A Prototype for Model-based On-board Diagnosis of Automotive Systems

Martin Sachenbacher [a], Peter Struss [b] and Claes M. Carlén [c]
[a] *Technische Universität München*
Department of Computer Science
Orleansstr. 34, 81667 München, Germany
E-mail: sachenba@in.tum.de
[b] *OCC'M Software GmbH*
Gleissentalstr. 22, 82041 Deisenhofen, Germany
E-mail: struss@in.tum.de
[c] *Volvo Car Corporation*
IT and Product Documentation
Department 98101, 40508 Göteborg, Sweden
E-mail: ex1.cmc@memo.volvo.se

Within the European "Vehicle Model based Diagnosis" (VMBD) project, demonstrator vehicles with built-in faults provided a serious challenge to model-based diagnosis techniques and a real-life test-bed for their evaluation. One of the guiding applications within VMBD was model-based on-board diagnosis of faults in a turbo diesel engine system with a focus on potential origins of increased carbon emissions. This paper focuses on the application aspects. We discuss the requirements imposed, the way they were addressed by the chosen solutions, and the results obtained by the on-board diagnosis prototype running on the demonstrator vehicle. The most important challenges of the demonstrator were to apply model-based diagnosis systems to dynamic systems with feedback, to handle systems without a rigorous mathematical model (such as a combustion engine), and to try to provide the response times required for real-time applications.

Keywords: model-based diagnosis, qualitative reasoning, real-world applications

1. Introduction

Research on model-based diagnosis (e.g. [5], [8]) has generated a number of well-founded theories and sophisticated prototypes of implemented diagnosis engines. However, many of these diagnosis systems have only been applied to toy examples or to problems that ignored the practical context of industrial applications. As a result, the transfer of the technology into practice is well behind the expectations, despite the fact that it promises to meet some crucial requirements of automated diagnosis for industrial needs.

Car industries provides a good example of such industrial needs. It is estimated that European passenger cars have an average yearly down-time of 16 working hours due to malfunctions and maintenance. This figure is even greater for commercial vehicles. For the European Community alone, this amounts to a total of over one billion hours for diagnosis and repair. At the same time, with increased environmental awareness, stricter constraints are imposed on the car manufacturers to develop clean cars, and also to keep them clean during their life cycle (see, for example, [13]). These growing constraints are reflected in increased requirements on on-board diagnostics development. For engine management control units, currently about one half of the software is dedicated to diagnosis, and this share is still growing.

In response to this situation, several car manufacturers and suppliers joined to launch the Brite-EuRam project VMBD (Vehicle Model Based Diagnosis) with the intention to promote the transfer of model-based diagnosis technology by the challenge of applying it to on-board and off-board diagnosis of passenger cars. The results and system performance were evaluated on real demonstrator vehicles. Within this project, Volvo Car Corporation, Robert Bosch GmbH, and OCC'M Software GmbH produced a model-based system that diagnoses problems related to increased carbon emissions of diesel engines, a problem of significant importance w.r.t. environmental impact and compliance with legal requirements. The work was built upon previous case studies in diagnosis and failure analysis of car subsystems carried out in collaboration between Robert Bosch GmbH and the Technical University of Munich ([22], [18]). The sys-

AI Communications
ISSN 0921-7126, IOS Press.

tem transforms the sensor signals that are available to the standard electronic control unit (ECU) on-board to a qualitative level and exploits them for detecting and localizing faults based on a model of the turbo control system. It has been installed on a Volvo demonstrator vehicle with a number of built-in faults.

The goal of this work was to face the requirements imposed by this kind of application. In particular, we expected answers to questions like: Are model-based diagnosis systems able to diagnose dynamic systems with feedback (like a turbo control system)? If they are, can they provide a sufficient response time (for fast processes like the ones in the engine)? And, even more fundamentally: Are they applicable if no rigorous mathematical models are available (as is the case for the combustion process)? These questions have triggered some interesting research ([12], [18], [20]). However, in this paper, we want to share our experiences regarding the application aspects.

Since one cannot learn about the problems to be addressed in the transfer of the technology into industrial applications without understanding their nature to some extent, we start with a brief explanation of the respective vehicle subsystem and then discuss the most relevant requirements to be addressed. In section 3, we outline the technical solutions we adopted and discuss how and to what extent they satisfy the requirements introduced. Finally, we describe the set-up of the experiment and summarize the results obtained.

2. The Application Domain

The demonstrator vehicle used in the VMBD project is a Volvo car equipped with a so-called distributor-type diesel injection (DTI) system ([3]). The DTI is an approved system which has been on the market for many years. However, increased legislative and customer demands have lead to new requirements especially for aspects related to emissions and performance of this system. Figure 1 shows the part of the system which is responsible for supplying air to the diesel engine. It can be decomposed into the exhaust gas recirculation (EGR) subsystem (upper part of figure 1) and the turbo control subsystem (lower part of figure 1).

The purpose of the exhaust gas re-circulation system is to return a certain amount of the ex-

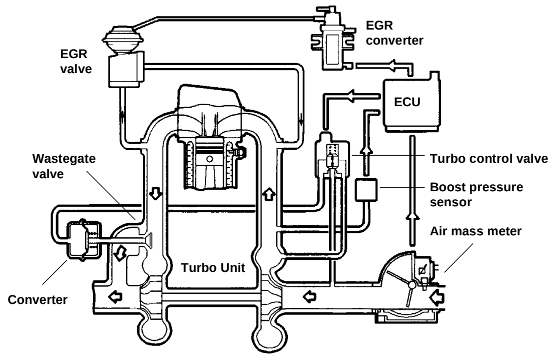

Fig. 1. Turbo control and exhaust gas recirculation subsystem of the DTI.

haust gas to the intake air to decrease the oxygen rate of the intake air and thus to reduce emission levels of the fuel combustion. Depending on driving conditions, the ECU governs the EGR converter to achieve a certain air pressure in a control pipe, which in turn sets the position of the exhaust gas re-circulation valve. The position of this re-circulation valve then determines how much of the exhaust gas is fed back to the air intake pipe.

The turbo control subsystem consists of a turbocharger turbine, which is driven by the engine's exhaust gas, for compressing (and thereby increasing the mass of) the air taken into the engine. The ECU controls the boost pressure (i.e. the pressure in the engine intake pipe) admitted in a certain driving situation by opening or closing the turbo control valve, which determines the position of a so-called waste-gate valve. The position of this valve determines how much of the exhaust gas drives the exhaust turbine of the turbo-charger.

The ECU not only issues commands to the actuators, but also monitors and checks the sensor values it receives from these systems. The goal of this so-called on-board diagnosis is to signal alarms to warn the driver in case of a failure and to generate fault codes that can be further used in the service bays to track down a failure.

