

From Electrics to Emissions: Experiences in Applying Model-based Diagnosis to Real Problems in Real Cars

Martin Sachenbacher, Andreas Malik and Peter Struss

Technische Universität München
Department of Computer Science
Orleansstr. 34, D-81667 Munich
Germany
{sachenba, malik, struss}@in.tum.de
<http://www9.informatik.tu-muenchen.de/MQM/>

Abstract

This paper reports on experiences and preliminary results gained in applying model-based diagnosis on diesel engine management systems in passenger vehicles. It states some major requirements that have been identified in this application. In particular, it is crucial to reason across different physical domains, to exploit emission-related symptoms like increased carbon emissions in the exhaust gas for identifying failures in upstream subsystems. The main foundation of our approach is to use qualitative deviation models, which serve as a coherent modeling paradigm for the different domains. Second, the models are used in a state-based diagnosis framework, which checks consistency of the model with observed states only, to avoid simulation of the system's behavior over time. We present preliminary results of putting this approach into practice for real problems in demonstrator vehicles. The evaluation provides empirical evidence for the utility of model-based techniques to improve the discriminative power of automated vehicle diagnosis both in on-board and off-board situations. The experiments also raised a number of open issues that are subject to further research.

Introduction

In a collaboration between Robert Bosch GmbH and the Technical University of Munich, a number of case studies on applying model-based diagnosis techniques to diagnosis of car subsystems have been carried out ([Struss et al. 95], [Struss et al. 97]). In our presentation, we report on work that has been done in the framework of the Brite-EuRam project VMDB (Vehicle Model Based Diagnosis) ([Tatar 97]). This project aims at demonstrating the utility of the technology for real problems both in on-board and off-board situations by installing model-based diagnosis systems on demonstrator vehicles.

Our work in VMDB focuses on diesel injection systems. The purpose of these systems is to supply the diesel engine with the right amount of air and fuel at the right time. To perform diagnosis on such systems, electrical sensors and actuators, air and fuel subsystems, the combustion process and, to some extent, the software of the electronic control unit have to be modeled. A feature of this application is that

there are very few sensors available for diagnosis, especially in on-board diagnosis. One of the focal points of work in VMDB is to investigate if better discrimination of faults in on-board diagnosis is possible in order to allow for a more appropriate choice of recovery actions. For example, if an implausible value of fuel pressure has been detected, the control unit currently has no choice than to stop the engine. However, if a more specific failure cause could be identified, less restrictive actions like limiting the engine speed could be taken, which maintain mobility and would allow at least for reaching the next repair workshop. We are concerned with two different types of injection systems in the VMDB project, which vary in the technical realization of certain subsystems: the so-called Common Rail diesel injection system and the distributor-type diesel injection. However, more important is the fact that these two applications pose different requirements for diagnosis and failure analysis.

The main foundations of our approach to automated diagnosis are

- qualitative deviation models, i.e. component models that state and propagate (directions of) deviations from some nominal or reference behavior (which is possibly left unspecified), even across different domains, and
- state-based diagnosis, i.e. even for dynamic systems, diagnosis is based on checking consistency of the model with observed states only, as opposed to simulating its evolution over time.

Our preliminary implementation is based on RAZ'R, a commercial diagnosis system produced by OCC'M Software GmbH. It provides a graphical user interface for modeling, prediction and diagnosis in the framework of GDE ([de Kleer and Williams 87]).

The next section describes the application domain in more detail and summarizes the key requirements that have been identified in this application. Then we list our choice of techniques to meet these requirements and present part of the models, focusing on pneumatic components and the diesel engine. The last section illustrates preliminary results of two case studies with scenarios from the air supply subsystem and the low-pressure fuel delivery subsystem.

The Application Domain

The automotive drive train is the essential part of a passenger car. It converts energy in chemical form into mechanical energy and provides the thrust and tractive forces required for vehicle motion.

