Proceedings of IMECE2001:
2001 ASME International Mechanical Engineering Congress and Exposition
November 11–16, 2001, New York, NY

IMECE2001/DSC10B-6

MULTI-MODAL DIAGNOSTICS FOR VEHICLE FAULT DETECTION

Matthew L. Schwall and J. Christian Gerdes

Design Division, Department of Mechanical Engineering, Stanford University
Terman Engineering, MC 4021, Stanford, CA 94305
E-mail: mschwall@stanford.edu Tel: (650)736-1504

Abstract

On-board vehicle diagnostic systems must have low development and hardware costs in order to be viable. Model-based methods have shown promise since they use analytical redundancy to reduce costly physical redundancy. However, these methods must also be computationally efficient and function accurately even with simple, low-cost models.

The approach presented in this paper uses multiple simple models to analyze dissimilar observable modes of a system. Residuals generated using the models are related and interpreted in a Bayesian network to determine fault probabilities and yield a diagnosis. The technique is demonstrated with a diagnostic system for automobile handling.

Keywords: Diagnostics, multi-modal, fault detection, Bayesian networks, vehicle handling.

I. Introduction

On-board vehicle fault detection and isolation (FDI) has received increased attention as market and safety demands have pushed for improved automotive maintainability and reparability. In conjunction with growing demands for better system performance and reliability, there are also strong constraints on the additional cost of such systems. This dictates that diagnostic 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 FDI methods have shown great promise in meeting many of these challenges. As shown in Fig. 1, model-based techniques consist of two stages [1]. 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. The use of analytical redundancy reduces the need for costly physical redundancy such as extra sensors, and has been shown to be successful in a wide variety of applications [2].

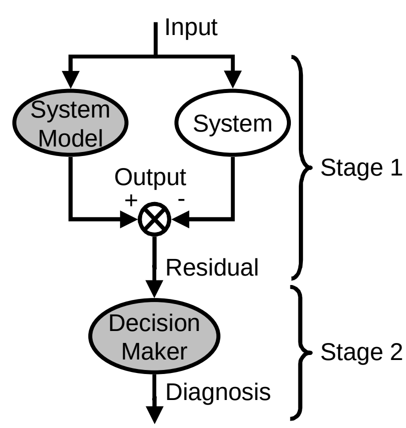

Figure 1 Generalized Structure of Model-Based Diagnostics

Most of the work in the field of quantitative, model-based, fault diagnosis is focused on the residual generation process with the assumption that making a decision based on very well designed residuals is relatively easy. Typically, the residuals are simply compared with a threshold to determine if a fault has occurred. However, due to the complexity of vehicle systems, the generation of well-behaved residuals often requires extensive modeling and high computational effort. This can translate into unacceptable engineering and on-board computer costs. Furthermore, standalone residual generation methods often lack a structure for setting thresholds. This paper illustrates a new methodology that recognizes the interdependence of the two stages. It includes a decision making stage that reduces the requirements on residual generation and thereby lowers the costs associated with the first stage. In this way, the total effort required to create a diagnostic system is reduced.

Multi-Modal Modeling

There are many different ways of modeling any complex system. Chittaro et al. [3] observed that models can differ in terms of ontology, representational assumptions, epistemological type, and level of abstraction. They showed that exploiting many diverse models can yield a better understanding of a system. This is the central idea behind Multi-Modal Diagnostics (MMD), the term for techniques that base diagnostic systems on multiple models from different domains.

In theory, there are many advantages to using a multi-modal approach. First, the multiple representations provide an easy means for incorporating system information provided from a variety of sources in the design process. Moyes et al [4] realized this advantage by combining design and operational knowledge. 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. For example, a holistic model could recognize that a wire diagnosed by the electrical system as being broken is incapable of transmitting signals for the control system. The challenge is to develop a framework for MMD that exploits all of these advantages.

Previous MMD Systems

In addition to the system mentioned above [4], a number of other MMD systems have been developed.

Pisu and Rizzoni [5] developed a method that represents a system using multiple modules consisting of one or more models. In doing this, they were able to demonstrate computational advantages of generating residuals using smaller, more manageable subsystem models. The method presented here differs from Pisu and Rizzoni's work in that our approach analyzes different observable modes of the same subsystem. In addition, the models are chosen based on availability, not component distinctions, thereby increasing the chance for model reuse and reducing modeling costs.

Douglas, Speyer et al [6] demonstrated the use of a neural network in an MMD system of an automated vehicle. Such a system can effectively analyze residuals but has practical limitations. Neural networks need to be retrained when modifications are made to the system or when more diagnostic models become available. In addition, as was observed by the authors, their system is inherently non-probabilistic. The advantages of a probabilistic system will be demonstrated later in this paper.

Yu, Biswas et al [7] developed a Multi-Level Diagnosis System for assisting airplane mechanics that integrates model-based diagnostics with an associational approach based on conventional expert system architecture. In the associational module, rule-based hypotheses are combined using the Dempster-Shafter probability model. The result is two modules that work together to guide a trained mechanic in determining the nature of a fault.