For failures which are considered critical, onboard diagnosis also aims at selecting appropriate recovery actions. The built-in recovery actions that will be performed depend on the assumed failure and the expected failure effects and range from minor performance reductions to full engine stop. They attempt to take the vehicle back into safe operational conditions (so-called limp-home modes), which allow the driver to reach the next service bay, for instance.

The design of such diagnosis capabilities has to cope with a number of fundamental challenges: Variant problem. The systems come in many different variants. The DTI system we dealt with is only one specific instance. The reason is that a supplier of automotive subsystems has to develop his products for many vehicle manufacturers and a lot of vehicle models, which all impose different requirements on the base system. The actual configuration may differ in the number of sensors, and redundant parts may be present or absent dependent on the specific car manufacturer. Also, the components themselves come with different constructive details. These modifications must be thoroughly handled in the diagnosis algorithms. However, generating specialized diagnostics for all variants by hand is a very expensive task.

Dynamic and controlled subsystems. The above device is a controlled system which has internal states depending on previous inputs. Hence, failures may be observable only in a subset of the operating modes of the vehicle (e.g. engine start, idling, take off phase, full acceleration, etc.) or transitions between operating modes. For example, the pressure in the air hose between the turbine outlet and the intake manifold of the engine varies, depending on driving conditions, from atmospheric low pressure (during idling) to about 2 bars overpressure (during full acceleration). A leakage at this point may therefore, depending on its size, only be perceivable during high power demands. Additionally, the control unit tends to compensate for failures during certain operating modes. As an example, figure 2 shows measurements taken for an electrical failure in the air mass sensor (which is a combination of air flow and air temperature meter). The failure has an effect only when the engine is idling, while it is compensated for during driving. This all becomes a problem especially in combination with low measurement granularity, that is, if measurements within feedback loops are sparse over time. Determining correct diagnostics that cover these situations is a complex task, often infeasible for hand-crafted diagnostic procedures which are based on predefined range or plausibility checks only.

Limited system information. The knowledge about the behavior of some vehicle components, especially components that involve several physical domains like the engine and its combustion process, is incomplete and precarious. Particularly in the case of a failure, quantities are not known exactly or even remain unknown. In contrast to electric or hydraulic components, there exist no general mathematical models one could start from. However, to exploit the kind of symptoms we are interested in, it is necessary to establish a link between emission- or performance-related phenomena and components in the various subsystems (e.g. a too high exhaust gas recirculation rate or a too low pressure of the intake air). To some extent, it is therefore necessary to reason about the behavior of such ill-specified components as well.

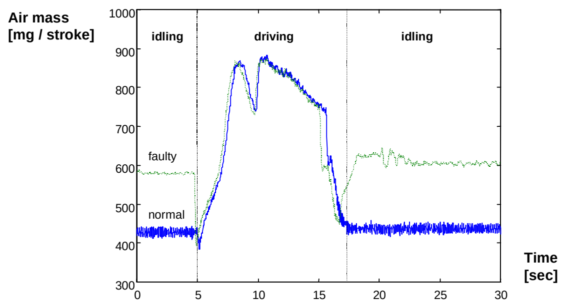

Fig. 2. Measurement taken for a failure which is compensated during driving and shows up in idle mode.

Limited observability. Very few sensors are available in car systems. This is true e.g. for the hydraulic part of the fuel injection system, which contains no sensor at all. In addition, the context in which a car is operated in (e.g. road and weather conditions, load) is highly dynamic, uncertain and often neither measurable nor reproducible. The main consequences are noisy signals and rather qualitative symptom descriptions. Diagnosis has to be capable of processing such information.

Real-time requirements. On-board diagnosis must come up with a conclusion before the system has to move to another state (e.g. shut off the engine) to prevent safety-critical situations or severe system damage, or to comply with legal restrictions. As a consequence, the computational requirements of on-board diagnosis functions must be relatively low to bring them into state-of-the-art ECUs.

Currently, the DTI control unit is equipped only with a restricted form of on-board diagnostic capabilities. It continuously monitors part of the sensor signals using predefined range and plausibility checks and is able to detect a limited number of faults on this basis. However, in many cases it fails to discriminate among the different possible causes that lead to the failure. Consequently, it sometimes applies a more restrictive recovery action than would actually be necessary. One way to improve this would be to exploit more of the interdependencies between different signals. But currently, this is not done in a general and systematic way, leading to sub-optimal solutions. This reflects the fact that at present the development ("design") of automotive on-board diagnosis does not follow precisely defined methodologies or criteria and is generally regarded as an addendum to the design of the system. On the other hand, a growing interest of car manufacturers lies in the possibility to utilize the same set of techniques for all phases of the design process (e.g. for pre-design, function-design and layout of the diagnostic concept), as well as for diagnosis and fault tracing in the service bay. Since car manufacturers work with a number of system and component suppliers, there is a high, but currently unfulfilled, demand for developing a standardized methodology among the parties involved.

In accordance with the overall thrust of the project, our goal thus was

- to produce a prototypical model-based diagnosis system that is capable of diagnosing faults in the diesel engine based on the sensor signals that are available to the ordinary ECU,
- to this end, generate a library of models of the relevant components, and
- to perform this task in a systematic way as a contribution to a general methodology for producing on-board diagnostics.

In other words, on the one hand the prototype had to solve the specific problems given by the fault scenarios on the real vehicle. On the other hand, the method for obtaining the solution should be applicable to a broader class of automotive subsystems and diagnosis problems. While the first task is certainly achievable with current practice in engineering, the second goal in our view promises merits to be earned by AI technology.

3. Technical Foundations of the Prototype

In the following subsections, we will outline the technical foundations of the implemented prototype and discuss in what way and to what extent they address the requirements described in the previous section. The goal of this paper is not to present all underlying theoretical results; for technicalities, the reader is referred to the referenced papers.

3.1. Consistency-Based Diagnosis

The choice of a model-based approach is founded on the recognition that most complex subsystems in a vehicle share the following features with respect to their function:

- there exists a natural decomposition into subsystems with only few component types,
- in most cases, malfunctions of the car or a subsystem are due to some component failure,
- component behavior can be described by relations among local variables and parameters,
- system behavior is established by the behavior of its components and their connections w.r.t. processing of material, energy or signals.

Therefore, we applied a knowledge-based systems approach to diagnosis which is based on the fundamental idea to separate and represent different elements of knowledge within a computer program. It consists of a declarative and modular representation of knowledge about a family of technical devices in terms of a library of component behavior models and, separated from this knowledge, a set of domain-independent diagnosis algorithms to exploit the models.

