

A Vehicle Health Monitoring System Evaluated Experimentally on a Passenger Vehicle

Hok K. Ng, Robert H. Chen and Jason L. Speyer

Abstract—A vehicle health monitoring system based on analytical redundancy is developed for automated passenger vehicles. A residual generator and a residual processor are designed together to detect and identify actuator and sensor faults of the Buick LeSabre rapidly. The residual generator includes fault detection filters and parity equations. It uses the control commands and sensor measurements to generate the residuals, which have a unique static pattern in response to each fault. Then, the residual processor interrogates the residuals by matching them to one of several known patterns. It computes the probability of each hypothesis conditioned on the history of residuals. The vehicle health monitoring system is evaluated in real-time on a Buick LeSabre. The vehicle sensor and actuator faults are simulated artificially by the computer or created manually by the driver. In one experiment, a real intermittent sensor fault occurred and was immediately detected and identified. The real-time evaluation demonstrates that the vehicle health monitoring system can detect and identify actuator and sensor faults under various disturbances and uncertainties with almost minimal detection latency.

Index Terms—Health monitoring, fault detection and identification, analytical redundancy, fault detection filter, sequential probability test

I. INTRODUCTION

A vehicle health monitoring system based on analytical redundancy is developed for automated passenger vehicles. Although we only present results for health monitoring system, our system view of vehicle health monitoring and management is depicted in Figure 1 indicating the importance of minimal detection and identification latency. The vehicle dynamics are driven by the throttle, brake and steering commands, and various unmeasured exogenous influences such as road noise and actuator faults. Sensors measure a possible nonlinear function of the dynamic states and are corrupted by noise, biases and faults of their own. The vehicle health monitoring system generates the probabilities of no fault and probabilities of each fault hypothesis conditioned on the residual sequence. The conditional probabilities are computed using residuals generated by the fault detection filters and the parity equations. These residuals have a unique static pattern with respect to a given fault or no-fault condition. Nominally, the residuals are nonzero only when a fault has occurred and is zero at other times. However, when they are driven by uncertainty and disturbances, the

This work was supported by California Department of Transportation TO4209. We would like to thank Dr. Xiao-Yun Lu and Paul Kretz at PATH (California Partners for Advanced Transit and Highways) for helping us conduct vehicle experiments.

The authors are with the department of Mechanical and Aerospace Engineering, University of California, Los Angeles.

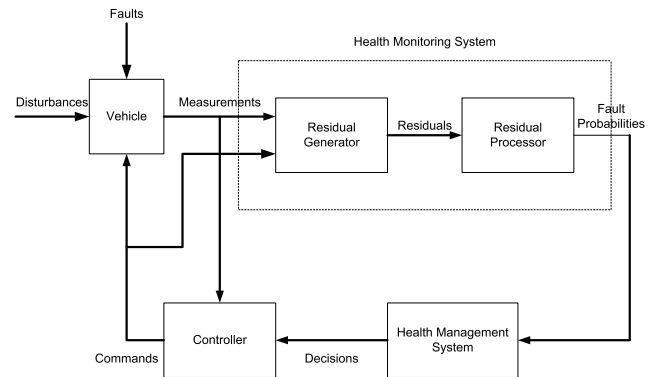


Fig. 1. A system view of vehicle health monitoring and management

residuals fail to go to zero even in absence of a fault. When this occurs, the residual pattern becomes fuzzy. Furthermore, the fault detection filters residuals responds arbitrarily to a fault that is undefined in the process of fault detection filters design. To enhance fault detection and identification, a residual processor interrogates the residuals from the fault detection filters and the parity equations by matching the residuals to one of several known patterns. The pattern matching is done with a probabilistically based algorithm, so the residual processor generates probabilities rather than a simple binary announcement. A simple threshold mapping could be added easily to produce a binary announcement, if that were needed. The conditional probabilities are passed to a vehicle health management system which determines the impact of the possible fault on safe vehicle operation and adjusts control laws if necessary to accommodate a degraded operating condition. The performance of the health monitoring system is only considered here and not the performance of the control loop shown in Figure 1.