New MMD System

This paper presents a framework for diagnostic decision-making in which residuals from multi-modal system models are interpreted using a Bayesian network. This probabilistic approach possesses the accuracy and efficiency benefits of model integration inherent in MMD methods. In addition, the technique has many features unique to probabilistic approaches, such as the ability to integrate system models generated with different modeling assumptions and confidence levels.

Because the technique integrates modeling forms from other diagnostic techniques, at the extremes it resembles such methods. For instance, in the absence of probability data, the Bayesian network functions like a rule based or truth table decision system. In the absence of sensor data, the system represents only design probabilities and field data.

It should be noted that Bayesian networks have been used previously in diagnosing complex dynamic systems [8]. However, such methods are not multi-modal, but rather model the system directly with a Bayesian network. This requires a transformation for each model and results in a highly complex network that poses challenging computational problems.

Section II of this paper introduces the Bayesian network MMD framework. Section III then discusses the steps used in developing a Bayesian network-structured diagnostic system and the necessary constraints this imposes. Section IV introduces several models of a car's wheels and steering system and describes the application of the multi-model diagnostic methodology on this platform. Section V presents experimental results from this system. Results show that the system can distinguish any of the faults analyzed and plots are shown for the diagnosis of an example fault.

II. Probabilistic Framework

The general problem of FDI is one where, given information from m sensors, any member of a finite set of faults $F_1,\ldots,F_n$, can be detected and identified. Multi-modal diagnostic methods use a set of models $M_1,\ldots,M_p$ each of which use information from one or more sensors to generate one or more residuals. The set of residuals is $R_1,\ldots,R_q$, where $q \geq p$.

In order to implement a MMD system, a framework for incorporating models and interpreting residuals is needed. This research has utilized a Bayesian network for this task.

Bayesian Networks

A Bayesian network (BN) is a causal graph consisting of nodes connected by directed edges. The relationship between two connected nodes is defined by a conditional probability distribution (CPD). Bayesian networks can be powerful artificial intelligence reasoning systems particularly when given large sets of evidence and many relationships. For this reason, they have been successfully employed in fields such as medical diagnosis.

Bayesian networks contain only probabilistic relations and are therefore cumbersome for representing dynamic equations. In order to overcome this, the evidence in the Bayesian network comes from sensor data processed by residual generation methods such as those already used in model-based diagnostic methods. The difference is that when used by a Bayesian network, the residual generators need not be as robust or as complex as those employed by conventional diagnostic methods.

The models $M_1,\ldots,M_p$ are not nodes in the network. Rather, the evidence set is comprised of the residuals $R_1,\ldots,R_q$. The rest of the nodes are either part of the query $F_1,\ldots,F_n$ or hidden nodes, whose use in facilitating reasoning within the network will be shown later. With this structure, the BN is simply a structure for interpreting residual values to determine the probability that a component has failed.

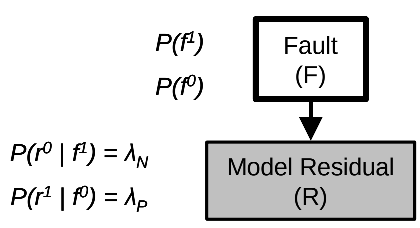

Figure 2 Diagnostic Conditional Probabilities

To demonstrate how residuals are used in BN-structured diagnostics, consider the example show in Fig. 2. This shows the most basic two-node BN with its associated CPD. In this example the fault can either be absent (f^0) or present (f^1). Similarly, the residual can have either an acceptable (r^0) or unacceptable value (r^1). In this paper, observable nodes, such as the model residual, are shown in shaded boxes. The CPD reveals that the residual has a finite chance of indicating either a false positive λ_P or a false negative λ_N. These parameters allow for model and sensor uncertainties to be taken into account.

This probability of the fault can be inferred given the state of the residual using Bayes' rule. More complex networks can be solved for the probability of any node using a variety of available exact and approximate techniques.

The BN above used binary, discrete variables. Other discrete, continuous, and hybrid distributions are also possible. Probabilities propagate in exactly the same manner though the solutions can be more complex.

Integrating Information

The driving advantage of structuring the MMD problem with a Bayesian network is the ability to integrate information without requiring an explicit analytical relation. This benefit of statistical modeling is useful when the observations are difficult or impossible to analytically relate; such as if one is a time-dependant measurement and the other exists in the frequency domain. An example of this synthesis of model types is shown in Fig. 3. This feature of BN-structured MMD is also beneficial when it is desirable to model the system at a level of abstraction where all observations are not included in the same model.

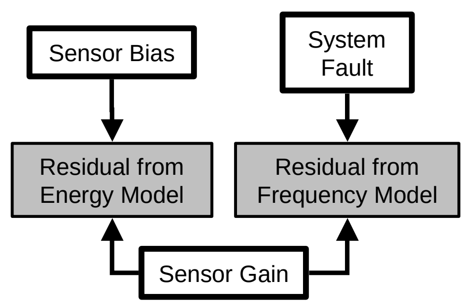

Figure 3 Example Bayesian Network Diagnostic System

By adding additional, non-physical, parameters to the Bayes net, extra information of diagnostic value can be included. For instance, an evidence node can be added to reflect whether or not the driver has reported a complaint about a given vehicle system. This information can add confidence to or reject a diagnostic hypothesis. This can be seen in Fig. 4, where either a faulty sensor or worn shocks can result in a residual, but a faulty sensor is unlikely to result in a customer complaint.

