A Probabilistic Vehicle Diagnostic System Using Multiple Models

Matthew L. Schwall
mschwall@stanford.edu

J. Christian Gerdes
gerdes@cdr.stanford.edu

Design Division
Department of Mechanical Engineering
Stanford University
Stanford, California 94305-4021

Bernard Bäker
bernard.baeker@daimlerchrysler.com

Thomas Forchert
thomas.forchert@daimlerchrysler.com

DaimlerChrysler AG
Research and Technology
Mechatronic and Diagnosis, REM/SD
Stuttgart, Germany

Abstract

In addition to being accurate, it is important that diagnostic systems for use in automobiles also have low development and hardware costs. Model-based methods have shown promise at reducing hardware costs since they use analytical redundancy to reduce physical redundancy. In addition to requiring no extra sensors, the diagnostic system presented in this paper also allows for high accuracy and low development costs by using information from multiple simple models. This is made possible by the use of a Bayesian network to process model residuals. A hybrid, dynamic Bayesian network is used to model the temporal behavior of the faults and determine fault probabilities. A prototype of the system has been implemented and tested on a Mercedes-Benz E320 sedan. This paper describes the prototype system and presents results demonstrating the system's advantages over traditional residual threshold techniques.

Introduction

The goal of on-board vehicle fault detection and isolation (FDI) is to identify faults before they damage the vehicle or create a dangerous situation for the occupants. In order to be practical to implement, such systems should require minimal additional hardware, such as sensors and computational power, and also have low complexity and development cost.

Model-Based Techniques

Model-based methods are the preferred means of achieving these goals. These methods use analytical redundancy in order to reduce the need for costly physical redundancy, such as extra sensors, and have been shown to be successful in a wide variety of applications (Isermann & Ballé 1996). As shown in Fig. 1, model-based techniques consist of two stages. In the first stage, measurements of system variables and parameters are compared with values predicted by dynamic models in order to generate a residual. This residual is then analyzed in the second decision-making stage to determine if a fault has occurred.

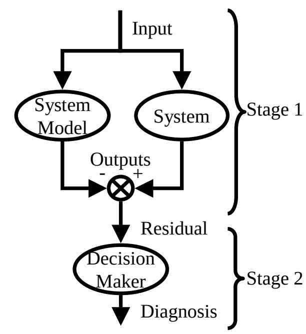

Figure 1: Generalized structure of model-based diagnostics

Multiple Model Diagnostics

For complex systems such as a vehicle, it is often impractical to represent the entire system with a single model. Instead, separate models are used for different subsystems and dissimilar observable modes such as temporal and frequency-based signals.

In theory, there are many benefits in using a multiple model approach. First, the many representations provide an easy means for incorporating system information provided from a variety of sources. Another possible advantage is that design models can be used directly and need not be adapted to different ontologies or assumption sets. Finally, integrating several forms of information yields a more complete diagnostic system that can take advantage of relationships between different modeling domains. The challenge is to develop a framework for integrating the models that exploits all of these advantages.

A challenge common to both single and multiple model techniques is how to distinguish between a fault and noise on the residuals. Many model-based diagnostic systems simply compare the residuals against a threshold to determine if a fault has occurred. This requires carefully designed, well-behaved residuals, which often requires extensive modeling and high computational effort. This can translate into unacceptable engineering and on-board computer costs. Furthermore, most techniques lack a structure for setting thresholds.

The system described in this paper uses a Bayesian network (BN) for residual processing. Because of the use of this AI technique, the system is capable of using less robust residuals. This lowers the costs associated with residual generation and reduces the total effort required to create a diagnostic system. In addition to enabling the integration of multiple models, this method permits the use of multiple modes of reasoning, and has therefore been referred to as Multi-Modal Diagnostics (MMD). The next section discusses the probabilistic Bayesian network framework. The following two sections introduce the vehicle system and describe the implementation of the diagnostic system. Experimental results are then presented and discussed.