In a nutshell, the standard, so-called consistency-based approach to diagnosis ([5]) can be described as follows (see figure 3):

- observations of the actual behavior of the system are entered.
- based on the device model, conclusions are computed about system parameters and variables (observed and unobserved). For each derived prediction, the set of component models involved in it is recorded. This information can be determined by the diagnosis system because the device model has a structure that reflects the device constituents.
- if a contradiction is detected, i.e. conflicting conclusions for a variable occur (fault detection), the set of components involved in it indicates which components possibly deviate from their intended behavior.

– diagnosis hypotheses are generated, i.e. sets of faulty components that account for all detected contradictions (fault localization).
– in case models of faulty behavior are provided, the same approach (checking consistency of a model with the observations) can be used to discard particular faults (fault identification) or to conclude correctness of certain components if the set of modeled faults is considered complete.

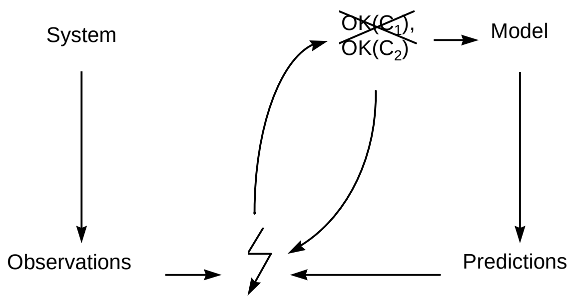

Fig. 3. Consistency-based approach to diagnosis.

This diagnosis framework has the desired property to systematically exploit the analytic redundancy among the available sensor signals. The model-based approach alone provides one answer to the methodological challenge, because its underlying principles (and the implementation) are independent of the particular subsystem and enable the re-use of the involved software components. Generating a specific diagnostic system is thus reduced to generating an appropriate model of the system to be diagnosed.

As stated above, component-oriented modeling is a natural approach in our application domain. Beyond this, it is the key to solving the variant problem, because the model of a subsystem is derived as the aggregation of standard building blocks. This is another element of a general methodology and enables the automated generation of a device model and, hence, of a tailored diagnosis system based on a structural description of the device only (which should be the natural output of a CAD system). A way of creating diagnostics for all variants of vehicle subsystems is thus obtained that is systematic as well as efficiently supported by computer tools. Figure 4 illustrates this idea.

For diagnostic purposes, faults can be described as certain component failures, and fault models associated with the respective components. This provides a principled way of capturing knowledge about faults in a modular way which contrasts other approaches in AI (based on storing associations between symptoms and faults for each device in terms of rules or cases) or engineering (trying to identify parameter deviations in a closed mathematical model of the entire device).

Since a component model is meant to be used within the contexts of various devices, it has to capture a behavior description which must not presume a specific context and, particularly, not the correct functioning of the rest of the device. The strict discipline in modeling required to achieve this goal is another important element of the methodology.

It is interesting to note that we need not to build a model of the control unit behavior itself, unless we want to detect faults in the ECU. Due to the fact that the model runs within an on-board environment, all the control unit's signals will be available for observation. Consequently, a behavior model of the control unit could never be part of a diagnostic hypothesis, and would therefore be useless.

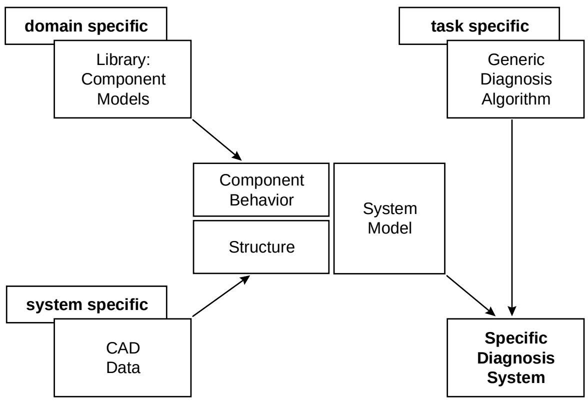

Fig. 4. Automated generation of model-based diagnostic systems.

3.2. State-Based Diagnosis of Dynamic Systems

As outlined above, consistency-based diagnosis requires checking whether the behavior that is predicted by the model is consistent with the observed behavior. Formally, if $model(mode)$ denotes the model of a behavior mode and OBS denotes a set of observations, this means checking the consistency of

$$model(mode) \cup OBS. \quad (1)$$

If we apply this diagnosis approach to dynamic systems like in our application, a crucial question is whether this requires prediction of behavior over time, i.e. simulation. Frequently, this is taken for granted when dynamic systems are concerned. However, it turns out that diagnostic results can be obtained based on checking consistency of observed and modeled states only, i.e. without performing simulation. Sometimes this can be done even without any loss at all (see the recent experimental and theoretical work in [7], [12], [23], [18]).

The idea of the state-based approach to diagnosis of dynamic systems is to separate the constraints of a dynamic model into constraints that restrict the possible states of the system at a point in time (*stateconstr*) and constraints that restrict the possible transitions from one state to another (*tempconstr*):

$$\begin{aligned} &model(mode) = \\ &stateconstr(mode) \cup tempconstr(mode). \end{aligned} \quad (2)$$

State-based diagnosis then checks only if

$$stateconstr(mode) \cup OBS \quad (3)$$

is consistent, that is, it is only determined if the set of observed states occurs in the set of possible states of the system. Since the actual order in which states occur is thus not taken into account, this might seem much weaker than a simulation-based approach which compares the actual behavior over time and the model behavior simulated over time. However, if temp-constraints contains only constraints which capture general laws of continuity, integration and derivatives that can never be violated (called CID-constraints in [18]), adding the temp-constraints to the consistency check will not yield additional contradictions, and they can be removed from the consistency check without affecting the result of diagnosis (see [18]).

Avoiding simulation and performing the consistency check for states only provides a great computational advantage, in particular, if the system needs not simulate possible faults, and is one contribution to achieving the required response times of the diagnosis system.

Another contribution from the algorithmic part to meet on-board time restrictions for diagnosing dynamic systems comes from techniques for reusing predictions over different time points (see [7]). On top of a state-based approach, this "temporal caching" enabled to run model-based diagnosis fast enough to cope even with real-time requirements, as can be seen from the example section 4.

3.3. Qualitative Models and Qualitative Deviation Models

A second key to provide the efficiency that enables to run model-based diagnosis even on-board a vehicle is to avoid the computational complexity of numerical modeling and simulation on the part of the model.

The consistency-based diagnosis principle is obviously independent of the specific modeling formalism and works, for instance, with numerical component models. But, as pointed out, if systems operate in an ill-specified environment, accurate component models might be unavailable and/or unnecessary, partly because of the nature of the involved processes (e.g. friction or combustion) or because parameters change due to wearing. Qualitative reasoning ([25]) has developed mathematical foundations and systems for representing partial knowledge about system behavior and computation of behavior descriptions based on this. Since qualitative modeling allows to cover classes of systems (by ignoring irrelevant differences in parameters) and classes of faults (e.g. leakages of different or unknown sizes), it also helps to keep the library of model fragments manageable.