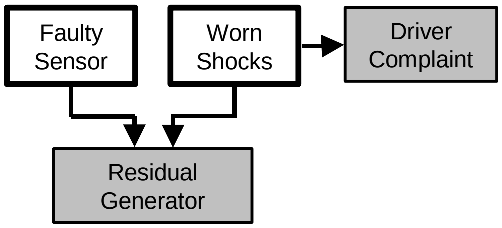

Figure 4 Bayesian Network with Customer Complaint Node

In making use of all available models, it is likely that some simplified models will not be applicable in certain operating conditions as a result of modeling assumptions. This can be overcome by using hidden nodes that indicate the status of the system. For a vehicle, these might take on values such as turning, braking or low speed maneuvering. The structure of a Bayesian network with such a parameter is show in Fig. 5 where the hidden node is shown in a dashed box. These parameters thereby enable the use of simplified models based on different modeling assumptions. An example of this is given in Section IV.

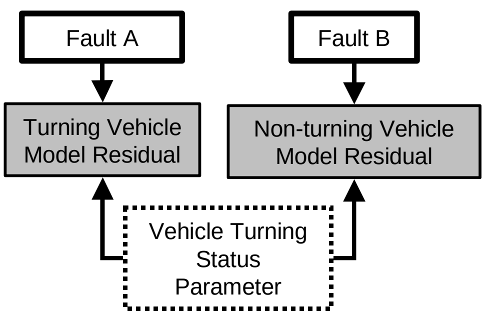

Figure 5 Bayesian Network with Operating Condition Node

Decision Making

The product of a diagnostic BN is a set of probabilities that each component has failed. Modeling uncertainties and sensor noise, which are captured in non-zero values for λ_N and λ_P, result in finite fault probabilities even when no fault exists. Therefore, in order to generate a deterministic diagnosis, threshold probabilities are needed. These thresholds are conceptually similar to those placed on residual values in conventional model-based diagnostic methods. Bayesian networks provide the information necessary to optimally set these thresholds. Recalling that the goal of diagnostics is to minimize risk and repair costs, the goal of setting thresholds is to minimize a cost function determined by these variables. For any single fault, there are four scenarios corresponding to the actual presence of the fault, f^0 or f^1, and whether a fault is declared or not, d^0 or d^1. The costs associated with each scenario are as follows:

$Cost(f^0,d^0)$ is zero since it is the cost of correctly inferring that the system is working.

$Cost(f^0,d^1)$ is the cost associated with a false positive diagnosis, which might include equipment downtime and maintenance costs necessary to discover the mistaken diagnosis.

$Cost(f^1,d^0)$ is the cost of missing a fault, which includes the eventual cost of repair plus any consequential damage caused by continued operation.

$Cost(f^1,d^1)$ is the cost of repairing a correctly identified fault.

The threshold should be set where the expected costs associated with d^0 are equal to those associated with d^1.

$$\begin{aligned} &P(f^0)\cdot Cost(f^0,d^0)+P(f^1)\cdot Cost(f^1,d^0)= \\ &\quad P(f^0)\cdot Cost(f^0,d^1)+P(f^1)\cdot Cost(f^1,d^1) \end{aligned} \tag{1}$$

Since $Cost(f^0,d^0)=0$, the equation can be rearranged as

$$\begin{aligned} &P(f^1)\cdot\left[Cost(f^1,d^0)-Cost(f^1,d^1)\right]= \\ &\qquad P(f^0)\cdot Cost(f^0,d^1) \end{aligned} \tag{2}$$

The optimal threshold is therefore

$$\begin{aligned} &P(f^1)_{Optimal}= \\ &\frac{Cost(f^0,d^1)}{\left[Cost(f^1,d^0)-Cost(f^1,d^1)\right]+Cost(f^0,d^1)} \end{aligned} \tag{3}$$

We note that in the term $\left[Cost(f^1,d^0)-Cost(f^1,d^1)\right]$, the cost of the repair falls out and leaving only the additional cost of consequential damage associated with missing the fault. Knowing the prior probabilities for the fault, we therefore have a closed-form solution for the optimal detection threshold given the cost of falsely identifying a fault and the additional cost of missing a fault.

In order to use these optimal thresholds in making a diagnosis, the diagnostic system continuously calculates the probability of each fault based on the current values of the residual. Each fault probability is compared against its optimal threshold, and if a threshold is exceeded, that fault is declared.

Thresholds are not only useful for detecting single faults. One of the unique features of a probabilistic network is the potential to determine joint probabilities in order to diagnose simultaneous faults. In most model-based diagnostic schemes, residuals are generated then compared with a threshold to determine whether or not a fault has occurred. These deterministic approaches then identify the fault based on a unique signature from multiple residuals. In order for these systems to work, the assumption is often made that faults occur only one at a time and are independent of each other. This assumption can be justified for certain applications where, in the event of a fault, the system can be reconfigured such that any other faults can subsequently be detected. However, in many vehicle systems, the ability to reconfigure is limited and scenarios involving multiple simultaneous faults can readily be foreseen. An example of this would be the presence of both low tire pressure and worn shocks. A BN-structured MMD system allows for an easy query of such joint probabilities in the form of $P(f_1^1,f_2^1|R)$.