Probabilistic Residual Processing

Inherent in all diagnostic methods is some uncertainty about the true state of the system, due to sensor noise and unmodeled dynamics. For this reason, all techniques must have some sort of threshold that determines the boundary between accepting a signal as the result of noise, and declaring a fault. The methodology presented here differs from previous techniques in that thresholds are not set on residuals, but on the probability that a fault has occurred. This threshold is an intuitive quantity that can be easily understood in terms of risks and costs. In addition, as shown in Schwall and Gerdes (2001), optimal thresholds can be determined given the cost of a false alarm and the additional costs associated with missing a fault.

Static Bayesian Network Representation

Bayesian (belief) networks are often used for diagnostics, and the technique presented here is not the first to apply them to diagnosing a dynamic system. However, previous methods, such as that presented by Lerner et al., model the system directly with a Bayesian network. This requires a complicated transformation for each model and results in a highly complex network that poses challenging computational problems.

Figure 2 shows how a BN can be used for processing multiple residuals. In this paper, observable nodes, such as the model residuals, are shown in shaded boxes. This BN relates the state of three faults to the probability distribution for two residuals. For example, the probability distribution for Residual 1 depends on the states of its parents, Fault 1 and Fault 2. The probability of either an individual fault, or the joint probability of the faults can be inferred given the state of the residuals.

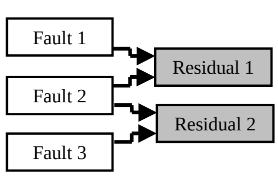

Figure 2: An example diagnostic Bayesian network

Faults can be assumed to have two states—fault and no-fault. Residuals, however, are continuous values, as they represent the difference between a measured and an estimated scalar value. Therefore, the fault nodes in the network are discrete, the observed residual nodes are continuous, and the resulting network is a hybrid Bayesian network. For a single residual R with a set of parent faults $\boldsymbol{F}$, $P(r \mid \boldsymbol{F})$ is represented by a continuous conditional probability distribution (CPD). It is convenient to represent $P(r \mid \boldsymbol{F})$ by separate probability density functions $P(r \mid \boldsymbol{f})$ for every possible state $\boldsymbol{f}$ of the parent faults $\boldsymbol{F}$. The Gaussian distribution is commonly used both because it is valid for a normally distributed noise assumption and also simplifies calculations in certain types of networks. However, for the diagnostic network presented here, there is little computational incentive to use one form of distribution over another, because residuals are the only continuous nodes, and they are observed.

Representation of Time Dependence

The appearance and persistence of faults is a temporal process. In order to represent this dependence, a dynamic Bayesian network (DBN) is used. Two significant assumptions must be made in order to compactly represent the DBN by a 2-time-slice Bayesian network (2-TBN).

The first is that the system is time-invariant. For our system this implies that the probability of a fault given the previous state of the fault, $P(f(t+1) \mid f(t))$, is constant for all t. For components that wear, this probability could slowly change over the life of the vehicle. Fortunately, such a process is slow enough that it could be updated each time the car is started, and does not impair our ability to use a DBN.

The second assumption is the Markov assumption, which states that the future is conditionally independent of the past given the present. This is reasonable when considering faults and implies that, given the status of the faults at time t, the probability of faults at time $t+1$ does not depend on faults prior to t.

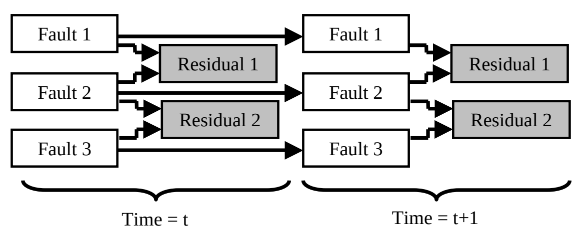

Figure 3: An example dynamic Bayesian network

The 2-TBN corresponding to adjacent time slices of the BN shown in Fig. 2 is shown in Fig. 3. The lack of edges from the residuals at t to the residuals at $t+1$ implies that the residuals at $t+1$ are conditionally independent of the residuals at t given the state of the faults at time t or $t+1$. This means that the only correlation between residuals at different time steps is through the mechanisms of the faults.