The design presented in this paper is used for online health monitoring of an automated Buick LeSabre. This automated vehicle is developed for the Automated Highway Systems (AHS), and provided by the California PATH (Partners for Advanced Transit and Highways). Studies of fault detection related to AHS can be found in [1] among other references. The work of [1] constructs a bank of residuals to identify sensor and actuator fault component using parity equations and several reduced-order nonlinear observers. Their performance is then evaluated in the computer simulation. The experimental results of range sensor fault detection is also

presented. Although the fault detection scheme successfully generates a bank of signals for fault diagnostics, the integration of the residual generation and residual processing to declare a fault has not been addressed. This is particularly important when the residuals are corrupted by uncertainties, and the time to identify a fault is critical. The present paper addresses this issue by presenting a vehicle health monitoring system, which is an integration of a residual generator and a residual processor. Moreover, the health monitoring system is evaluated in real-time in an experimental vehicle having real disturbances and uncertainties.

In Section II, the background of the fault detection filter and residual processor is briefly discussed. In Section III, the vehicle model of the Buick LeSabre is introduced. The vehicle health monitoring system is designed in Section IV. The performance of the vehicle health monitoring system is evaluated in real-time on a Buick LeSabre provided by the California PATH in Section V.

II. FAULT DETECTION AND IDENTIFICATION BACKGROUND

A popular approach to fault detection filter design is the unknown input observer (UIO) [2], [3], [4] which divides the faults into two groups: a single target fault and possibly several nuisance faults. The nuisance faults are placed in an invariant subspace which is unobservable to the residual. Therefore, the residual is only sensitive to the target fault, but not to the nuisance faults. Although only one fault can be monitored in each unknown input observer, there are some benefits. For example, one gains additional flexibility which can be used to improve robustness and time-varying systems can be treated [5], [6], [7]. The design of fault detection filter and parity equation allows constructing a residual pattern which uniquely response to each possible fault. An associated hypothesis can then be defined according to each of the residual pattern.

The essential feature of a residual processor is to analyze the filter residual by matching the residual to one of several hypotheses. By considering the residual processor design as a hypothesis testing problem, the Multiple-hypothesis Shirayev sequential probability test (MHSSPT) [8] can be adopted. This test is a generalized result of Speyer and White [9] who approached the change detection problem based on the results of Shirayev [10], but used a dynamic programming formulation. Under certain condition, this Generalized Shirayev SPRT detects the change in hypothesis (i.e., occurrence of a disruption) in a sequence of conditionally independent measurements in minimal time.

III. VEHICLE MODEL

Since the fault detection filters are designed using linear vehicle models, linear models are derived from a high-fidelity six degree-of-freedom nonlinear vehicle model using a central differences method [11]. Two linear vehicle models are derived at different nominal operating points where the vehicle is travelling straight ahead on a flat road. The linearization is done at 20 *m/s* (45 *mph*) and 24 *m/s* (54

mph) respectively. Although only one linear model is used in each fault detection filter design, linear models derived at different operating points are obtained to gain more insight into the vehicle dynamics and provide more freedom in the design process.

Since the vehicle is travelling straight ahead, the linear longitudinal dynamics decouple completely from the linear lateral dynamics. We use the decomposition to design separately the longitudinal and lateral fault detectors, and apply them even when the vehicle travels on curves. The fast modes of the linear longitudinal model and linear lateral model are removed by a model reduction technique using dynamic truncation with a steady-state correction [12]. The reduced-order longitudinal vehicle model has ten states (manifold air mass, engine speed, longitudinal velocity, vertical position and velocity, pitch angle and rate, sum of front suspension lengths, sum of rear suspension lengths and brake state), two control inputs (throttle and brake) and five measurements (manifold pressure sensor, engine speed sensor, longitudinal accelerometer, sum of front wheel speed sensors and sum of rear wheel speed sensors). The reduced-order lateral vehicle model has six states (lateral velocity, roll angle and rate, yaw rate, difference of front suspension lengths and difference of rear suspension lengths), a steering control and four measurements (lateral accelerometer, yaw rate sensor, difference of front wheel speed sensors and difference of rear wheel speed sensors).

