtering is a recursive estimation procedure using sequential measurement data sets. Prior knowledge of the state (expressed by the covariance matrix) 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 [8] 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. [11] used the same technique (with an adaptive control strategy) to control F-15 aircraft. Further in [12], [10] the multiple model adaptive estimation (MMAE) technique was used to reliably detect and identify sensor and actuator failure for aircraft.

In recent years Kalman filter based localization has become common practice [2] in the robotics literature. Since the MMAE technique relies upon a bank

of Kalman filters it seems natural to apply it to fault detection and identification in mobile robot systems. The basic philosophy 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 two failures as examples. A future challenge is to develop algorithms that can discriminate between the time profiles of multiple faults as they appear simultaneously in the residual.

The work reported in this paper is largely fault detection and a basic example of fault identification. The overall architecture is a two-component system comprised of 1. fault detection and identification and 2. fault accommodation. The architecture is ultimately intended for larger scale application. In this work four filters were designed and implemented. The filter bank used for fault detection and identification consisted of three of them (the first one was not used). The preliminary filter developed is an adaptive estimator used to estimate the values of two parameters. The basic observation is as follows: in the absence of any fault when the left and right wheels of the robot are commanded to rotate at fixed, equal speeds, the resultant trajectory is not a straight line. This is primarily due to two reasons 1. actuator mismatch and 2. unequal wheel radii. Measurements confirmed that the second factor was not significant. However, modeling actuator mismatch is quite complex. An alternative method is to build a filter with a kinematic model of the robot but to allow the radii of the two wheels to be parameters which can be estimated online. Even though the actual values of the two wheel radii are not (significantly) different the actuator mismatch is effectively modeled using them as parameters. Once the two radii parameters are learned they are used as base values in the other filters. The first filter in the bank uses two measurements from the wheel encoders, a chassis yaw rate and a kinematic model of the robot. This is called the base or nominal filter and it uses the learned radii values as constants in its design making it a good predictor of the nominal behavior of the robot. The second and third filters were each built with an embedded fault model. Filter two in the bank uses a reduced value for the left wheel radius as a model for a flat tire fault on the left wheel. Filter three in the



Figure 1: The Pioneer I

bank uses a reduced value for the right wheel radius as a model for a flat tire fault on the right wheel.

In the experiments performed two faults are examined. The first (referred to henceforth as **right-tire-flat**) is a decrease in the radius of the right tire. The second fault (referred to henceforth as **periodic-bump**) was created by attaching a small object to the left wheel of the robot. A data stream from the robot sensors with these faults absent at first and present subsequently is used to illustrate the response of the three filters. The rest of this paper is organized as follows: first the robot model and the filter residuals are presented. A block level view of the fault tolerant architecture is then discussed. This is followed by a summary of the experimental results, conclusions and a discussion of future work.

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 castor. The kinematics of the Pioneer I are given in Equations 1-2.

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

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

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 . The 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




Figure 2: The Robot Kinematics

with minimal measurement residual. 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. 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 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. The second stage of the proposed work is fault diagnosis. In the current work we are able to diagnose one fault apart from the other rather trivially since one of the filters designed minimizes the residual for a particular fault.

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 equations (with different parameters) is used in both filters. All three filters used in the experiments reported here fall into that category. If the first entry is unequal the two filters have different kinematic equations of the system. Each filter produces a residual \mathbf{r}_{ij} . The nominal model results

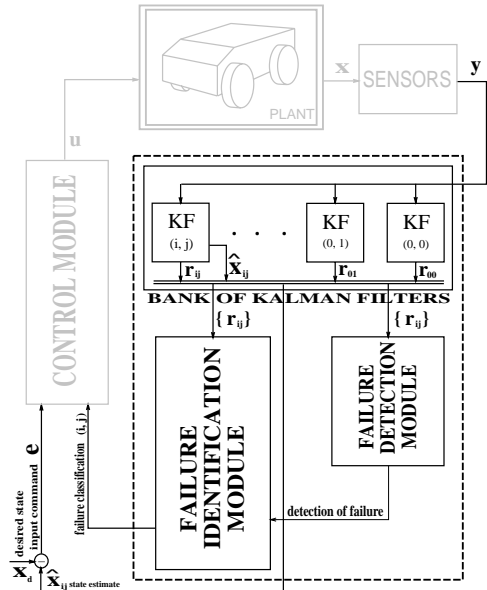


Figure 3: Failure Detection and Identification

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 three 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 correct 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.