In order to assess whether or not this is a valid assumption, it is necessary to analyze the mechanisms that cause residuals. As was shown in Fig. 1, residuals exist because the model predicts a different value than what a sensor indicates. Expanding on this representation with the use of a Box-Jenkins model structure (Box, Jenkins, & Reinsel 1994), Fig. 4, shows that there are two possible sources of a

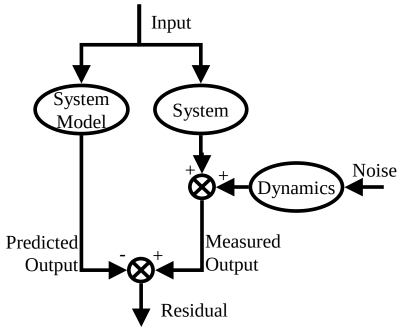

Figure 4: Residual generating system structure

residual. One is due to differences between the system and the system model, and the other is due to noise.

Errors in the system model can either be undesirable dynamics and parameter errors, which are the faults, or dynamics that are a part of the functioning system but are not modeled due to their complexity. The components of a residual due to undesirable errors are correlated through the mechanism of the fault and are therefore conditionally independent given the fault state. Residuals due to other unmodeled dynamics and parametric errors in the system model will be correlated with the system input. These correlations must be accounted for in order to maintain the accuracy of the solution. As will be seen below, they can be dealt with by using an extra node called a status parameter provided the type of input that excites the unmodeled dynamics is known.

The other contribution to residuals is noise. Noise is introduced into the system as an independent random input. However, after this input has been shaped by the noise dynamics, its contribution to the residual is autocorrelated. Thus, in order for our assumption to be valid, the BN must sample the residuals slower than the noise dynamics.

Prototype System

A prototype of the diagnostic system has been implemented on a Mercedes-Benz E320 sedan in order to diagnose the car's handling system. As the focus of this paper is the implementation of the AI system, the models are only briefly described. For more detail regarding the models, see (Schwall & Gerdes 2001). The section after this will describe the hardware and software implementation.

The input to the diagnostic system comes from six sensors located on the vehicle: four wheel speed sensors, a yaw rate gyro and a steering angle sensor. For the purpose of this example, ten faults will be considered. The faults consist of a failure of any of the four tires, as well as faults of any of the six sensors. In order to diagnose the system, three simple models are used; the bicycle model of vehicle handling, a model of yaw rate given left and right wheel speeds, and a model for predicting longitudinal slip of the driven wheels.

Vehicle Models

The classic bicycle model of vehicle cornering consists of a pair of dynamic equations relating steering angle to yaw rate and sideslip angle. The model assumes a constant longitudinal velocity.

By comparing the yaw rate predicted by the bicycle model with the value measured with the gyro, a residual can be generated which will indicate a fault in one of the parameters used in the model.

$$R_1 = r_{expected(bicycle\ model)} - r_{measured(gyro)} \quad (1)$$

When turning, the outside wheels of a car rotate faster than the inside wheels. The second model compares the left and right wheel speeds in order to estimated the yaw rate of the vehicle (Frank, Palkovics, & Gianone 2000).

Estimates of the yaw rate can be computed using data from both the front and rear wheel pairs and compared with the yaw rate measured by the yaw rate gyro. Thus:

$$R_2 = r_{expected(front\ wheels)} - r_{measured(gyro)} \quad (2)$$

$$R_3 = r_{expected(rear\ wheels)} - r_{measured(gyro)} \quad (3)$$

Wheel slip is defined by the relative difference in velocity of the vehicle and the wheel's rotational rate. For a two-wheel drive vehicle that is not braking, the undriven wheels have zero slip. Assuming that both wheels on the same side of the vehicle have the same effective radius, the wheel speeds of the driven and undriven wheels are sufficient to calculate the actual wheel slip

Carlson and Gerdes (2002) demonstrated that a simple longitudinal model with linear tire stiffness can reasonably approximate the slip arising from low levels of acceleration. Therefore, given the car's mass and an experimentally determined longitudinal tire stiffness, slip can be estimated from the vehicle's acceleration, which is calculated by differentiating the undriven wheel speeds.