The examples in this section have demonstrated the power of even very simple Bayesian networks. This results in scalability that has implications for the use of BNs during the design process. When using such a methodology, engineers need not wait until a diagnostic system is complete to make use of its reasoning power. For example, before a design has been finalized, an engineer can query the probability of a fault or the likelihood of detecting it. This can enable informed decisions regarding system robustness or sensor choice early in the design cycle.

III. Diagnostic System Generation

The task of creating Bayes net-structured MMD systems consists of three steps. The first is to gather existing models and transform them into a form compatible with the Bayes net. Next, the directed graph is created, and finally the conditional probability distributions associated with each node are created.

Model Preparation

In order for a model to be used in the Bayesian network, some information must be known about it. Specifically, one must know what variables influence the model and on what underlying assumptions it is based. For example, is the model valid under all operating conditions, or will qualifiers be needed? After this information has been gathered, converting design models to a form compatible with the Bayes net is in most cases a straightforward task. The Bayes net can accept any scalar residual so most existing residual generation techniques can be used without alteration. This includes all observer-based methods such as robust unknown input observer methods as described in Chen and Patton [1] and optimally robust frequency domain techniques such as those proposed by Frank and Ding [9].

The condition that inputs into the Bayes net must be a scalar poses a challenge for techniques that generate vectored residuals. This includes structured residuals [10] and those techniques based on the influence matrix approach of Ono et al [11]. With some manipulation, however, the vectors can be transformed into useable scalars that contain all the fault-specific information of the vectors. As defined by Poshtan, Doraiswami, and Stevenson [12] the influence matrix Ω is made up of individual influence vectors $\Omega_1,...,\Omega_j$ corresponding to j parameters. The feature vector deviation $\Delta\Theta$ is a directional residual. In order to transform this vector into scalar quantities corresponding to each fault, we find the projection of the feature vector deviation on each influence vector:

$$P_j = |\Omega_j^T \Delta\Theta| \tag{4}$$

Although $P_1,..,P_j$ are generated from the same model, they can now be treated as individual scalar residuals $R_i,\ldots,R_{i+j}$ and be represented by j distinct nodes.

Graph Generation

After any necessary transformations have been performed on the residuals, the next step is to define the directed graph that forms the structure of the Bayesian net. While the residual-generating first stage of the diagnostic process can use any one or more ontologies, the structure used in creating the graph structure for residual interpretation is based on system behavior. A connection is drawn to a residual from each fault or operating status parameter that affects the value of the residual. If multiple faults affect a residual through the same mechanism or parameter, a hidden variable can be added to in between the faults and the residual. This is done in Section IV with the vehicle's cornering stiffness parameters.

There are several advantages afforded by the nature of the network's structure. First is that creating the net is straightforward and doesn't require extensive knowledge about the physical structure of the system. The engineer only needs to know the function of each components and what other functionalities its behavior depends on. This method also results in a more efficient diagnosis, because in tracing a fault, only those suspects that contribute to the symptom need to be considered. Perhaps more important is that this ontology makes the network easy to update as the system changes, such as if a different component design or manufacturer is used. Functional design models are reusable for any devices that share common functionality. For example, switching actuator designs might affect the physical structure of the system, but would not change the functional design model and therefore the Bayes net structure.

Probability Distribution Generation

The final step in creating a BN-structured MMD system is to develop the conditional probability distribution associated with each node of the Bayesian network. This requires more in-depth knowledge than creating the physical structure of the network. The purpose of a CPD is to define the relationship between residuals and the faults they indicate. This may not only reflect the functional connection between the residual and a given fault, but also the degree of confidence in the residual value. For example, a residual that is generated using an inexact model or noisy sensor might be a weaker indicator of a fault than another residual. These relations can be expressed with terms such as λ_P and λ_N from the example in Section II.

Because of the amount of information contained within the CPD, approximate values can be used initially, and then adjusted later. The initial numbers can come from engineering estimates of simple scenarios. For more complex scenarios, such as a node with many parents, one approximation is to specify the probabilities given individual causes and then use stereotypical combination rules, such as a noisy-or, for combining them when more than one is present.

After approximate conditional probability distributions have been developed, simulation results, test data, or operating data can be used to refine them. The example presented in this paper is based entirely on estimated CPDs. However, the results are very encouraging and these distributions could be refined through learning if sufficient data were available.

IV. Example Application – Car Steering System

In order to demonstrate the concepts described above, this section describes a Bayesian network-structured MMD system applied to a car's handling system.

The input to the diagnostic system comes from six sensors $S_1,\ldots,S_6$ located on the vehicle:

$S_1,\ldots S_4$	4 wheel speed sensors
S_5	Yaw rate gyro
S_6	Steering angle sensor

For the purpose of this example, faults $F_1,\ldots,F_{10}$ will be considered:

Physical faults	
$F_1,\ldots,F_4$	Tire failure on any wheel
Sensor faults:	
$F_5,\ldots,F_8$	Wheel speed sensor error
F_9	Yaw gyro error
F_{10}	Steering angle sensor error

In order to diagnose this system, three simple models (M_1,M_2,M_3) 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.