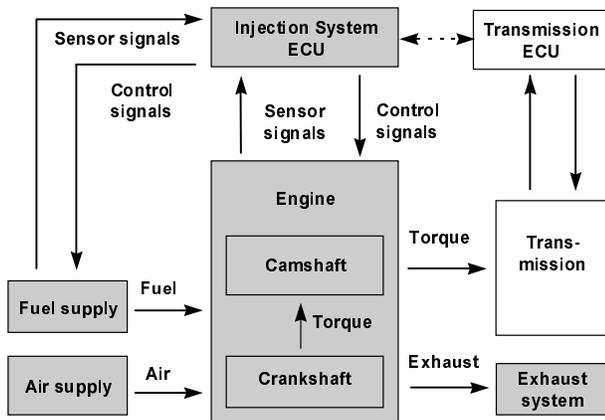


Figure 1: Schematic view of the vehicle drive train

The drive train consists of the engine, the engine management system and the automatic transmission, which is the subject of work of another subgroup in VMBD. Figure 1 shows a schematic view with the elements of the injection system highlighted. The diesel injection is a complex, modular system which consists of electronic, electrical, mechanical, hydraulic and pneumatic components. The electronic control unit (ECU) receives signals from various sensors, including engine speed, air and fuel temperature. Based on this information, it controls actuators in hydraulic and pneumatic subsystems which provide air and fuel as inputs to the engine. The combustion in the engine produces torque but also exhaust gases as outputs.

Diagnostic Tasks and Goals

The project distinguishes between on-board and off-board diagnosis, which pose different requirements on fault localization and identification and differ significantly in terms of available measurements and reproducibility of conditions:

On-board diagnosis. On-board diagnosis aims at selecting appropriate recovery actions. The injection system continuously monitors part of the sensor signals. Current on-board diagnosis can detect faults on the basis of predefined range and plausibility checks for signals. It will then perform built-in recovery actions that range from minor performance reductions to full engine stop and depend on the assumed failure and the expected failure effects. However, due to the scarcity of sensors, in most cases the control unit fails to discriminate among the different possible causes that lead to the failure.

Consequently, the system often applies a more restrictive recovery action than would be necessary.

One of the aims in the VMBD project is to use model-based diagnosis to achieve better fault discrimination and to determine more appropriate recovery actions, which helps to improve the availability of the vehicle.

Off-board diagnosis. Off-board diagnosis aims at repair and, more specifically, replacement of components. In contrast to on-board diagnosis, the task is to localize the fault down to the smallest replaceable unit. Therefore, off-board diagnosis requires a different level of accuracy (granularity) for fault localization. Off-board diagnosis can make use of more observations, using advanced service testers and human observations. Off-board diagnosis can also propose distinguishing tests, in order to partition the system, and use more information about the application context (e.g. engine load). In contrast to on-board diagnosis, where measurements "come for free" (at least at runtime), obtaining proper observations in the workshop has its, often non-negligible, price. Saving such costs for testing and, for instance, disaggregation of systems, establishes a justification for using more elaborate diagnostic algorithms such as fault identification based on fault models.

Requirements and Problems for Modeling and Diagnosis

This application domain provides a number of challenges for modeling and automated diagnosis (see also [Sachenbacher and Struss 97] for a description of similar problems arising in off-board diagnosis for an anti-lock braking system):

Multi-domain model. The diesel injection system, like most car subsystems, involves different domains of physical systems: electrical, electronic, hydraulic, pneumatic and mechanical ones. This poses the problem of creating models of such components that can be used in one common framework.

Engine model. To establish a link between symptoms such as increased emissions or reduced performance and failures in the subsystems, it is necessary to model the behavior of the engine. However, the available knowledge about its behavior is usually incomplete and precarious. Other than for standard electric or hydraulic components, there exist no general mathematical models one could start from.

Control unit model. Even if faults in the control unit itself are not considered, the behavior of the control unit has to be modeled in some cases, since it is part of the feedback loop comprising sensors and actuators.

Limited measurability. Very few sensors are available in these systems. This is especially true for the high pressure hydraulics of the Common Rail system, which contains only one pressure sensor. In addition, the context in which a car is operated (e.g. road and weather conditions, load) is highly dynamic, uncertain and often neither measurable nor reproducible in the workshop. The main consequences are

noisy signals and rather qualitative and vague symptom descriptions. Models supporting diagnosis have to be capable of processing such information.

Dynamic systems. The injection system has internal states depending on previous inputs, thus being an example of a dynamic system.

Controlled, feedback systems. The effects of faults may be compensated by control. For example, the control unit of the Common Rail injection system may, to a certain extent, compensate for leakages in the high-pressure hydraulics. Therefore, faults may be observable in transition phases only (e.g. through changed time constants or oscillatory behavior). It becomes a problem especially in combination with low measurement quality, that is, if measurements within feedback loops are sparse or over time.