The actual and estimated slip for each of the driven wheels can be compared to yield one residual for each of the right and left wheel pairs.

$$R_4 = S_{expected(acceleration)} - S_{measured(right\ wheels)} \quad (4)$$

$$R_5 = S_{expected(acceleration)} - S_{measured(left\ wheels)} \quad (5)$$

Since all four wheels slip during braking, these residuals are only valid when the vehicle is not braking. In the future, the model could be extended to cover braking if desired.

Bayesian Network Structure

Figure 5 shows a single time slice of the DAG that results from drawing the causal relationships implied by the models. Here, the five residuals are shown in shaded solid boxes. In addition to the residuals and faults, the graph contains discrete "hidden" nodes shown in dashed boxes. Some of these nodes, such as $C_{\alpha f}$ and $C_{\alpha r}$ representing the front and rear cornering sicknesses, are included in order to make the BN easier to generate and understand. They are used when multiple faults effect a residual through the same mechanism or parameter.

The observed (shaded) "hidden" nodes such as *Not Braking* and *Constant Velocity* are status parameter nodes that are necessary to represent the assumptions used when generating the residuals. They function because of the structure of the CPDs of their child residuals. The CPDs indicate that in the presence of a vehicle status parameter, there may be a high likelihood of an abnormal residual even in the absence of any faults. This provides a means to "explain away" a large residual by the presence of a status parameter.

The DBN is based on this single-slice BN with interconnections between each fault at time t and the same fault at time $t+1$. Therefore, in order to solve the DBN exactly, we need to solve for and maintain the joint probabilities of all possible fault combinations. For the 10 binary faults, this distribution contains 2^{10} probabilities. This distribution was calculated and maintained at each time step for the system described in this paper, and is acceptable for diagnostic systems of similar scope given modern computing power. However, for larger systems such as might eventually be used on production vehicles, approximate solution methods would be valuable and are an area of active research.

Conditional Probability Distributions

As seen in Schwall and Gerdes (2002) , the noise on residuals generated with automotive dynamic models does not perfectly fit a normal distribution, though the approximation is quite close. For the prototype system, the residual distributions were approximated by normal distributions, both in the fault and no-fault cases.

Estimating the mean μ and variance σ^2 of the residual distributions in the no-fault case is straightforward, since this data is readily available. These parameters were estimated by taking data with the fully functioning vehicle and are shown in Table 1. The residuals' variances are high due to the simplicity of the models and demonstrate the potential of the method when using low-cost models.

Table 1: No-fault residual distribution parameters

Residual	$\boldsymbol{\mu}$	$\boldsymbol{\sigma}$
Bicycle Model (rad/s)	0.005	0.020
Front Wheel Speed Yaw Rate (rad/s)	-.014	0.015
Rear Wheel Speed Yaw Rate (rad/s)	-.008	0.018
Right Wheel Slip	0.0	0.004
Left Wheel Slip	0.0	0.004

Data was not taken in the presence of faults; rather, the faults evaluated were simulated. The residual distributions in the presence of faults were estimated as having the same mean as the no-fault distribution and larger variances. In this case, the standard deviations were estimated as 10 times that of the no fault cases. The variance for the residual with multiple faults was made the same as that with one fault. This was done so that exceptionally large residuals would not necessarily suggest multiple faults.

The *a priori* fault probabilities are based on very rough approximations of vehicle component life. They correspond approximately to one failure of each tire every 30,000 miles and one failure every 100,000 miles for the sensor faults.

System Implementation

The research prototype system has been implemented in the vehicle in a manner driven by the desire for easy development, frequent changes, and future upgrades. For this reason, the hardware and software are considerably different than what would be used in a production application, where cost would be a driving factor and economies of scale would allow for custom hardware and optimized software.