Bicycle Model M_1

Using the classic bicycle model of vehicle cornering, dynamic equations relating steering angle to yaw rate and sideslip angle can be derived. This model assumes a constant longitudinal velocity and is captured in a pair of dynamic equations.

$$\dot{\beta} = \frac{-C_0\beta}{mV} - \left(\frac{C_1}{mV^2}+1\right)r + \frac{C_{\alpha f}}{mV}\delta \qquad (5)$$

$$\dot{r} = -\frac{C_1}{I_z}\beta - \frac{C_2}{I_z V}r + \frac{aC_{\alpha f}}{I_z}\delta \qquad (6)$$

where:

$$C_0 = C_{\alpha f} + C_{\alpha r}$$
$$C_1 = aC_{\alpha f} - bC_{\alpha r}$$
$$C_2 = a^2 C_{\alpha f} + b^2 C_{\alpha r}$$

and

a = Distance from front axle to CG
b = Distance from rear axle to CG
m = Vehicle mass
V = Longitudinal velocity
β = Sideslip angle
r = Yaw rate
δ = Steering angle
$C_{\alpha f}$ = Front cornering stiffness
$C_{\alpha r}$ = Rear cornering stiffness

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.

$$R_1 = |r_{expected(bicycle\ model)} - r_{measured(gyro)}| \qquad (7)$$

Yaw Rate and Wheel Speeds M_2

By comparing the left and right wheel speeds, the yaw rate of the vehicle can be estimated [13].

$$\text{Yaw Rate} = \frac{r_{tire}}{t}(\omega_r - \omega_l) \qquad (8)$$

Where:

r_{tire} = Radius of tire
t = Track width
ω_r = Right wheel speed
ω_l = Left wheel speed

Estimates of the yaw rate can be computed using data from both the front and rear tire pairs. When compared with the yaw rate measured by the yaw rate gyro, a residual can be generated for each pair indicating a fault in one of the parameters estimated or sensors used.

$$R_2 = |r_{expected(front\ wheel\ speeds)} - r_{measured(gyro)}| \qquad (9)$$
$$R_3 = |r_{expected(rear\ wheel\ speeds)} - r_{measured(gyro)}| \qquad (10)$$

Longitudinal Slip Model M_3

Miller et al [14] demonstrated identification of a longitudinal tire model of the driven wheels based on vehicle acceleration and wheel slip. The SAE definition of wheel slip is

$$S = \frac{V - R_e \omega}{V} \tag{11}$$

where V is the longitudinal speed of the wheel center, ω is the angular speed of the tire, and R_e is the effective tire radius, which is defined as the radius of the tire when free rolling.

For a two-wheel drive vehicle that is not braking, the undriven wheels are free rolling and by definition have zero slip. Assuming that both wheel centers on the same side of the vehicle have the same longitudinal velocity and effective radius, the slip on a driven wheel is given by:

$$S = \frac{\omega_{undriven} - \omega_{driven}}{\omega_{undriven}} \tag{12}$$

where ω_{driven} and $\omega_{undriven}$ are the wheel speeds of two wheels on the same side of the vehicle. This is a simplified model that could be improved by integrating information from the gyro. However, one of the advantages of the diagnostic structure is that it can tolerate simplified models such as this.

Miller et al's research showed good results with a linear tire stiffness and approximating the longitudinal forces on the vehicle as being entirely due to vehicle acceleration. Assuming that the force on the wheels is evenly shared between both left and right driven wheels:

$$F_x = \frac{1}{2} m a_x = C_x S \tag{13}$$

where F_x is the longitudinal force on one tire, a_x is the vehicle's longitudinal acceleration, and C_x is the longitudinal tire stiffness. Acceleration can be determined by differentiating the undriven wheel's speed, and C_x can be determined experimentally.

This allows for an estimation of S for both driven wheels. This can be compared with the actual slip determined from Equation 12. This difference yields one residual for each of the right and left wheel pairs.

$$R_4 = |S_{expected(acceleration)} - S_{measured(right\ wheel\ speeds)}| \tag{14}$$
$$R_5 = |S_{expected(acceleration)} - S_{measured(left\ wheel\ speeds)}| \tag{15}$$

Because of the assumptions made in deriving Equation 12, these residuals are only valid when the vehicle is not braking.

Bayesian Network Graph

Having gathered models M_1,M_2,M_3 with their associated assumptions, identified faults $F_1,\ldots,F_{13}$ and defined residuals $R_1,\ldots,R_5$, the Bayesian network can be constructed. Fig. 6 shows this graph, where the residuals generated by models are shown in shaded boxes and hidden nodes are shown in dashed boxes.