IV. VEHICLE HEALTH MONITORING SYSTEM DESIGN

The vehicle health monitoring system is designed in this section. In Section IV-A, the robustness of fault detection filter is enhanced by considering nonlinearity as a nuisance fault. In Section IV-B, the configuration of fault detection filters are determined. The fault detection filters and parity equations are designed in Section IV-C and Section IV-D, respectively. In Section IV-E, the design of the residual processor is given.

A. Fault Detection Filter Robustness Enhancement

The fault detection filter is designed based on the linear model linearized from the nonlinear model at a single nominal point. Hence, the filter might not be robust when the vehicle is operating far from the nominal point. The robustness of the fault detection filter is enhanced by considering the nonlinearity as a fault and included in the nuisance faults when designing the unknown input observer [7]. The nonlinearity fault direction is determined by using a state estimator and modeling the nonlinearity as a slow time varying one-dimensional or two-dimensional additive term in the linear state equation [4]. Then, the UIO can place the nonlinearity approximately into an invariant subspace that is unobservable to the residual. Therefore, the robustness of the fault detection filter can be enhanced.

The throttle actuator, manifold pressure sensor and nonlinearity faults are not output separable. This indicates that the nonlinearity affects the longitudinal dynamics of the vehicle in a way similar to the throttle actuator and manifold pressure

sensor faults. Therefore, the throttle actuator fault is detected by using parity equations. The manifold pressure sensor fault is detected by forming a nonlinear parity equation based on vehicle dynamics, and can be found in [13].

B. Fault Detection Filter Configuration

Since there are five measurements associate to the vehicle longitudinal mode, the five faults to be detected and identified by the fault detection filters and the nonlinearity fault are not output separable. Therefore, they are grouped into three sets where the faults in each set are output separable. In fault detection filter set no.1, the engine speed sensor fault is paired with longitudinal accelerometer fault and nonlinearity fault. The filter set no.2 includes the sum of front wheel speed sensors fault, sum of rear wheel speed sensors fault and nonlinearity fault. The filter set no.3 includes the brake actuator fault, sum of rear wheel speed sensors fault and nonlinearity fault.

There are five faults to be detected and identified by the fault detection filter for the vehicle lateral mode. Since there are only four measurements associate to the vehicle lateral mode these five faults are not output separable and are grouped into three sets. No nonlinearity fault is required to be included in the nuisance faults when fault detection filters are designed using the linear lateral model. The filter set no.4 includes the steering actuator fault and difference of front wheel speed sensors fault. The filter set no.5 includes the yaw rate sensor fault and difference of rear wheel speed sensors fault. The filter set no.6 includes the lateral accelerometer fault and difference of rear wheel speed sensors fault.

C. Fault Detection Filter Design

Fault detection filters are designed for the six sets of faults determined in Section IV-B using the design algorithm in [7]. Two UIOs are designed for each set of faults. For example, when designing the first UIO for fault detection filter set no.1, the target fault is the engine speed sensor fault and the nuisance faults are nonlinearity and longitudinal accelerometer fault. For the second UIO, the target fault is the longitudinal accelerometer fault and the nuisance faults are nonlinearity and engine speed sensor fault. The UIOs for the other five sets of faults are designed similarly.

The UIO places the nuisance fault into its detection space which is unobservable to the residual and the eigenvalues associated with the unobservable subspace go to $-\infty$ in the limit. When it is not in the limit, the nuisance fault is partially blocked and the UIO has some fast eigenvalues associated with some weakly observable states. These weakly observable states approximate the detection space of the nuisance fault. Therefore, model reduction is needed to reduce the order of the UIO when it is not in the limit by truncating the weakly observable states or using balance realization [14].