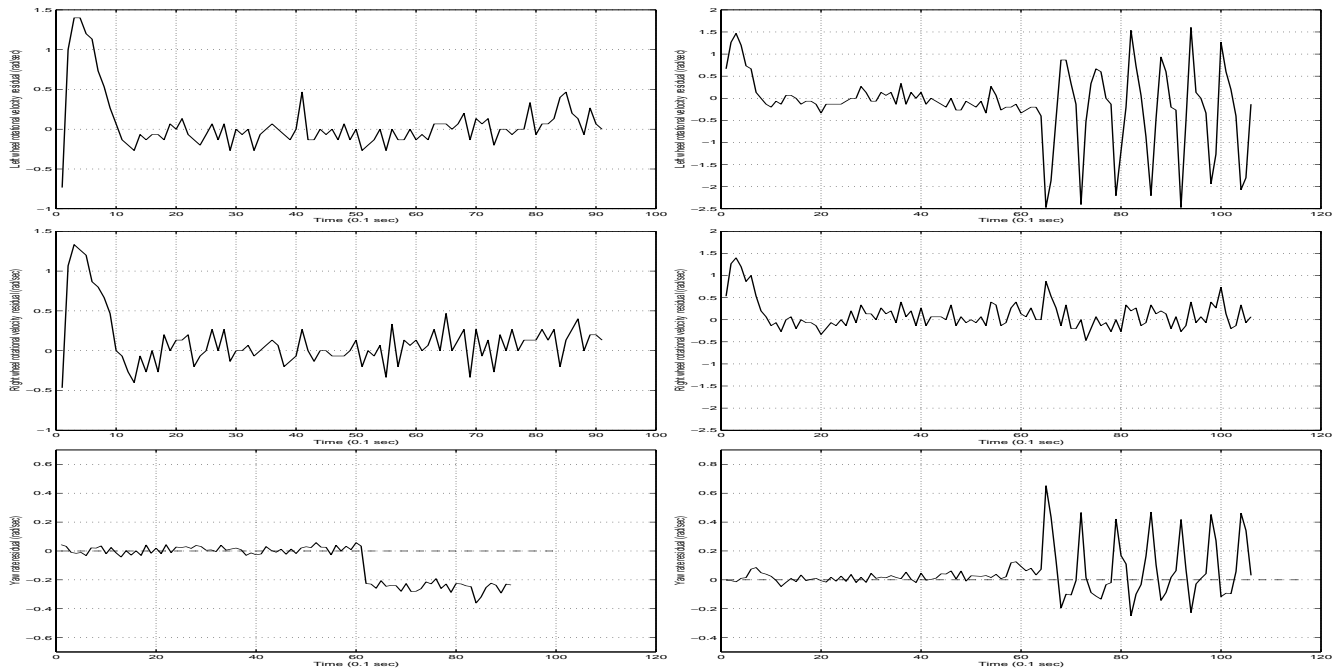


Figure 4: Each column of figures above shows the three **residuals for the nominal filter** (which has no failure model). The left column shows the residuals for the **right-tire-flat** fault. The right column shows the residuals for the **periodic-bump** fault.

4 Experimental Results

In the experiments reported here, two faults were considered and three 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. Each of the three filters used the three component measurement vector given in Equation 3 thus producing a three component residual vector.