System Hardware and Sensors

Figure 6: Mercedes-Benz E320 research vehicle

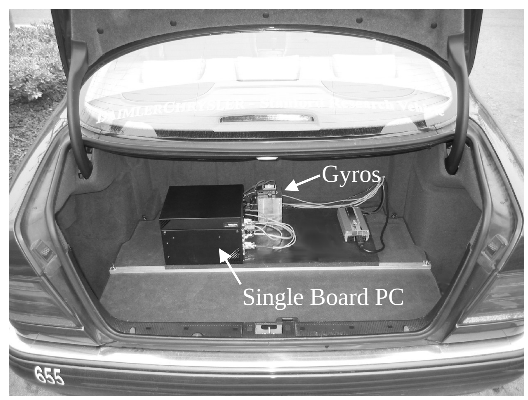

Figure 7: Trunk area showing single board computer system and gyros

The diagnostic system consists of two computers. One is a Versalogic VSBC-6 single board PC, configured with an additional Versalogic VCM-DAS-1 data acquisition module and Softing CAN-AC2-104 Controller Area Network (CAN) board. As show in Fig. 7, this system is mounted in the trunk of the vehicle. It communicates over ethernet with a laptop computer inside the vehicle.

The research vehicle is a 1999 Mercedes-Benz E320 sedan donated to Stanford University by DaimlerChrysler. As standard features, the vehicle has Anti-lock Braking System (ABS) and Electronic Stability Program (ESP). As part of the sensor package for these systems, the vehicle has low-resolution wheel speed sensors on all four wheels, a steering angle sensor, and a yaw rate gyro.

The wheel speed signals are available on the vehicle's CAN bus and are read using the Softing board. The steering angle value is not available over CAN, but instead is read

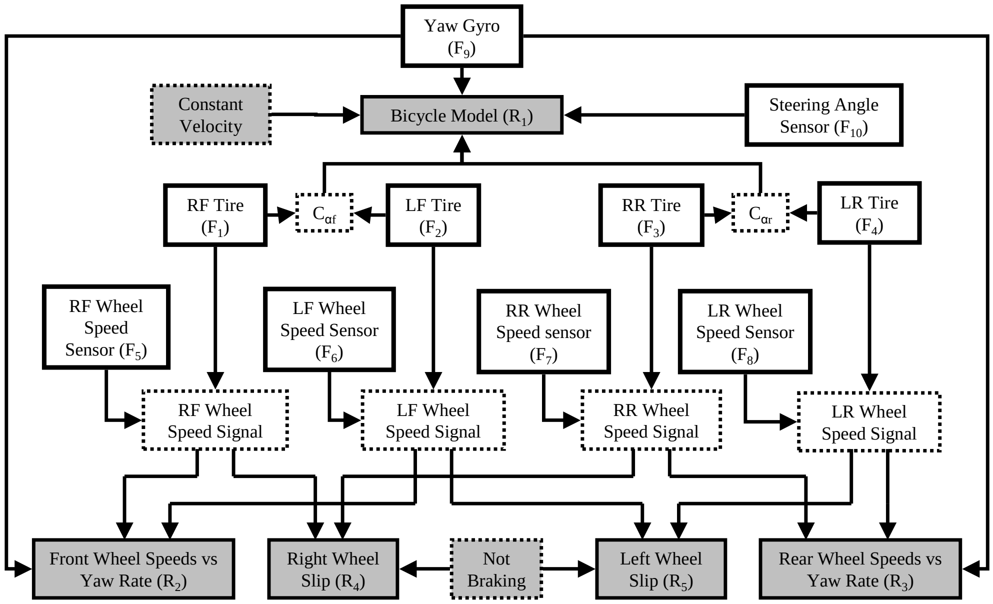

Figure 5: Bayesian network structure for vehicle stability control system

directly and interpreted using a custom designed board. The yaw rate signal is also not available over CAN, and therefore a separate Bosch automotive-grade yaw rate gyro has been installed in the trunk and its signal is read using the data acquisition module. For the production implementation of such a diagnostic system, all of the sensor values could be made available over CAN and no extra sensor hardware would be required.

Matlab Simulink and xPC Environment

In order to facilitate rapid development and frequent changes, the MathWorks' Matlab Simulink and xPC Target prototyping environments were used. The laptop inside the vehicle acts as the host computer and runs Simulink while the single board PC runs xPC Target.