D. Parity Equation Design

Parity equations are derived to detect the throttle actuator, throttle sensor, brake actuator and brake sensor faults. There

Residual \ Fault	Fault Detection filter #1		Fault Detection Filter #2		Fault Detection Filter #3		Parity Equation #1	Parity Equation #2
	Engine Speed	Long. Accelerometer	Front Wheel Speed	Rear Wheel Speed	Brake Actuator	Rear wheel Speed		
Engine Speed		—				—	—	—
Long. Accelerometer	—		—	—	—	—	—	—
Front Wheel Sensor		—		—			—	—
Rear Wheel Sensor		—	—		—		—	—
Throttle	—	—	—	—	—	—		—
Brake Actuator	—	—				—	—	
Brake Sensor	—	—	—	—	—	—	—	

Fig. 2. Residual patterns of the longitudinal vehicle generator

are two parity equations derived for the Buick LeSabre. The first parity equation is a function of the throttle command and throttle measurement. The residual is zero when there is no throttle actuator or throttle sensor fault. The residual becomes nonzero when any of these two faults occurs. Therefore, this parity equation can detect the throttle actuator and throttle sensor faults, but cannot identify these two faults. The second parity equation is a function of the brake command and brake measurement. This parity equation can detect the brake actuator and brake sensor faults, but cannot identify these two faults. However, by combining the parity equations with the approximate unknown input observers in Section IV-C, unique residual patterns are presented allowing each fault to be identified.

E. Residual Processor Design

The residual processor is designed to compute conditional probability of each possible failure disruptions based on the residual pattern constructed by the residual generator. Since the longitudinal and lateral vehicles dynamics are decoupled, and the residual generator is designed for the vehicle longitudinal and lateral mode separately, the design of the residual processor is also divided into two parts respectively. The longitudinal residual processor processes the residuals generated by the fault detection filters and parity equations designed for the longitudinal vehicle dynamics. The lateral residual processor responses to the residuals of fault detection filters designed for the lateral vehicle dynamics. Furthermore, the norm of the residual of each filter and parity equation is used to reduce the complexity of the problem. Dividing the residual processor into two parts and using norm of residual can reduce the dimension of the input vector and thus, shorten the computation time. Another advantage of using just the norm is that only one hypothesis is enough when defining a failure hypothesis. Without the norm operating on the residual, two hypotheses are needed to define one fault. One hypothesis is associated to positive fault scenario and the another hypothesis is associated to the negative fault scenario.

There are fifteen hypotheses in total defined when designing the longitudinal residual processor. The null hypothesis is defined for the case before any failure disruption. A pair

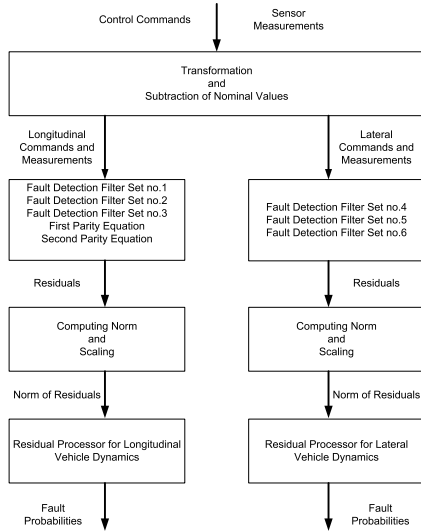


Fig. 3. Fault Probabilities generated by using vehicle data

of hypotheses is defined for each failure in throttle (actuator or sensor), brake actuator, brake sensor, engine speed sensor, longitudinal accelerometer, sum of front wheel speed sensors and sum of rear wheel speed sensors, respectively. For each pair of hypotheses, one is defined for a small failure magnitude and the another is defined for a larger failure magnitude. The posterior probability of each failure is the sum of the probability of the two hypotheses. The lateral residual processor is designed similarly and computes the conditional probability of no failure disruption and failure in steering actuator, lateral accelerometer, yaw rate sensor and difference of front wheel speed sensors and difference of rear wheel speed sensors, respectively. When defining the conditional probability density function for each hypothesis, it is assumed that the input residual vector generated at each time instant is Gaussian distributed. Note that other probability density functions can be chosen and the derivation of the Shirayev test still remains valid. Each conditional probability density function is determined based on its associated residual pattern. Figure 2 shows the residual patterns associated with the longitudinal residual generator. Note that, each row of Figure 2 is a pattern of the input residual vector under a particular scenario. The magnitudes of the residuals depend on the size of the fault, which is not known a priori. Thus, the mean of the input residual vector is assumed to be an uniform random vector, and is uniformly distributed in a range determined by simulating the fault with a variety of magnitudes. The corresponding magnitudes and variances of the residuals are recorded in the computer simulation. The range of the mean and the covariance of the conditional probability density function can be determined accordingly.