Real-time requirements. On-board diagnosis must come up with a conclusion before a shut-down of the injection system is necessary to prevent safety-critical situations or severe system damage or to comply with legal restrictions. Although this problem is also addressed in the context of VMBD, we will not further discuss it here.

Variant problem. Like many other automotive subsystems, diesel injection systems come in different variants. They 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 in different constructive details. Generating specialized diagnostics for all variants demands for computer support or even an automated solution.

Modeling methodology. Building the appropriate models is still carried out mostly by computer scientists with a background in engineering. However, modeling techniques should be designed so as to make them accessible to a wider class of users, in particular to engineers in industry, and they should allow for the derivation of suitable models from existing mathematical models.

Our Answers

Our choice of techniques in order to meet some of these requirements is as follows:

Model-based diagnosis. One reason for the restrictiveness of recovery actions is that current on-board diagnosis cannot precisely identify faults because the analytic redundancy in the signals is not used systematically. A model of the system captures such redundancies, and it can make maximum use of this information for discriminating failures.

Compositional models. Automotive subsystems can be decomposed into components and aggregates. Compositional modeling provides a way for generating the diagnostic system (or, at least, a basis for it) from a structural model (a "blueprint") automatically, which helps remedying the variant problem.

Qualitative models. Qualitative descriptions reflect the nature of available observations, e.g. "black smoke" symptoms in the distributor-type injection system. By covering entire classes of behaviors, they also help to keep the library of model fragments manageable.

Order-of-magnitude reasoning. In addition to distinctions expressed in the quantity space of variables, it also turned out necessary to distinguish different orders of magnitude of the relevant physical quantities. As an ad-hoc solution, we introduced a special type of separation component. It explicitly declares two quantities to be of different order of magnitude and thus constrains the way they can influence each other. This component can be interposed e.g. at junctions of pipes with significantly different diameters (Figure 3 shows an example of its application).

Deviation models. In some cases, especially on-board, it appears not to be 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, it may suffice to explain why the pressure in the hydraulic system is (much) higher than it should be, regardless of the actual value of the pressure. Based on a theoretical foundation of such deviations, models have been developed that state and propagate deviations from some nominal or reference behavior (which is possibly left unspecified), even across different domains.

Fault models. One way of trying to compensate for the limited observability is to use a strong model, or more specifically, to use fault models. The idea behind this is to improve fault localization by means of fault identification, which means exonerating certain components through refutation of all faults they can possibly exhibit.

Control unit model and engine model. Qualitative (deviation) models of the ECU software and the diesel engine behavior have been developed to provide the connection between different subsystems.

State-based diagnosis. Since the diesel injection is a dynamic system with several feedback loops, a crucial question is whether diagnosis demands for prediction of behavior over time, i.e. simulation. Frequently, this is taken for granted when dynamic systems are concerned. However, diagnostic results can be obtained based on checking consistency of observed and modeled states only, i.e. without performing simulation, sometimes even without any loss at all (see the recent experimental and theoretical work in [Dressler 96], [Malik and Struss 96], [Struss et al. 97], [Struss 97]).

Data acquisition and interpretation. The key for the state-based approach to diagnosis being successful is to have appropriately complete state descriptions at hand. This contrasts with the fact that there is hardly any information about the dynamics available. It has been shown in [Struss et al. 97], [Struss 97] how in such a situation constraints capturing laws of continuity and integration (called CID-constraints) can be exploited to enhance the observations of system variables by information about derivatives.

[Loeser et al. 98] presents a related approach for the task of verification. The idea is to apply general, device-independent mathematical laws, e.g. the mean value theorem, to construct intermediate states that the system must pass through in order to reach an observed state (or specified state, in the case of verification). These intermediate state descriptions can be checked for consistency with the system model, as opposed to using the laws to simulate the model.

Modeling with Deviation Models

In previous work ([Malik and Struss 96]), models have been developed which capture the deviation of an actual value from its reference. For each variable, this deviation can be represented as $\Delta x := x_{act} - x_{ref}$. This is the foundation of the approach we used in our case studies.