It turns out that in some cases it is not even relevant to reason in terms of the actual values of quantities. Rather, it can be sufficient to reason in terms of (qualitative) deviations from nominal values only. For example, if the EGR valve closes (much) slower than normally, the oxygen rate of the intake air will be (substantially) lower than normally. It may suffice to explain why the oxygen rate is lower than it should be, regardless of the actual value of the oxygen rate. Descriptions of deviations can reflect the fact that it may be unnecessary or impossible to specify the normal behavior exactly and numerically.

In [12] and [18], we described such models which capture the deviation of an actual value from some reference. For each variable, this deviation can be represented as

$$\triangle x := x_{act} - x_{ref}. \tag{4}$$

In particular, the reference can be given by the normal behavior of the respective system, i.e. by the condition that all components operate correctly:

$$\triangle x := x_{act} - x_{corr}. \tag{5}$$

The key idea to use this is that the models describe how deviations from the normal behavior are generated or propagated by the individual components. For instance, a faulty valve with increased friction does not open (fast) enough, i.e. it creates a negative deviation of the cross-sectional area, A (or the derivative of A). Still working as a valve, it thus causes a negative deviation of the amount of flow through the valve. A (correctly working) pipe connected to the valve will simply propagate this deviation in flow. It is important that this kind of modeling can be performed without leaving the firm ground of first principles. Qualitative deviation models can be obtained from the physical laws in a canonical way. For instance, consider the equation

$$A = \beta \cdot s^2 \tag{6}$$

where $s > 0$, and b is a positive constant. Then, if [.] denotes sign abstraction,

$$\begin{aligned} &[\triangle A = \\ &= [A_{act} - A_{corr}] = \\ &= [\beta \cdot s_{act}^2 - \beta \cdot s_{corr}^2] = \\ &= \beta[(s_{corr} + \triangle s)^2 - s_{corr}^2] = \\ &= [s_{corr}^2 + 2s_{corr} \cdot \triangle s + \triangle s^2 - s_{corr}^2] = \\ &= [\triangle s \cdot (2s_{corr} + \triangle s)] = \\ &= [\triangle s] \end{aligned} \tag{7}$$

where the last step follows, because $2s_{corr} + \triangle s > 0$.

Actually, we could have obtained this result in a simpler way, because for any monotone function f,

$$\begin{aligned} &y = f(x) \\ &\Rightarrow [\triangle y] = [f(x_{act}) - f(x_{corr})] = [\triangle x]. \end{aligned} \tag{8}$$

This simple example illustrates the desired feature that using deviation models potentially avoids the necessity to specify the correct behavior of components explicitly. Based on a theoretical foundation of such deviations that is not restricted to simple functions as in the example ([16]), component models have been developed that declare and propagate deviations from some nominal or reference behavior (which is possibly left unspecified), even across different domains.

For hydraulic and pneumatic components in our application domain, the starting point is a number of well-known equations to describe their behavior. We complemented the qualitative versions of the equations with deviation models as described above.

On the other hand, the engine model, particularly the combustion process itself, is an example of a component where it is not possible to generate the model from a set of equations. However, even for such an ill-specified component, knowledge might exist concerning the qualitative effects of certain variable variations. For example, a decreased amount of injected fuel will lead to a decrease in engine performance, provided that all other inputs are the same. We can capture such types of relations using descriptions of qualitative deviations. For example, in the case of a too low intake air oxygen rate AO (e.g. due to a too high exhaust gas recirculation rate), the resulting combustion energy E and the exhaust gas oxygen rate EO will be lower, but the carbon emissions EC higher than normally (due to incomplete combustion in this case). The engine model in our prototype encodes various such situations, but is incomplete, as only the effects of certain variable deviations are considered.

Our notion of a deviation does not fix the reference which is used for comparison. Using the nominal behavior as a reference is a natural option for diagnosis, but not the only possibility. What is exploited to construct deviation models from the underlying equations is that the same set of constraints holds for both the actual and the reference values. This is why we can also interpret two time points t_0, t_1 in such a way that previous observations at t_0 represent a reference value and the actual observations at t_1 are to be compared with this reference:

$$\triangle x(t_1, t_0) = x(t_1) - x(t_0). \tag{9}$$

Although the interpretation of a "deviation" is now a (significant) change in time, the deviation models hold for any pair of time points t_0, t_1. Capturing such temporal changes in the behavior models and the observations of the device is a way to avoid having to calculate derivatives based on noisy signals, which are present in our application. The idea is that while derivatives at a single point in time might be unobservable, deviations of signals over some (significantly large) time interval can be observed. If we have a model that relates such deviations of variables (or their integrals) over time, this can be used for diagnosis in the state-based framework.

This presupposes that we are adequately supplied with observations about deviations over time. In our prototype, such deviations over time are computed from the observed signals. Currently, this is based merely on a pre-defined, restricted schema for determining past time points t_0 as suitable reference. Section 5 discusses possible improvements.

4. Evaluation on the Demonstrator Vehicle

4.1. Prototypic On-board Diagnosis System

The software for the on-board diagnosis prototype consists of the following components:

- a module for the conversion of raw signals into qualitative observations, and
- a model-based run-time system that performs diagnosis on the basis of these observations.

The former point comprises a component for the conversion of quantitative signals into qualitative values and qualitative deviations as described in section 3.3. Each time a change of observed variables or their deviations occurs on the qualitative level, a new vector of observations is created and handed to the diagnosis engine.

For the latter point, we used components of the commercial RAZ'R system ([14]) that offers a development environment for diagnostic models as well as a run-time version of a consistency-based diagnosis engine. The diagnosis engine performs behavior prediction using the qualitative observation vectors and the model in a state-based framework, where it re-uses, if possible, existing predictions from previous observation vectors (section 3.2). The diagnostic result is derived as the combination of the results obtained for each individual observation vector.

4.2. Demonstrator Car Set-Up

In the VMBD project, a Volvo 850 TDI demonstrator car was made available for hands-on experimentation with the DTI application. Failures can be induced in the car during various operational conditions of the engine with the model-based diagnosis system running, and the results can be compared with the conventional diagnostic capabilities of the control unit. The various failures in the demonstrator car can be adjusted by potentiometers and triggered by switchboards from inside the passenger compartment (see figure 5). A pneumatic leakage, for example, is simulated by additional valves opened and closed by electrical switches.

Fig. 5. View of the Volvo Demonstrator Car showing the notebook connected to the ECU. The glove compartment (behind) contains the switchboard for controlling the built-in faults.