V. VEHICLE HEALTH MONITORING SYSTEM EVALUATION

In this section, the performance of the vehicle health monitoring system, which includes residual generator and

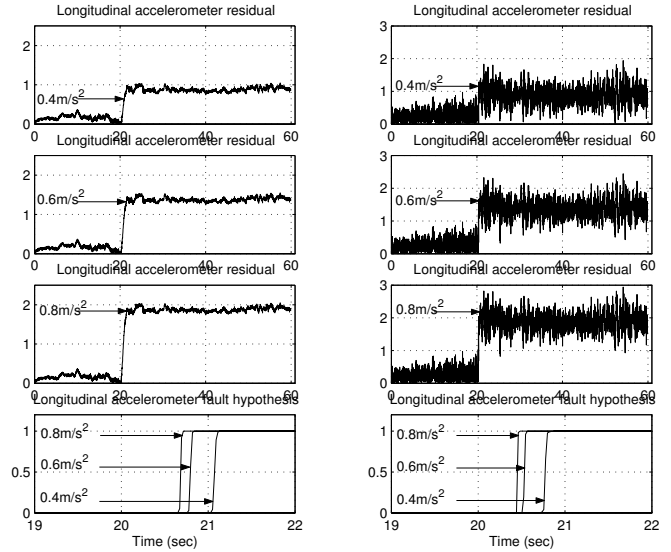


Fig. 4. Residuals and probabilities when longitudinal accelerometer fault occurs

residual processor, is evaluated in real-time on a Buick LeSabre. Since the simulation does not match the empirical data on some control commands and measurements, a transformation is derived to approximately match the simulation and empirical data. Then, the transformation is applied to the control commands and measurements of the empirical data before the fault detection filters use these data to generate the residuals. The fault detection filters are designed using linearized vehicle dynamics derived from the nonlinear vehicle dynamics at certain nominal operating point. Thus, the nominal values of the control commands and measurements have to be subtracted from the vehicle generated data before the fault detection filters can use these data to generate the residuals. Furthermore, a normalization scheme is used for the presentation of the performance of the fault detection filters. Since the residuals are zero vectors when there is no fault and nonzero vectors when their associated faults occur, the performance of the fault detection filters is examined by checking the norms of the residuals. When evaluating the fault detection filters, the norms of the residuals are scaled to one when their associated faults of certain magnitudes occur. The purpose of the scaling is to present the performance of the fault detection filters in a clearer fashion, i.e., zero residuals represent no fault and residuals of magnitude one represent the occurrence of their associated faults. Note that the residuals can be scaled with respect to any other fault magnitudes if that were desired. The steps of generating posterior probabilities of each possible fault using vehicle data during the evaluation of the vehicle health monitoring system is summarized in the block diagram shown in Figure 3.

The vehicle health monitoring system is evaluated off-line using empirical data recorded at Crow's Landing, California. The first column of Figure 4 shows the longitudinal accelerometer residuals and the probabilities of longitudi-