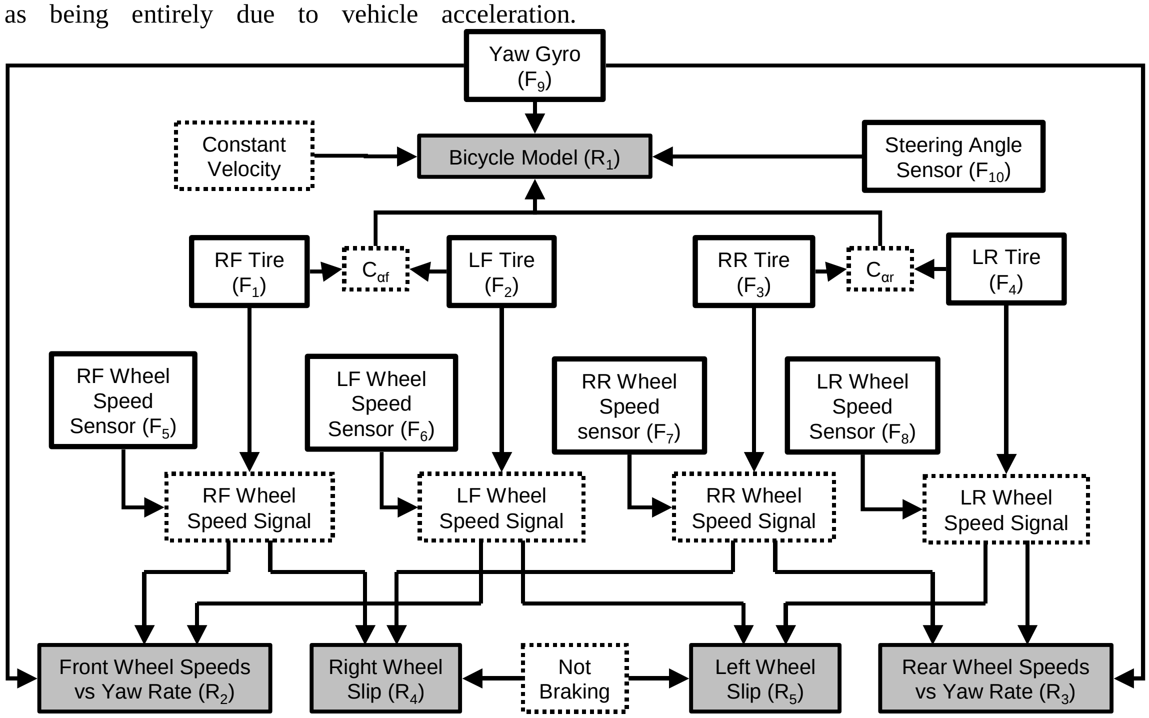

Figure 6 Bayesian network structure for vehicle handling system

With the structure of the BN in place, all that remains is to define the probabilistic relationships relating the nodes. As described above, this is initially done using engineering estimates and simplified scenarios. This example demonstrates the effectiveness of the diagnostic structure even in the absence of probabilistic data with which to refine the CPDs.

As an example, consider residual R_2, which is generated using the front wheel speed and yaw rate comparison model. This residual is affected by the node's parents in the graph; the yaw gyro F_9, right front wheel speed ω_{rf} and the left front wheel speed ω_{lf}. Given that one or more of these faults exists, R_2 is likely to be present (r_2^1). Otherwise, it is expected that the residual will be normal (r_2^0). Conceptually R_2 is therefore similar to an or-gate, and this basic understanding is sufficient to develop the conditional probability distribution for this node. This information is similar to what is contained in residual signature tables used in conventional FDI, except the CPD is probabilistic.

The CPD for R_2 is shown in Figure 7 where a 1 indicates the presence of a fault or abnormality of a parameter and 0 indicates the opposite. Instead of a strict or-gate, a noisy-or is used and the CPD reflects the small probability that the residual will exist even when none of the three faults are present; $P(r_2^1|f_9^0,\omega_{rf}^0,\omega_{rf}^0) = \lambda_P = .10$. This accounts for conditions such as an icy road where the model might not be valid. Also, there is a small chance ($\lambda_N = .05$) that even when one or more faults exist, the residual may not be present.

State of Parent Nodes			$P(r_2^0)$	$P(r_2^1)$
F_9	ω_{rf}	ω_{lf}		
0	0	0	. 90	.10
0	0	1	.05	.95
0	1	0	.05	.95
0	1	1	.05	.95
1	0	0	.05	.95
1	0	1	.05	.95
1	1	0	.05	.95
1	1	1	.05	.95

Figure 7 Conditional probability distribution for R_2

The numbers in the CPD are approximate, but they do not need to be exact in order generate useful inferences from the BN. In addition, the probabilities can be further refined using operating data as it becomes available.

V. Experimental Results

Experimental data was taken on a Mercedes E320 wagon. Wheel speed and steering angle information was obtained from factory standard sensors used by the electronic stability program. Yaw rate was measured using an automotive-grade gyroscope mounted inside the vehicle.

Numerous data sets of 30 second to 2 minute duration were analyzed. Driving types ranged from parking lot maneuvers to highway lane changes. The tests were performed in the absence of any faults. Faults were added later in simulation and all diagnosis was done in post-processing.

Multiple data sets were analyzed in order to set appropriate thresholds for the residuals. The bicycle model residual threshold was set at 0.04 rad/s. The thresholds for the differential wheel speed yaw rate models were set at 0.05 rad/s. The longitudinal wheel slip residual threshold was set at 0.7%.