For these experiments, additional interfaces and devices had to be installed. Given that the ordinary control unit and its diagnosis algorithms should work in parallel without interruption (for safety reasons), and the model based diagnosis prototype needs the same information about the sensor signals as the current ECU, the architecture of the measurement facilities was chosen as shown in figure 6.

At present, control units still have rather limited computing power which prevented us from integrating the model-based diagnosis system within the ECU software. To circumvent this restriction, a so-called application control unit was used in the demonstrator. Application control units are normally used for calibration of ECU software for a specific vehicle type and are equipped with special dual-ported memory chips such that in principle all variables and signals of the control unit are accessible in real time, without interfering its normal operation. The data of the vehicle is interfaced to the model-based diagnosis prototype, which is running on a portable PC inside the passenger compartment. In figure 6, ETK is a hardware interface closely attached to the application ECU providing access to its controller bus. MAC is a protocol conversion box which stores the information gathered from the ETK, while VS100 is a commercial tool that car suppliers use for acquisition, storage, interpretation and display of control unit data. It runs on the same portable PC as the on-board diagnostic prototype. The AD-Scan device and the PC Tester allow to read in further signals (dotted lines) from additional sensors or workshop equipment for the purpose of off-board diagnosis in the VMBD project.

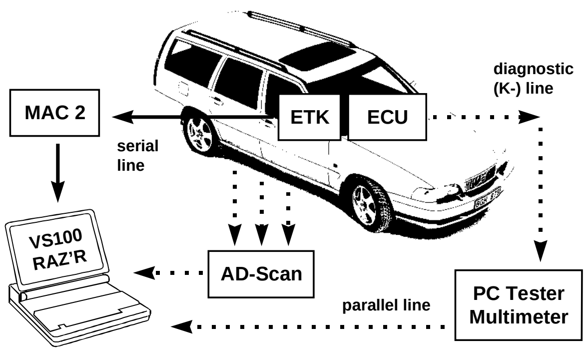

Fig. 6. Architecture for data acquisition in the demonstrator car.

Although this means that the model-based diagnostic software is not really running on-board within the ECU, we consider this solution adequate for our case studies since it provides all important constraints except the space and computing power limitations of the ECU. This aspect is beginning to be more and more relaxed in practice, anyway.

4.3. Diagnostic Scenarios

We were particularly interested in failures that cannot be captured or are hard to capture by traditional on-board diagnosis. Since increased legislative and customer demands have lead to new requirements especially for aspects related to emissions and performance of the system in the Volvo car, we concentrated on effects that involve incomplete fuel combustion and increased carbon emissions due to an excessive quantity of fuel injected or insufficient airflow to the engine (called "black smoke" problems).

One scenario in the demonstrator car consists of a leakage in the air hose between the turbine outlet and the engine intake manifold. The scenario was realized in the car by installing an electric motor which opens a valve to release pressure from the inter-cooler system via a 12mm opening. If the leakage is opened, air (oxygen) mass is lost after having passed the air mass sensor. The fuel quantity calculated by the control unit which is based on this signal will therefore be too high for the actual amount of oxygen in the combustion chamber. This leads to incomplete combustion of the diesel fuel, which causes increased carbon emissions in the exhaust gas (due to non-burnt particles) and reduces the torque of the engine. This effect is, depending on the driving condition, perceivable for the driver as black smoke emerging from the exhaust system.

In another scenario, a wrong flow from the exhaust gas re-circulation (EGR) system occurs due to a faulty signal or mechanical failure in the EGR valve. The real fault installed in the car consists of a switch used to control a magnetic valve that allows ingress of atmospheric pressure in the EGR valve, thus causing it to open outside its normal operating region. The rest of the scenarios involved faults in the boost pressure sensor, airflow sensor and engine temperature sensor. These faults are injected in the car by electrically manipulating the respective signal to the control unit.

4.4. Measurements

From the available control unit data, the following subset of signals was fed to the prototype for diagnosing the described scenarios:

- atmospheric pressure sensor signal
- boost pressure sensor signal
- mass airflow sensor signal
- engine speed sensor signal
- duty cycle of the turbo control valve
- current fuel quantity injected

The on-board diagnosis prototype uses only these control unit signals, and no further signals from additional sensors. The frequency at which the control unit reads the signals from the sensors varies with the speed of the engine, therefore the time points at which observations occur are not evenly distributed.

4.5. Diagnostic Results

In our experiments, we have injected the above failures in various operating modes of the car, such as idling, driving with constant speed, full acceleration and stalling. We present results for the leakage scenario in more detail. The leakage has an effect only if the pressure in the hose (i.e. the boost pressure) is significantly different from the pressure outside (i.e. the atmospheric pressure), which means that the failure is not visible e.g. during idling.

Figure 7 shows the diagnostic results for a slowly opening leakage during stalling the engine. The upper part of the window shows the control unit signals listed above. The measurement runs for 9.75 s and yields 1064 quantitative observation vectors.

The signal transformation module reduces them to only 12 qualitative observation vectors (indicated by the small "peaks" at the base of the signal window). Based on these observations and the model, the runtime system successively reveals three sets of conflicting assumptions:

- Junction1, Intake Turbine, Junction3, Engine, Airflow Sensor, Junction2, Pressure Sensor,
- Junction1, Intake Turbine, Junction3, Engine, Airflow Sensor, Junction2, Speed Sensor,
- Junction3, Engine, Speed Sensor, Pressure Sensor.

Here, the component names stand for the assumption that the respective component is working correctly; i.e. at least one component in each of the above sets must be faulty. The three different conflicts combine to two single fault hypotheses and a number of multiple fault hypotheses:

- Junction3
- Engine
- Pressure Sensor, Junction1
- Pressure Sensor, Intake Turbine
- Pressure Sensor, Airflow Sensor
- Pressure Sensor, Junction2

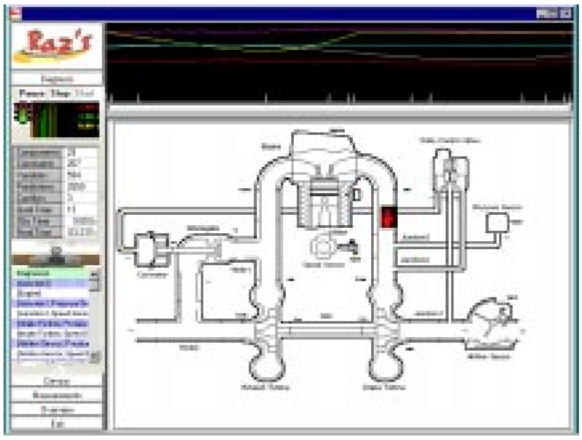

Fig. 7. Screenshot of the model-based run-time diagnosis prototype for the DTI turbo control subsystem.