nal accelerometer fault when a bias with a magnitude of $0.4m/s^2$, $0.6m/s^2$ or $0.8m/s^2$ is added to the longitudinal accelerometer, respectively. It can be seen that the size of the residuals is proportional to the magnitude of the fault. The delay in detecting the longitudinal accelerometer fault is smaller when the magnitude of the fault is bigger. To shorten the fault detection time, the longitudinal accelerometer fault detection filter is re-designed by tuning the design weightings. The residual responses to the fault is faster at the expense of being more sensitive to the noise. The residual processor is also tuned by using a larger variances for the longitudinal accelerometer residual. The results are shown in the second column of Figure 4. The delay in detecting the longitudinal accelerometer fault is reduced even though the residual is much noisier. Without the residual processor, it would be difficult to detect the fault in the noisy residual. Therefore, designing the residual generator and residual processor together can enhance the fault detection time. Note that the covariance of each conditional probability density function is larger, when the residuals are noisier. Thus, a fault with smaller magnitude can be announced faster, but with a possible larger false alarm rate. The integration of residual generator and residual processor allows a design tradeoff between fault detection time and minimum detectable fault magnitude. Furthermore, it enhances the decision making in the presence of uncertainties, because the decision of announcing a fault can be made by considering the rate of false alarm and miss detection.

In Section V-A, the experiment setup (i.e., the real-time testing environment) is discussed. In Section V-B, the health monitoring system is evaluated in real-time on a Buick LeSabre under different scenarios.

A. Experiment Setup

The Buick LeSabre used for experiment is provided by PATH. The vehicle can be operated under automatic control by a computer. The health monitoring system is written in C code and executed in real-time on the vehicle. When evaluating the performance of the vehicle health monitoring system, each sensor faults is simulated by superimposing a fault onto the measurement, i.e., the simulated faulty measurement is the sum of the correct measurement and the sensor fault. The real brake actuator fault is imposed by manually stepping on the brake pedal when the vehicle is under automatic control. Similarly, the real steering actuator fault is imposed by manually turning the steering wheel. The magnitudes of the two actuator faults are not known in these cases because the additional brake pressure and steering angle are not measured. Since this additional brake pressure or steering angle is added only to the vehicle but not to the control commands used by the health monitoring system, this creates a real brake actuator or steering actuator fault for the evaluation of the vehicle health monitoring system.

B. Real Time Evaluation

The vehicle health monitoring system is evaluated in real-time on the Buick LeSabre at Crow's Landing, California.

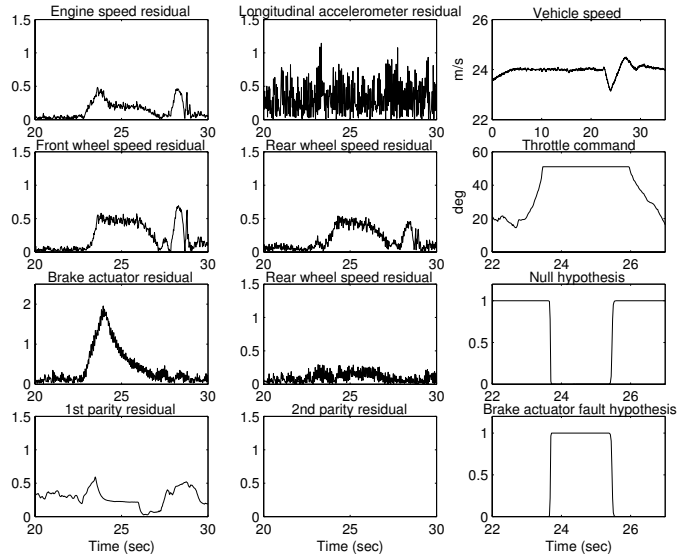


Fig. 5. Residuals and probabilities when brake actuator fault occurs

In each trial, the vehicle is first driven to $20m/s$ before entering the curve track which is built on a slight slope. Then, the evaluation is performed when the vehicle travels on the curve track in different constant and variable speeds. Due to the page limit, the performance of the vehicle health monitoring system is shown here when a brake actuator fault, steering actuator fault, and a real intermittent front wheel speed sensor fault occurs. Please see [13] for the complete evaluation of the monitoring system when other fault occurs.