The idea underlying deviation models is to describe deviations of variables which are consistent with a certain behavior model (Figure 2). Deviation models can be derived for any behavior model given as a relation. We assume that the behavior of a component can be specified in terms of local variables $\underline{v} = (v_1, v_2, \dots, v_k)$. Let $DOM(\underline{v}_i)$ denote the domain of a variable v_i . Then

$$DOM(\underline{v}) = DOM(v_1) \times DOM(v_2) \times \dots \times DOM(v_k)$$

defines the space for describing the component behavior. A behavior model can be represented as a relation $R(\underline{v}) \subseteq DOM(\underline{v})$. Let $CMPR_i$ be a space for comparison for each $DOM(v_i)$. A comparison measure δ_i on $CMPR_i$ can be defined as

$$\delta_i : DOM(v_i) \times DOM(v_i) \rightarrow CMPR_i$$

For example, δ_i could be a (qualitative abstraction of) distance: $\delta_i : (x_1, x_2) \mapsto [x_1 - x_2]$, where $[.]$ denotes sign abstraction. We extend the comparison measure to the vector \underline{v} by defining

$$\begin{aligned} CMPR &:= CMPR_1 \times CMPR_2 \times \dots \times CMPR_k \\ \delta &: DOM(\underline{v}) \times DOM(\underline{v}) \rightarrow CMPR \\ \delta((x_1, x_2, \dots, x_k), (y_1, y_2, \dots, y_k)) &:= \\ &(\delta_1(x_1, y_1), \delta_2(x_2, y_2), \dots, \delta_k(x_k, y_k)) \end{aligned}$$

The deviation model of R , denoted R_Δ , describes consistent deviations of the variables v_i :

$$\begin{aligned} R_\Delta &\subseteq DOM(\underline{v}) \times DOM(\underline{v}) \times CMPR \\ R_\Delta &:= \{(\underline{v}, \underline{v}', \underline{\Delta v}) : \underline{v}, \underline{v}' \in R, \underline{\Delta v} \in CMPR, \underline{\Delta v} = \delta(\underline{v}, \underline{v}')\} \end{aligned}$$

The projection of R_Δ on the Δv_i , denoted $proj_{\Delta v_i}(R_\Delta)$, can be viewed as a "pure" deviation model which relates deviations of variables independent of their actual and reference value. This is meaningful only in some cases, for example, if the relation describes a monotonic function such as $x + y = z$. In general, at least information about the actual value will be necessary.

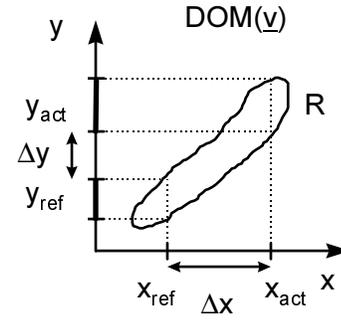


Figure 2: A relational behavior model imposes constraints on the deviations of variables

Applying the Techniques

In previous work ([Struss et al. 95]), we have developed models of electrical components. They are used for describing the behavior of the electrical subsystem of the injection system and provide the link between the control unit and electro-hydraulic or electro-pneumatic sensors and actuators.

For hydraulic and pneumatic components, a number of well-known equations is available to describe their behavior. We complemented the qualitative versions of the equations with deviation models as described above. The domain of each variable is limited to signs, and derivatives of time are added to capture changes in direction. By performing these operations, models capturing the relationships of qualitative values and qualitative deviations are generated from the equations.

Table 1 shows an example of a pneumatic valve with two pneumatic terminals T_1 , T_2 and one mechanical terminal M_1 . In the equations, q stands for flow, p for pressure, s for offset position, and A denotes the cross-sectional area of the valve, whereas α and β are positive constants. The shape of the component is used for composing the model graphically.

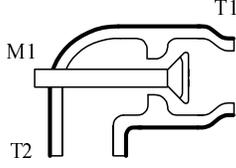
Quantitative Equations	Qualitative Model Fragments
	$T_1.[q] \oplus T_2.[q] = 0$ $T_1.[q] = [A] \otimes (T_1.[p] \ominus T_2.[p])$ $[A] = M_1.[s]$ $T_1.[\Delta q] \oplus T_2.[\Delta q] = 0$ $T_1.[\Delta q] =$ $[A] \otimes (T_1.[\Delta p] \ominus T_2.[\Delta p]) \oplus$ $[\Delta A] \otimes (T_1.[p] \ominus T_2.[p]) \ominus$ $[\Delta A] \otimes (T_1.[\Delta p] \ominus T_2.[\Delta p])$ $[\Delta A] = M_1.[\Delta s]$
$T_1.q + T_2.q = 0$ $T_1.q = A \sqrt{ T_1.p - T_2.p } \cdot$ $\alpha \cdot \text{sign}(T_1.p - T_2.p)$ $A = \beta \cdot (M_1.s)^2$	