The system was tested using simulated sensor faults including biases, offsets, and drifts. Results have been very promising. Provided that the faults are of sufficient magnitude that the residuals exceed their thresholds, the Bayesian network is able to discern any single fault, and some simultaneous faults. In order to convincingly diagnose a wider range of simultaneous fault combinations, more instantiated nodes are required. This would require more residual-generating models.

As an example of the system's capabilities, a single data run with a simulated fault is shown below. The data set shown was taken on a secondary road at speeds ranging from 40-60 km/hr. At 45 seconds, the sensitivity of the left front wheel speed sensor was decreased by 1%.

Figure 8 shows that the bicycle model residual is unaffected by the fault and remains below the threshold for the entire run.

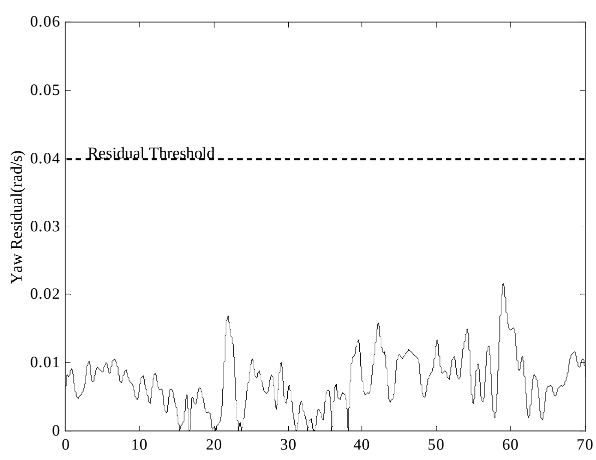

Figure 8 Bicycle model residual

The fault affects the front wheel speed yaw rate residual and, as seen in Figure 9, the residual exceeds its threshold immediately after the fault is initiated at 45s. However, the rear wheel speed yaw rate residual is unaffected and remains below the threshold for the entire run.

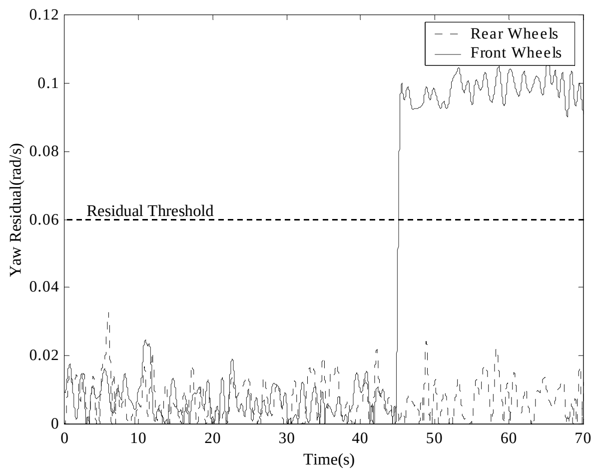

Figure 9 Yaw rate and wheel speed residuals

As seen in Figure 10, the left front wheel speed sensitivity fault results in a large residual from the left wheel slip model. This residual also exceeds the threshold at 45 seconds.

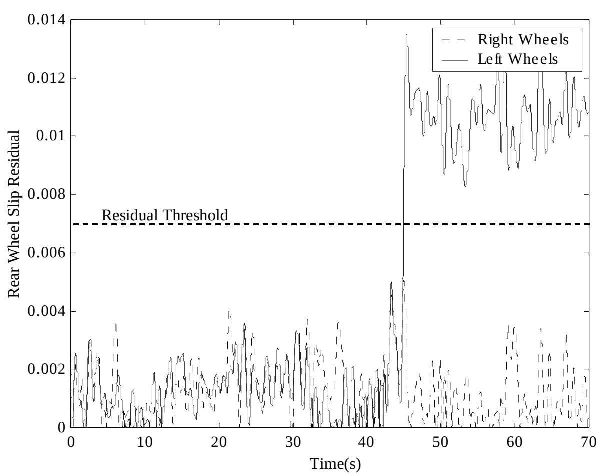

Figure 10 Longitudinal slip residual

The binary values of the residuals are used as evidence in the Bayesian network. The network was solved using junction tree exact inferencing and yielded the fault probabilities shown in Figure 11. Before the fault is initiated, the probability of all faults is insignificant. After 45s, the probability of a fault in the right front wheel speed sensor is 78%. There is also a 6% probability of a fault in the right front tire. This slight uncertainty is a result of the confidence levels in the sensors and models and is directly affected by the choice of λ_N and λ_P. The probability of other faults is very small.