The residual-generating system models are created in Simulink and downloaded to the single board computer where they process the sensor signals and send this information over the ethernet connection to the laptop. The models run at 100Hz, but the laptop only samples the residuals at 1Hz. The high update rate on the models is done in order to maintain the accuracy of the dynamic models. The residual sampling rate reflects the update rate of the dynamic Bayesian network running on the laptop. 1Hz was chosen in order to sample slower than the noise dynamics and fulfill the assumption that the residuals are not autocorrelated.

The DBN is programmed into Matlab using the Bayes Net Toolbox written by Kevin Murphy (Murphy 2001). Neither the programming language nor the code are optimized for speed, however, they allow for easy changes and can solve the BN at the required 1Hz rate. The result is a real-time diagnostic output available on the laptop screen inside the car.

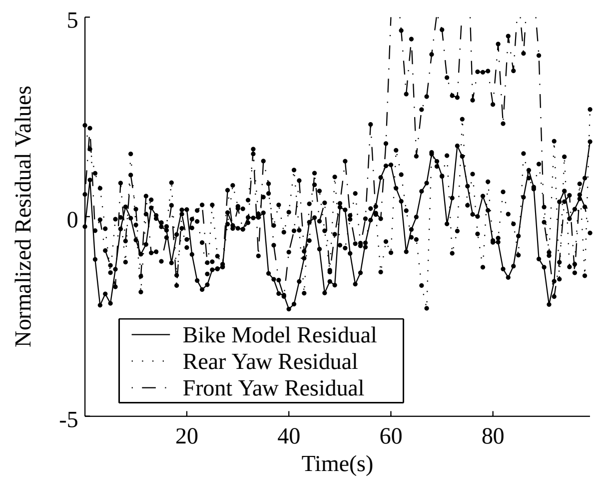

Figure 8: Residuals with simulated wheel speed sensor fault

Results

Testing is ongoing, and this section presents representative results showing the performance of the system. The faults examined have primarily been small magnitude sensor faults simulated in software, as these pose a greater challenge to the system than obvious physical faults such as a short circuit. Tire failures have not been tested due to the difficulty and danger in creating such a fault.

In one test, data was taken for 100 seconds of driving at speeds of 20-50 km/h. During the interval from 60 to 90 seconds into the run, a 1% bias was added to the left front wheel speed signal. This simulates a fault in the wheel speed sensor, such as the loss of a tooth on the sensor. Figure 8 shows the output of the bicycle model and two wheel speed yaw rate estimation models. The residual values have been