The brake actuator fault is created after the vehicle speed reaches $24 m/s$ on the curve track. The driver created the brake actuator fault after $23^{rd}s$ by stepping on the brake pedal manually, when the vehicle is still under automatic control. The first two columns of Figure 5 show the norms of the residuals generated by the longitudinal residual generator when the brake actuator fault occurs. Since the brake command is associated only with the vehicle longitudinal dynamics, the results of the residual generator and processor designed for the vehicle lateral dynamics are not shown because they do not respond to the fault. Through the plot of the throttle command in the third column of Figure 5, the vehicle controller tried to compensate this fault by increasing the throttle angle to keep the desired vehicle speed. The probabilities of brake actuator fault go to one in about 0.5 second after the fault is created. The probabilities of other faults are zero and not shown here.

The steering actuator fault is created manually when the vehicle travels on the curve track at $20m/s$. The driver introduces the steering actuator fault after $3^{rd}s$ by turning the steering wheel manually when the vehicle is still under automatic control. The vehicle is deviated from the original maneuver by this interruption. The first two columns of Figure 6 show the norms of the residuals generated by the lateral residual generator when the steering actuator fault occurs. The results of the longitudinal residual generator and processor are not shown because they do not respond

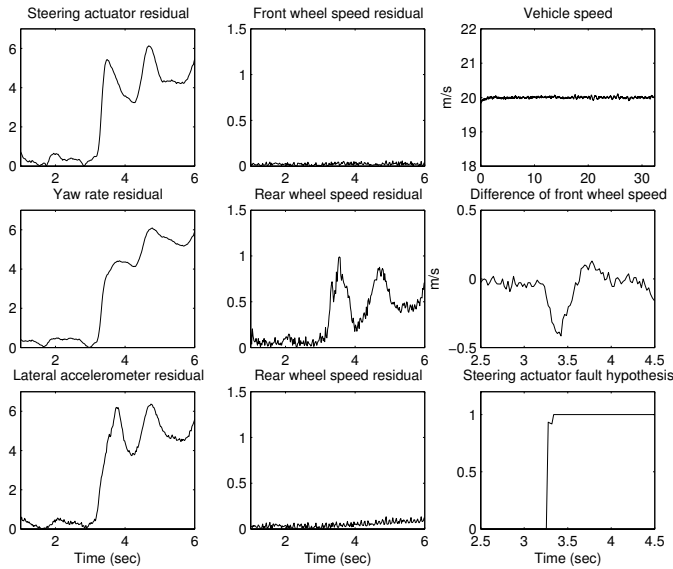


Fig. 6. Residuals and probabilities when steering actuator fault occurs

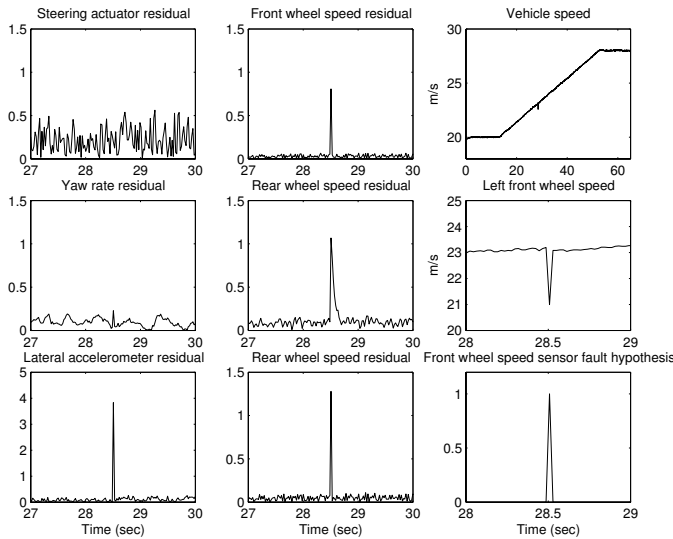


Fig. 7. Residuals and probabilities when real intermittent wheel speed sensor fault occurs

to the fault. The third column of Figure 6 shows the vehicle speed and the difference of the front wheel speed sensors measurements which indicates the presence of the steering actuator fault. The probabilities of steering actuator fault jump to one almost immediately after the fault is created.