Consider Figure 4. These are the residuals from the nominal filter (which has no failure model). The left column shows the residuals for the **right-tire-flat** fault. The right column shows the residuals for the **periodic-bump** fault. In the case of the **right-tire-flat** fault there is no observable change in the first two components of the residual since the (angular) rotational speeds of the wheels do not depend on their radii. However the vehicle does begin to yaw as seen by the yaw rate residual which jumps to a mean value well below zero at $t = 6$ s when the fault occurs. In the right column of Figure 4 one sees the residual response to the **periodic-bump** fault on the left wheel. The rotational velocity of the left wheel as well as the yaw rate residual respond with oscillations of large amplitude.

Consider the graphs shown in Figure 5. As before, the left column shows the residuals for the **right-tire-flat** fault and the right column shows the residuals for the **periodic-bump** fault. The filter is however tuned to the **right-tire-flat** fault and this is immediately seen in the yaw rate residual of the left column - the residual is high before the fault occurs ($t < 6$ s) and drops sharply to zero as soon as the fault occurs signaling the identification of this particular failure. The filter is not tuned to the **periodic-bump** fault but the residual signature clearly shows that the fault occurs at $t = 6.5$ s. The yaw rate residual switches from a steady bias of approximately 0.25 rad/sec to a high amplitude oscillation.

Lastly, consider Figure 6. The residuals in this figure are from the filter tuned to the fault where the left tire goes flat. This fault was **not** considered in the experiments. The filter residuals are again sensitive to both faults as seen in both the left and right columns. However in neither case is the filter able to bring the residual back to zero. This filter would be able to zero out the residual (and thereby identify) for the case where the left tire goes flat.

In the examples given here the failure identification is done very simply using thresholding. Also

the residual signature for the filter tuned to the **right-tire-flat** fault is distinctive. So the problem of detecting which fault has occurred is easily automated. Ideally one would like to design filters for every fault so that after a fault has occurred there is at least one filter whose state estimate is reliable and can be used for control. In this work we relax that condition.

5 Conclusion and Future Work

In this paper a multiple model based technique to detect and identify faults in mobile robotic systems was presented. The technique is based on using a bank of Kalman filters in parallel. Detection and identification of faults is done by analyzing the signature of the residual produced by each filter. In this paper we showed the application of this methodology to the case of a Pioneer I robot. Two faults were considered. The first was a flat tire and the second was an object stuck on the tire. The bank of filters was able to detect each of the two faults as evidenced in the residuals. Further, the flat tire fault is distinguishable by the low residual in the filter tuned to this particular fault. Without doing sophisticated residual analysis we are able to do detection and identification of two faults. It should be noted that the architecture presented here is generalizable to more faults and increasingly sophisticated filters and residual postprocessing. The method can easily be applied to other 3-wheeled mobile systems because it relies on a simple kinematic description.

In the future we plan to concentrate on these issues as well as sensor failure, actuator failure and other mechanical failures. Recovery from failure by altering the control strategy is also the subject of future work. The easiest (and least autonomous) solution is to stop, flag a fault and await human help. Other strategies include making guarded motions, changing the covariance matrices that characterize the sensors or switching to an entirely new control methodology.