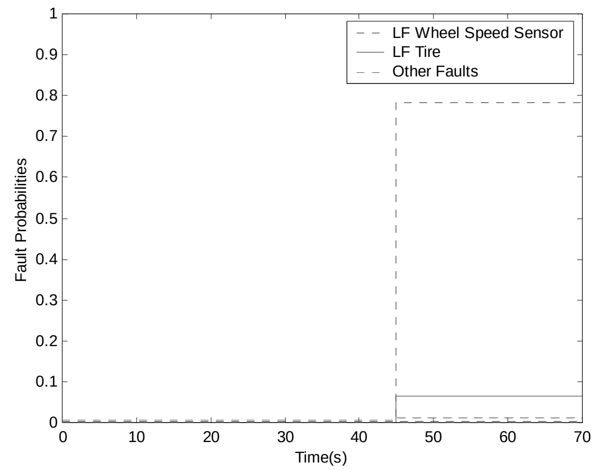

Figure 11 Fault probabilities for simulated fault in left front wheel speed sensor

In order to determine whether or not a 78% chance of a faulty wheel speed sensor is sufficiently large to alert the driver, the probability is compared against the optimal threshold for this fault. Calculating this threshold requires knowing the costs of falsely identifying the fault and of missing the fault. In this case, one would need to know the cost to check out a falsely diagnosed sensor and the cost of continuing to drive with a faulty wheel speed sensor (and therefore a faulty ABS system). For example, if the first cost were $100 and the second were $1000, using Equation 3 the optimal threshold to declare a fault would be 9%. This would indicate that the fault should be declared.

VI. Conclusions

This paper has detailed a new method for combining models for Multi-Modal Diagnostics. The Bayesian network framework enables the use of multiple system models with varying confidence levels, ontologies and assumption sets. Therefore, models created during the design process can be integrated. In addition, this structure facilitates the inclusion of all available information including the state of the vehicle and customer complaint information. The result is a methodology capable of generating low-cost diagnostic models early in the design cycle.

This paper has detailed the application of this methodology to a vehicle's handling system. The resulting Bayesian network is both computationally manageable and powerful. Results showed that the system is highly effective for diagnosing any single fault and select simultaneous faults.

In the example given with a small number of faults and few, high confidence residuals, the Bayesian network decision system begins to functionally resemble a truth table. At this extreme, the system becomes comparable to the lateral diagnostic model developed by Shrivastava and Rajamani [15]. The power of BN-structured MMD, however, comes in its flexibility and ability to integrate models into a complex probabilistic structure.

In order to build upon the system developed here, more models could be added and CPDs could be learned from repair history data. In addition, a dynamic Bayesian network could be used. Such a network would take advantage of the temporal correlation between faults. This has the potential to provide filtering such that residual thresholds could be lowered in order to detect smaller faults.

References

[1] Chen, J. and Patton, R. J., "Robust Model-Based Fault Diagnosis for Dynamic Systems," Kluwer Academic Publishers, 1999.

[2] Isermann, R. and Ballé, P., "Trends in the application of model based fault detection and diagnosis of technical processes." *Proc. IFAC 13th Triennial World Conference,* 1996.

[3] Chittaro, L. Guida, G. Tasso, C. and Toppano, E., "Functional and Teleological Knowledge in the Multimodeling Approach for Reasoning About Physical Systems: A Case Study in Diagnosis," *IEEE Trans. On Systems, Man and Cybernetics,* Vol. 23, No. 6. November/December 1993

[4] Moyes, A. et al, "Combining design and operational knowledge to enhance generator plant diagnostics," *IEE Proceedings Generation Transmission Distribution*, Vol. 143, No3, May 1996.

[5] Pisu, P. et al, "Model Based Diagnostics for Vehicle Systems," *Proc. of ASME Int. Mechanical Engineering Congress and Exposition – Dynamic Systems and Control Division, 2000*.

[6] Douglas, R., Speyer, J., et al. "Fault Detection and Identification with Application to Advanced Vehicle Control Systems," *California PATH Research Report* UCB-ITS-PRR-95-26.

[7] Yu, W., Biswas, G., Weinburg, J., "MDS: An Integrated Architecture for Associational and Model-based Diagnosis" *Applied Intelligence*, March/April 2001; v.14, no.2, p.179-195.

[8] Lerner, U., Parr, R., Koller, D., and Biswas, G., "Bayesian Fault Detection and Diagnosis in Dynamic Systems," *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-00).*

[9] Frank, P. and Ding, X., "Frequency Domain Approach to Optimally Robust Residual Generation and Evaluation for Model-based Fault Diagnosis," *Automatica* v 30 n 5 May 1994. p 789-804.

[10] Gertler, J., *Fault Detection and Diagnosis in Engineering Systems*, Marcel Dekker, New York, 1991.

[11] Ono, T. et al, "Influence Matrix Approach to Fault Diagnosis of Parameters in Dynamic Systems," *IEEE Trans. on Industrial Electronics*, Vol. IE-34, No 2, May 1987.

[12] Poshtan, J. Doraiswami, R. and Stevenson, M., "Influence Matrix Based Real-Time Fault Diagnosis," *1996 IEEE Int. Conf. on Control Applications*.

[13] Frank, P., Palkovics, L., and Gianone, L., "Using Wheel Speed and Wheel Slip Information for Controlling Vehicle Chassis Systems," *5th Int'l Symposium on Advanced Vehicle Control (AVEC) 2000*.

[14] Miller, S. et al, "Calculating Longitudinal Wheel Slip and Tire Parameters Using GPS Velocity, " *Proceedings of the American Controls Conference*, 2001, pp 1800-1805.

[15] Shrivastava, A., and Rajamani, R., "Fault Diagnosis for GPS-Based Lateral Vehicle Control," *Proceedings of the American Controls Conference* 2001, pp 31-36.