In addition to the intended experiment, an unexpected real intermittent fault accidentally occurs in the left front wheel speed sensor when we attempted to evaluate the lateral residual generator and processor in response to the steering actuator fault. The first two columns of Figure 7 show the norms of the residuals generated by the lateral residual generator when the wheel speed sensor fault occurs. It is shown in the third column that an intermittent fault is present after 28ths. The probabilities of front wheel speed sensor fault jump from zero to one immediately when the fault

occurs and then drop back to zero when the fault disappears.

VI. CONCLUSION

A vehicle health monitoring system is developed by integrating the design of fault detection filters, parity equations and residual processor together. It is evaluated in real-time on a PATH Buick LeSabre at Crow's Landing, California. The real-time evaluation shows that the health monitoring system can detect and identify most of the vehicle actuator and sensor faults almost immediately after the fault occurs. The fault detection filters are shown to have good performance in arbitrary vehicle motion even though they are designed using a linear vehicle model associated only with the longitudinal or lateral vehicle dynamics. The residual processor correctly matches the residuals to one of the hypothesized residual pattern in the presence of real disturbances. The design as presented is intended to be packaged as a module to be used by the vehicle health management system under development by the California PATH.

REFERENCES

- [1] R. Rajamani, A. S. Howell, C. Chen, J. K. Hedrick, and M. Tomizuka, "A complete fault diagnostic system for automated vehicles operating in a platoon," *IEEE Transactions on Control System Technology*, vol. 9, no. 4, pp. 553–564, July 2001.
- [2] M.-A. Massoumnia, G. C. Verghese, and A. S. Willsky, "Failure detection and identification," *IEEE Transactions on Automatic Control*, vol. AC-34, no. 3, pp. 316–321, Mar. 1989.
- [3] P. M. Frank, "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy - a survey and some new results," *Automatica*, vol. 26, no. 3, pp. 459–474, May 1990.
- [4] R. J. Patton and J. Chen, "Robust fault detection of jet engine sensor systems using eigenstructure assignment," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 15, no. 6, pp. 1491–1497, Nov-Dec 1992.
- [5] W. H. Chung and J. L. Speyer, "A game theoretic fault detection filter," *IEEE Transactions on Automatic Control*, vol. AC-43, no. 2, pp. 143–161, Feb. 1998.
- [6] R. H. Chen and J. L. Speyer, "A generalized least-squares fault detection filter," *International Journal of Adaptive Control and Signal Processing - Special Issue: Fault Detection and Isolation*, vol. 14, no. 7, pp. 747–757, Nov. 2000.
- [7] R. H. Chen, D. L. Mingori, and J. L. Speyer, "Optimal stochastic fault detection filter," *Automatica*, vol. 39, pp. 377–390, 2003.
- [8] D. P. Malladi and J. L. Speyer, "A generalized shiryayev sequential probability ratio test for change detection and isolation," *IEEE Transactions on Automatic Control*, vol. AC-44, no. 8, pp. 1522–1534, Aug. 1999.
- [9] J. L. Speyer and J. E. White, "Shiryayev sequential probability ratio test for redundancy management," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 7, no. 5, pp. 588–595, Sep-Oct 1984.
- [10] A. N. Shiryayev, *Optimal Stopping Rules*. Springer-Verlag, 1977.
- [11] R. K. Douglas, J. L. Speyer, D. L. Mingori, R. H. Chen, D. P. Malladi, and W. H. Chung, *Fault Detection and Identification with Application to Advanced Vehicle Control Systems*. California PATH Research Report UCB-ITS-PRR-96-25, 1996.
- [12] R. Prakash, "Properties of a low-frequency approximation balancing method of model reduction," *IEEE Transactions on Automatic Control*, vol. AC-39, no. 5, pp. 1135–1141, May 1994.
- [13] R. H. Chen, H. K. Ng, J. L. Speyer, and D. L. Mingori, *Testing and Evaluation of Robust Fault Detection and Identification for a Fault Tolerant Automated Highway System*. California PATH Research Report UCB-ITS-PRR-2004-46, 2004.
- [14] B. C. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Transactions on Automatic Control*, vol. AC-26, no. 1, pp. 17–32, Feb. 1981.