Table 1: Derived qualitative model fragments for a pneumatic valve

In addition to such canonically derived model fragments, in some cases we had to take into account further distinctions. For example, thin hoses with only some millimeters in diameters branch off the air intake pipe of the engine, which measures several centimeters in diameter. If the same generic model fragments were used for both types of pneumatic lines, this would result in predicting that e.g. a leakage in the thin hose leads to a pressure drop in the intake pipe. In principle, the inference is correct. However, the effect is insignificant in reality. The order-of-magnitude component mentioned above is applied in this situation.

Qualitative Model of the Diesel Engine

The diesel engine takes air and fuel provided by the hydraulic and pneumatic subsystems as inputs and produces torque and exhaust gases as outputs. Effects such as increased emissions in the exhaust gas or reduced torque form a major part of the observations for the injection system. To exploit these symptoms and to establish a link to failures in the electrical, hydraulic or pneumatic subsystems (e.g. too low pressure of the intake air), it is necessary to model the behavior of the engine.

The engine model is an example of a component where it is not possible to generate the model from a set of equations. This is especially true for the combustion process itself. Its behavior in terms of inputs and outputs can be characterized by a plethora of variables, including fuel mass (MF), air mass (MA), air oxygen rate (AO), degree of fuel atomization (AF), start of injection relative to crankshaft position (IS), duration of injection (ID), combustion energy (E), exhaust gas oxygen rate (EO), nitrous oxides emissions (NO), hydrocarbon emissions (HC), carbon emissions (EC), etc.

There exists no complete set of equations that relates these parameters appropriately, nor would it be particularly useful to describe the engine behavior independently of a specific fuel injection system. However, knowledge exists concerning the 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. Using the qualitative deviation approach, we can capture such types of relations.

ΔAF	ΔMF	ΔMA	ΔAO	ΔE	ΔEO	ΔNO	ΔHC	ΔEC
[0]	[0]	[0]	[-]	[-]	[-]	[-]	[+]	[+]
[0]	[0]	[0]	[+]	[0]	[+]	[+]	[0]	[0]
[0]	[0]	[-]	[0]	[-]	[-]	[0]	[+]	[+]
[-]	[0]	[0]	[0]	[-]	[+]	[0]	[+]	[+]
...

Table 2: Part of the qualitative combustion model

Table 2 shows part of the qualitative combustion model. It relates deviations of the input variables to deviations of the output variables. For example, the first row states that 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). Our model encodes various such situations, but is still rather incomplete, since up to now only the effects of certain variable deviations have been considered.

This level of representation of the diesel engine behavior was not only useful for establishing a link between the hydraulic and pneumatic subsystems. An interesting experience was that it also appears to be a practical level to communicate with experts in the domain.

Qualitative Model of the Control Unit

The Electronic Control Unit (ECU) comprises the control hardware (micro-controller), the control software as well as the evaluation and excitation logic. ECU processes are too complex to be modeled at transistor level. Furthermore, faults in the ECU itself are not taken into account. We thus characterize the ECU behavior using the logical relation between sensor readings and actuator settings as they appear in the ECU software specification. Typical rules which are part of the modeled ECU behavior are

$$\begin{aligned}
 [U_{\text{bat}}] = [\text{src}] \wedge [U_{\text{gnd}}] = [\text{gnd}] &\Rightarrow \text{ECU-state} = [\text{on}] \\
 \text{ECU-state} = [\text{on}] \wedge [U_{\text{ign}}] = [\text{src}] &\Rightarrow \text{ECU-wakeup} = [\text{on}] \\
 \text{ECU-wakeup} = [\text{on}] \wedge [U_{\text{bat}}] = [\text{src}] &\Rightarrow [U_{\text{EKPrelay}}] = [\text{src}]
 \end{aligned}$$

This states that if the ECU is supplied with voltage then the ECU state will be active. If the ECU is active and the ignition signal is on, the ECU is in wake-up mode. When the ECU is in wake-up mode and it is supplied with voltage then the low-pressure pump (EKP) is activated.