- Speed Sensor, Junction1
- Speed Sensor, Intake Turbine
- Speed Sensor, Airflow Sensor
- Speed Sensor, Junction2
- Speed Sensor, Pressure Sensor

The two single fault hypotheses contain the component where the failure was actually induced ("Junction3", see mark in 7 within the window depicting the system structure). The runtime for the example (on a Windows/Pentium PC) is 25.85 seconds without using temporal caching, and 3.91 seconds if temporal caching is activated. This means that, for this example, the performance of the on-board system is in the order of magnitude of real-time.

Similar results where achieved for the rest of the scenarios. Table 1 summarizes the results. Note that the current control unit software, based on the same signals, is not able to detect any of the above failures. Because some failure effects are noticeable only during certain operating conditions, the diagnosis system cannot always determine a unique diagnosis, but rather yields a number of hypotheses as in the example above. E.g. for the scenario with the boost pressure sensor out of tune, the diagnosis system yields two conflicts and outputs a list of three single faults which contain the boost pressure sensor as one possible candidate, but also other components that together could account for the same symptoms.

In these cases, knowledge about the behavior of faulty components, i.e. fault models, could be used to further constrain the set of diagnostic candidates. So far, only models of correct behavior have

Table 1

Summary of diagnostic results for the on-board prototype (showing typical instances of measurements).

Scenario	Fault detected	No. of (single) component fault hypotheses	Realtime	No. of quantitative vectors	No. of qualitative vectors generated	Runtime
Air intake pipe leakage	yes	2	9.75 sec	1064	12	3.91 sec
Boost pressure signal too high	yes	3	12.20 sec	1417	10	5.04 sec
Airflow sensor signal too high	yes	6	14.40 sec	1887	22	6.65 sec
EGR valve opens outside normal operating region	yes	7	18.61 sec	1676	4	2.43 sec
Engine temperature signal too low	no	–	21.08 sec	2003	6	2.22 sec

been used for the diagnostic experiments. At least in some cases, there is evidence that fault models could be useful to partially compensate for the limited observability, and thus to further restrict the diagnostic candidates.

5. Discussion

5.1. The Success Story

The demonstrator illustrates that model-based techniques are suited to address several requirements that are important to, but not limited to, car diagnosis:

- the variant problem by compositional modeling and model-based generation of diagnostic solutions,
- safety criticality by a systematic approach and completeness of results w.r.t. the model,
- dynamic systems by modeling changes over time as temporal deviations and by performing state-based diagnosis (enabled by sufficiently dense measurements),
- limited system information by modeling complex elements like the combustion process at a qualitative level and by an appropriate qualitative abstraction of noisy signals,
- real-time requirements by processing only qualitative changes in the measurements and by applying state-based diagnosis with temporal caching.

Which perspective is opened by the demonstrator? Of course, memory on current control units is limited. But this is going to change very soon, and the concept of a so-called “car PC” ([10]) may be approaching even faster.

But the work on the prototype as part of our efforts towards deploying model-based diagnosis as a technology in the automotive industry has increased our awareness of a number of theoretical, methodological, and organizational problems that will have to be solved in order to improve the utility of our technology. In the following, we discuss some problems and shortcomings related to the demonstrator and the development process before trying to draw some general conclusions about the current state and the future of model-based diagnosis.

5.2. Diagnosis - More than Candidate Generation

What our current demonstrator, like most of the implemented diagnosis systems, does is candidate generation: Given a system model and a set of observations, what are the diagnostic hypotheses that can be derived from them?

More specifically, it performs fault localization, i.e. diagnostic hypotheses are (sets of) possibly faulty components. If we add fault models in a new version of the demonstrator, it will realize fault identification: hypothesizing possible (sets of) specific component faults. While this is a useful goal to achieve, there is still something crucial lacking: an explicit representation and exploitation of the goal of the candidate generation step, namely therapy: Re-establish the functionality of the respective system if or as far as possible within given conditions.

Again like many other diagnostic systems, the demonstrator terminates fault identification when it has generated a set of assignments of behavior modes to components that are the most plausible ones under some criterion. This appears to be fairly appropriate if therapy consists of component replacement, as in workshop diagnosis. But on-board diagnosis has a totally different goal: to determine the appropriate recovery actions. Today, there are a number of predefined reactions of the control unit, ranging from stopping operation of the vehicle via special control schemes for modified continued operation to just turning on a warning lamp. Which one is appropriate is not so much dependent on the defect component (if any), but on the type of fault or disturbance. In particular, this means that the severity of the fault (in terms of potential damage and threat to safety) rather than its likelihood usually determines whether or not it has to be considered which is likely to be in conflict with focusing and control schemes of current diagnosis systems. Discrimination has to pursue a specific goal, and our current theories and systems have no systematic way of expressing this and allowing to control the diagnosis process under this goal.

Interpreting this finding in a more general way, we have to notice that even fairly well-established pieces of our technology have limited means for reflecting the practical context and conditions of the diagnostic tasks they might be used for.

5.3. Modeling - More than Libraries

Building the model library was one of the major tasks and achievements in the project. For our demonstrator, it was mainly carried out by AI researchers with input from engineers from the automotive companies. Actually, we consider it a remarkable aspect of the model-based technology that it enabled non-experts in car technology to produce a diagnosis system whose results conform with conclusions of experts!

The resulting models seem to do their job for the demonstrator. But to what extent will they do their job as re-usable elements of a model library when applied to a different problem?

In our solution, every non-zero value for $x_{act} - x_{corr}$ is considered a deviation. However, what we are interested in are significant deviations, where "significance" is determined by the absence or existence of functional disturbances. We reckon that it is not possible to derive well-defined and useful deviation models capturing this intuition purely within the qualitative domain of signs. The reason is that significance in the sense stated above will vary with the function and system considered. A certain deviation of an internal parameter may significantly disturb one function or operating mode, while being irrelevant to another one.

For example, the connections to the turbo control valve branching off the air intake pipe of the engine are thin hoses with only some millimeters in diameters, whereas the intake pipe itself measures several centimeters in diameter. A leakage in the pipe connecting to the turbo control valve, therefore, would not affect the boost pressure in a significant way. We can express this by allowing a positive pressure deviation in the small pipe to be compatible with a zero pressure deviation in the bigger pipe. We introduced in our model a special type of separation component that describes such constraints, which can be interposed e.g. at junctions of pipes with significantly different diameters.

Clearly, this is only an ad-hoc solution. Yet, the problem to characterize what makes a certain distinction to be significant is a fundamental issue, as using qualitative models (and signal abstractions) is crucial to achieving the performance required for on-board diagnosis. Whether or not a distinction in the domain of a variable or parameter is significant cannot be determined locally, but depends on the context, for instance on whether or not it causes another component to change its operating mode or leads to a disturbance of the desired function.