Acknowledgments

The authors would like to thank Barry Werger and Dani Goldberg for help with data collection. The authors also thank S. Hayati, G. Rodriguez, R. Volpe, C. Weisbin and B. Wilcox for several useful 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] F. Caccavale and I. Walker. Observer-based fault detection for robot manipulators. In *Proc. 1997 IEEE International Conference on Robotics and Automation*, pages 2881–2887, April 1997.
- [4] C. Ferrell. Failure recognition and fault tolerance of an autonomous robot. *Adaptive Behavior*, 2(4):375–398, 1994.
- [5] M. S. Grewal and A. P. Andrews. *Kalman Filtering, Theory and Practice*. Prentice Hall, 1993.
- [6] 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.
- [7] R. E. Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering*, 86:35–45, 1960.
- [8] D. T. Magill. Optimal adaptive estimation of sampled stochastic processes. *IEEE Transactions on Automatic Control*, AC-10(4):434–439, 1965.
- [9] P. S. Maybeck. *Stochastic Models, Estimation and Control*, volume 1. New York: Academic Press, 1979.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] M. Soika. Grid based fault detection and calibration of sensors on mobile robots. In *Proc. 1997 IEEE International Conference on Robotics and Automation*, pages 2589–2594, April 1997.
- [15] 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 5: Each column of figures above shows the three **residuals for the filter which is tuned for the right tire flat model**. The left column shows the residuals for the **right-tire-flat** fault. The right column shows the residuals for the **periodic-bump** fault.

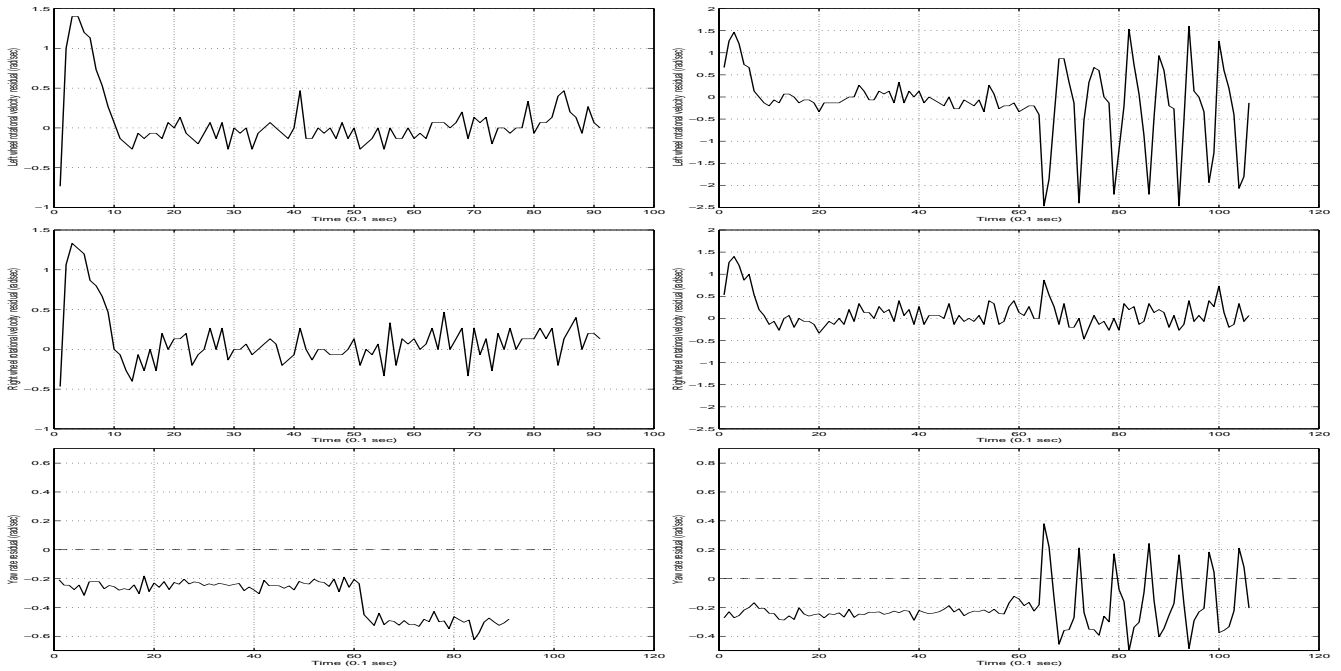


Figure 6: Each column of figures above shows the three **residuals for the filter which is tuned for the left tire flat model**. The left column shows the residuals for the **right-tire-flat** fault. The right column shows the residuals for the **periodic-bump** fault.