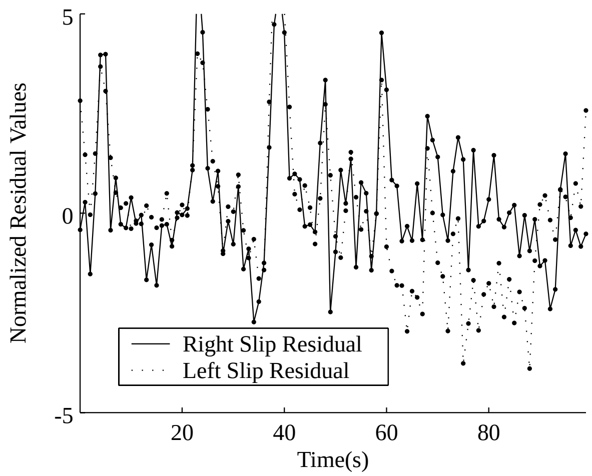

Figure 9: Residuals with simulated wheel speed sensor fault

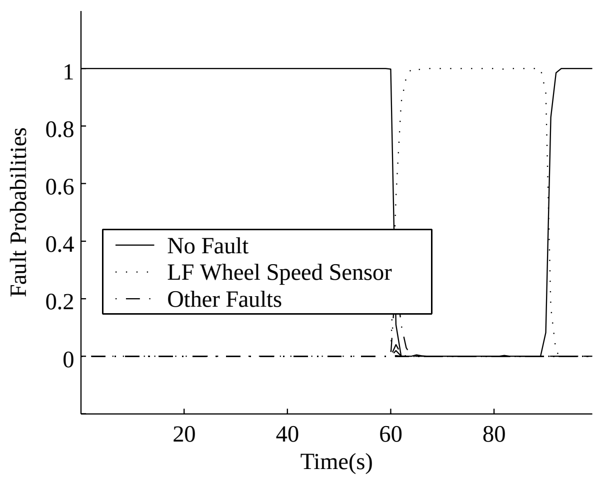

Figure 10: Diagnostic system output for simulated wheel speed sensor fault

normalized by their standard deviations. As is expected, the front wheel speed yaw rate residual was affected by the fault, while the rear wheel speed yaw rate residual and bicycle model were not.

Figure 9 shows the two slip model residuals for the same test. The left slip residual is influenced by this fault. It is interesting to note a few areas, such as around 25 and 40s, where some of the slip residuals are rather large even though there is not fault. This is where the brakes were applied, influencing the slip residuals.

Figure 10 shows the DBN's assessment of fault likelihood. The probability of no fault is very close to 1 except during the fault from 60 to 90 seconds. During this interval, the system accurately isolates the failure as left front wheel speed sensor. The system was not fooled by the high slip residuals during the periods of braking, because the braking status parameter "explained away" the large residuals.

Both the Bayesian network and a trained eye can detect that between 60 and 90 seconds, two of the five residuals are behaving abnormally. However, the deviation is slight and attempts at placing residual thresholds would have failed. There is no value that the residuals consistently exceed during the fault that they do not exceed during regular driving. Therefore, a system that relied on thresholds could not have detected this fault.

Conclusion

The prototype diagnostic system implemented demonstrates the power of multiple model diagnostics when residuals are integrated using AI. The use of a hybrid, dynamic, Bayesian network enables the detection of faults that could not be detected using techniques based on residual thresholds. In addition, the network provides an easy framework for including information such as modeling assumptions.

Due to the availability of models and sensor information, automotive systems are ideal applications for this technology. While the prototype presented examines only a small number of faults, these faults effect safety critical systems including ABS and ESP. If put into production, the additional cost of such a system would be limited to the processing power necessary for the models and Bayesian network since no extra sensors are required. Future advances in approximate DBN solution methods would help enable the use of cheaper processors and mitigate these costs.

References

Box, G.; Jenkins, G.; and Reinsel, G. 1994. *Time Series Analysis, Forecasting and Control*. Englewood Cliffs, New Jersey: Prentice Hall.

Carlson, C. R., and Gerdes, J. C. 2002. Identifying Tire Pressure Variation by Nonlinear Estimation of Longitudinal Stiffness and Effective Radius. In *Proceedings of AVEC 2002* 6^{th} *International Symposium of Advanced Vehicle Control*.

Frank, P.; Palkovics, L.; and Gianone, L. 2000. Using Wheel Speed and Wheel Slip Information for Controlling Vehicle Chassis Systems. In 5^{th} *International Symposium on Advanced Vehicle Control (AVEC)*.

Isermann, R., and Ballé, P. 1996. Trends in the Application of Model Based Fault Detection and Diagnosis of Technical Processes. In *Proc. of IFAC* 13^{th} *Triennial World Conference*.

Lerner, U.; Parr, R.; Koller, D.; and Biswas, G. 2000. Bayesian Fault Detection and Diagnosis in Dynamic Systems. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI)*, 531–537.

Murphy, K. 2001. The Bayes Net Toolbox for Matlab. In *Computing Science and Statistics*, volume 33.

Schwall, M. L., and Gerdes, J. C. 2001. Multi-Modal Diagnostics for Vehicle Fault Detection. In *Proceedings of IMECE 2001*, DSC–24600.

Schwall, M. L., and Gerdes, J. C. 2002. A Probabilistic Approach to Residual Processing for Vehicle Fault Detection. In *Proceedings of the American Controls Conference*, 2552–2557.