There is a related problem. The device under consideration (like other automotive systems) comprises a discrete systems (the control unit software), standard physical components (electrical, hydraulic, pneumatic ones), and elements for which no rigorous mathematical model can be derived (the combustion engine). Consequently, models of the various parts come at different granularity, ranging from continuous-valued variables governed by a set (differential) equations for the physical part to discrete control signals and characteristic lines or characteristic maps for components or parts of the system for which no rigorous mathematical model exists. What we would like to have is support to the composition and smooth integration of such independently developed, "hybrid" component models, for instance by determining tailored qualitative distinctions in the domain of continuous variables reflecting the given distinctions in other parts of the model and its specific structure.

From a broader perspective, the problem of deriving models that perform exactly the relevant inferences is not only a matter of the granularity of the variable domains. More generally, it has to be decided which set of phenomena have to be covered. For instance, models of some pipes and hoses in our application example need to capture the transportation of oxygen, carbon oxide, etc. However, this is irrelevant e.g. to the pipe that connects the turbo control valve to the waste gate valve, for which only pressure matters. Again, which features to include in a model, depends on the context and the task and is not straightforward. Is it necessary or not to propagate information about air temperature and oxygen rate from the air intake to the turbo control valve?

We have been able to develop satisfactory models for our demonstrator. We can do it for another one and might also be able to improve their utility for a broader class of applications. However, it would be our task and a time-consuming one. This is another serious obstacle in the process of deploying the model-based diagnosis technology.

5.4. Signal Interpretation - More than Signal Pre-processing

The potential of the temporal deviation models we used and the algorithms for the computation and exploitation of temporal deviations need to be further explored, both empirically and in their mathematical foundations. In particular, the possibilities of deriving such deviation models from arbitrary differential equations in a systematic way, as well as the selection of appropriate time points for comparison, still need theoretic investigation.

To use such types of models has an impact on the signal transformation component, too, since it requires this component to compute the deviations of (integrals of) signals also.

Under a broader perspective, this highlights a fundamental issue. On the one hand, the theory of signal processing offers a wide variety of filters and algorithms in order to interpret measured data. On the other hand, employing this in the context of qualitative models raises several problems. First, one is often left with the problem which kind of method and which parameters to choose for a specific application and model. A second, deeper problem is the fact that these methods hardly reflect the goal the resulting signal will be used for. Model-based reasoning is different from the traditional applications of these methods, like e.g. smoothing a signal such that its display on the dashboard appears less erratic, or fault detection based on predefined thresholds of signals or residuals. It is different because it becomes necessary to express an explicit “belief” in the observations once they are used for sophisticated logic-based reasoning, and in particular for refuting behaviors in diagnosis.

As a consequence, if we want to move away from hand-crafted solutions of coupling model-based systems with sensor measurements, we must find ways to (automatically) harmonize the granularity of the observations, the granularity of the model, its underlying assumptions and the task it is used for. While approaches exist in each of these directions, a general answer to this problem is still beyond the current state of the art.

5.5. Transfer to Industry - The Inertia of Current Practice

So far, we pointed out some theoretical and technical problems that need to be addressed to promote technology transfer. But even if this were not the case, actually installing the technology in control units on vehicles does not appear to be the next feasible step. There are major obstacles arising from organizational and “cultural” problems. As we pointed out, our technology is offering a number of very attractive and cost-effective features and benefits to industry. But with this offer, we are not entering a “white spot” on the map of current industrial work processes during the product life cycle. There are probably hundreds of thousands of experts on car diagnostics successfully solving problems, including mechanics in the workshops and engineers designing on-board systems, and they do so in established work processes, with a certain education, using particular tools. In our domain, we are facing mainly experts with an engineering education and equipped with some experience.

The point is that we are not simply offering a new tool that simplifies or improves some tiny step in their work, such as a better data base interface, a faster version of a simulator, something they can easily switch to and benefit from immediately. Model-based systems interfere with their knowledge and crucially with the content of their expertise. That these systems change the work process dramatically, is actually one of their great potentials, but first, exactly this forms a serious obstacle to their introduction. Of course, engineers are accustomed to perform modeling. But even this may turn out as a problem rather than an advantage, because it usually means handcrafting special-purpose, often black-box, mathematical and numerical models. Building model libraries with context-free behavior models at a conceptual level for automated model composition is a quite different task.

As a result, the confidence in model-based systems, the preconditions in terms of education and organization of work have yet to be developed, and this is something that needs time, even if there is a serious commitment of management.

The consequence is that the solutions-in-principle found in the demonstrators will have to be turned into tools that support the actual work process, for instance for analysis of diagnosability, sensor placement, and FMEA. This also involves the integration with other tools, such as simulation systems and CAD tools.

6. Future Work

For the involved car supplier, the results give rise to expectations of improving the diagnosis in vehicles in essentially three respects:

- in automatically generating fault candidates and decision trees for diagnosis and repair in the workshops,
- in assisting the design phase of automotive subsystems by providing methods and tools for verification, testing and design for diagnosability,
- in integrating model-based diagnosis techniques in the on-board equipment, as soon as sufficiently powerful processors in the ECU are available.

For OCC'M Software, the attempt to really put the technology to work generated a number of requirements that will be considered in the future development of the RAZ'R software. This includes support during the modeling phase (e.g. re-use of model fragments below the level of component models) and features of the runtime system (e.g. the interface to the measurement acquisition module).

The results so far give Volvo reason to expect a commercialization of the developed techniques supported by commercial tools supplied by professional software companies. Despite the deep involvement of the companies, the success of the case studies in this project heavily relied on the availability of AI researchers who are familiar with the technology. For a really broad application within an industrial setting, in particular the task of modeling cannot be left to software engineers. Rather, it has to be integrated with the extensive work on modeling carried out by the engineers during different phases in the product life cycle. A coherent framework for different types of modeling and product representations is a worthwhile target, and could to a certain extent fulfill the company's needs for some kind of standard in the domain. Volvo has reason to believe that the techniques developed in this project have the potential to become the basis for such a standard, thus making a contribution to corporate knowledge management in this area. The solutions-in-principle found in the demonstrators will have to be turned into tools that support the actual work process, for instance for analysis of diagnosability, sensor placement, and FMEA. This also involves the integration with other tools, such as simulation systems and CAD tools. In the near future, Volvo expects to have to create new work processes internally as well as new methods for cooperation with suppliers to make full use of the possibilities.

Actually, the results of VMBD lead to a follow-up project that again joins a number of car manufacturers, a supplier of vehicle subsystems, a software company (OCC'M Software), and research groups from European universities. It aims explicitly at turning the technology into tools that can be used in the current design process of on-board control systems (including hardware and software). This appears to be a natural next step in deploying the technology in the automotive industry in the near future, rather than trying to install model-based diagnosis directly on the car at this time.

7. Summary

The demonstrator described in this paper proved the feasibility of the technology in the automotive domain. In particular, it can provide a basis for a systematic and cost-effective approach to creating diagnostics for car subsystems, and has the potential to improve the quality of diagnostics by handling fault situations that are not covered by current on-board diagnostics, or by improving fault identification up to a point where more specific recovery actions can be chosen. Of course, such conclusions are still preliminary and have to be confirmed and exploited by further studies.

The other two demonstrators that have been realized in the VMBD project have focused on different problems in connection with vehicle diagnosis, namely off-line compilation of a model to use it in an on-board environment and experimenting with additional sensors in the diesel injection system ([4]) and using a numerical approach for off-board diagnosis of faults in the hydraulic and mechanical parts of an automatic transmission system ([2]).

These outcomes of the project as a whole represent a major step in the transfer of model-based diagnosis techniques into the automotive industry. Focusing on the application also shed a light on preconditions and tasks of the technology transfer process. The experience gained provides a basis for decisions about the introduction of the technology by the industrial users, and helps the technology supplier, OCC'M, to improve and extend the tools.

Acknowledgements

We would like to thank our colleagues, especially Oskar Dressler, Ulrich Heller, Andreas Malik, and Björn Svensson, and all partners in VMBD for their valuable contributions. This work was supported in part by the Commission of the European Union through the BriteEuRam III Programme (♯ BE 95/2128) and by the German Ministry of Education and Research (♯ 01 IN 509 41).

References

[1] S. Addanki, R. Cremonini, J. S. Penberthy: Graphs of models, *Artificial Intelligence* 51(1-3):145–177, 1991

[2] P. Bidian, M. Tatar, F. Cascio, D. Theseider-Dupré, M. Sachenbacher, R. Weber, C. Carlén: Powertrain Diagnostics: A Model-Based Approach, Proceedings of ERA Technology Vehicle Electronic Systems Conference '99, Coventry, UK, 1999

[3] Robert Bosch GmbH (ed.): Automotive Handbook, 4th edition, Society of Automotive Engineers, Warrendale, PA, 1997

[4] F. Cascio, L. Console, M. Guagliumi, M. Osella, A. Panati, S. Sottano, D. Theseider-Dupré: Generating on-board diagnostics of dynamic automotive systems based on qualitative deviations, *AI Communications* 12(1):33–44, 1999

[5] L. Console, W. Hamscher, J. deKleer (eds.): Readings in Model-based Diagnosis, Morgan Kaufmann Publishers, San Mateo, CA, 1992

[6] H. de Jong: Computer-Supported Analysis of Scientific Measurements, Ph.D. Thesis, University of Twente, Netherlands, 1998

[7] O. Dressler: On-line Diagnosis and Monitoring of Dynamic Systems based on Qualitative Models and Dependency-based Diagnosis Engines. In W. Wahlster, Editor, In Proceedings of the European Conference on Artificial Intelligence (ECAI-96), pages 461–465, John Wiley & Sons, 1996

[8] O. Dressler, P. Struss: The Consistency-based Approach to Automated Diagnosis of Devices. In: Brewka, G. (ed.), Principles of Knowledge Representation, CSLI Publications, Stanford, pp. 267–311, 1996

[9] B. Falkenhainer, K.D. Forbus: Compositional modeling: finding the right model for the job, *Artificial Intelligence* 51(1-3):95–143, 1991

[10] Intel Connected Car PC Technology Website, see `http://intel.com/technology/carpc/index.htm`

[11] E. Loiez, P. Taillibert: Polynomial Temporal Band Sequences for Analog Diagnosis, International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan, p. 474–479, 1997

[12] A. Malik, P. Struss: Diagnosis of Dynamic Systems Does Not Necessarily Require Simulation. In Workshop Notes of the 10th International Workshop on Qualitative Reasoning (QR-96), AAAI Press, 1996

[13] California's OBD–II regulation, section 1968.1, title 13, California code of regulation, Resolution 93–40, 1993

[14] Occ'm Raz'r Version 1.6, Occ'm Software GmbH, see `http://www.occm.de/`

[15] M. Sachenbacher, P. Struss: Fault Isolation in the Hydraulic Circuit of an ABS: A Real-World Reference Problem for Diagnosis. In Working Papers of the 8th International Workshop on Principles of Diagnosis, Mont-Saint-Michel, France, 1997

[16] M. Sachenbacher, A. Malik, P. Struss: From Electrics to Emissions: Experiences in Applying Model-based Diagnosis to Real Problems in Cars, Working Papers of the 9th International Workshop on Principles of Diagnosis (DX-98), Cape Cod, USA, pp. 246–253, 1998

[17] P. Struss: What's in SD? Towards a Theory of Modeling for Diagnosis. In: Hamscher et al. (eds.), Readings in Model-based Diagnosis, Morgan Kaufmann Publishers, pp. 419–450, 1992

[18] P. Struss: Fundamentals of Model-Based Diagnosis of Dynamic Systems. International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan, pp. 480–485, 1997

[19] P. Struss, O. Dressler: Physical Negation: Integrating fault models into the general diagnostic engine, Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-89), Detroit, USA, pp. 1318–1323, 1989

[20] P. Struss, M. Sachenbacher: Significant Distinctions Only: Context-Dependent Automated Qualitative Modeling, Working Papers of the 13th International Workshop on Qualitative Reasoning (QR-99), Loch Awe, Scotland, 1999

[21] M. Sachenbacher, P. Struss: Automated Qualitative Model Abstraction – Theoretical Foundations and Practical Results, to appear in Working Papers of the 14th International Workshop on Qualitative Reasoning (QR-00), Morelia, Mexico, 2000

[22] P. Struss, A. Malik, M. Sachenbacher: Qualitative Modeling is the Key to Automated Diagnosis, Proceedings of 13th World Congress of IFAC, San Francisco, CA, USA, Pergamon, 1996

[23] P. Struss, M. Sachenbacher, F. Dummert: Diagnosing a Dynamic System with (almost) no Observations. In Workshop Notes of the 11th International Workshop on Qualitative Reasoning (QR-97), Cortona, Italy, pp. 193–201, 1997

[24] D. Theseider Dupré, A. Panati: State-based vs. simulation-based diagnosis of dynamic systems. In Workshop Notes of the 9th International Workshop on Principles of Diagnosis (DX-98), Cape Cod, USA, 1998

[25] D. Weld, J. de Kleer (eds.): Readings in qualitative reasoning about physical systems, Morgan Kaufmann Publishers, San Mateo (CA), 1990