Data Acquisition and Interpretation

At present, electronic control units still have rather limited computing power (typically 256 to 512 Kbytes memory and about 10 to 20 MHz clock speed). This does not suffice to integrate a model-based diagnosis system within the ECU software. To circumvent these restrictions, so-called application control units are used in the VMBD project. Application control units are equipped with special dual-ported memory chips and are normally used for adjustment and optimization of ECU parameters for a specific vehicle type. Thus, in principle all variables and signals of the control unit are accessible in real time, without interfering the normal operation of the ECU.

The data of the vehicle is interfaced to the model-based diagnosis system, which is running on a portable PC inside the passenger compartment. The conventional signal range and plausibility checks of the ECU are used to determine the initial deviations of signals. As described above, CID-constraints are applied to these observations. We used instances of CID-constraints to relate deviations of pressure and flow variables to deviations of their derivatives. This turned out to be crucial to compensate for the limited observability in the hydraulic and pneumatic subsystems.

Experimental Evaluation

Two demonstrator vehicles are available in the VMDB project, one for the common rail injection system and one for the distributor-type injection system. For each of the two guiding applications, failures are induced in the cars, the model-based diagnosis system will be run and results are compared with the conventional diagnostic capabilities of the injection system control units. The various failures in the demonstrator cars can be injected by switchboards from inside the passenger compartment. For example, the pneumatic leakages are simulated by installing additional valves controlled by electrical switches.

The project is currently preparing the technical preconditions for interfacing the model-based diagnosis system with the demonstrator vehicles. At this stage, the examples below have been based on manually derived observations only. Currently, the diagnosis system exploits only models of correct behavior.

Case Study 1: Common Rail Injection System

The Common Rail is a newly developed injection system that can maintain a high injection pressure independently of the engine speed ([Bosch 96]).

In its hydraulic subsystem, an electrical fuel pump conveys fuel from the tank and feeds it to the high pressure pump. The high pressure pump delivers the fuel to a common fuel supply, the rail. The fuel pressure in the rail, which can reach up to 1300 bar, is measured by a pressure sensor. The ECU compares its signal to the set point which is determined according to the actual operating conditions. If the measured value and the set point are different, an overflow orifice in the pressure regulator on the high pressure side is opened or closed. To inject a certain amount of diesel fuel into the cylinders of the engine, the ECU triggers solenoid valves which open and close the injectors. The duration of injection and the fuel pressure in the rail determine the injected fuel quantity.

Scenarios Possible faults in the Common Rail system that are interesting to detect and discriminate are failures of the electrical fuel pump, the pressure regulator and pressure sensor, and fuel injection problems caused by faulty injectors. All these failures can occur intermittently, making them hard to diagnose, off-board as well as on-board. The ECU currently shuts down the system in any of the above situations, although this is not always necessary. A fault in the pressure sensor, for example, could be handled by switching to open loop control of the pressure in the rail. A malfunctioning pressure regulator, on the other hand, could lead to a critical fuel pressure in the rail and, hence, really requires to stop the engine immediately.

Examples In the Common Rail application, many faults cause a drop in the rail pressure. In one of the chosen scenarios, the task was to find sensors that could help to discriminate less critical faults in the low pressure part of the fuel supply, like a loose contact of the electrical fuel pump power line. We used the models of electrical and

hydraulic components and the ECU to forecast the effects of such failure causes. The system predicted, besides a drop in the rail pressure, a decrease of pressure in the supply for the high pressure pump. We tested the discriminative power that would be gained if an additional pressure sensor was placed at this location. To this end, we used the predicted sensor data as input to model-based diagnosis. The set of diagnostic candidates inferred consisted of electrical faults in the power supply of the ECU and the electrical fuel pump, together with some less safety critical hydraulic faults, like an empty fuel tank or a clogged filter. None of these faults are considered severe enough to require a complete system shut down. This result indicates that in some situations, the additional sensor can enable on-board diagnosis of the Common Rail to take less restrictive recovery actions. This information is useful for the developer of on-board recovery algorithms. One perspective is to provide a test-bed, which enables to check the failure causes which lead to specific on-board symptoms. This, of course, requires a model of the whole diesel injection system and not just the subsystems covered in our experiment.

Case Study 2: Distributor-type Injection System

The other type of injection system we are concerned with is the distributor-type injection (DTI, [Bosch 96]). The main difference to the Common Rail system lies in the high-pressure hydraulic subsystem. In the DTI, there is no common high-pressure fuel reservoir for the injections. Instead, a high-pressure pump generates a certain amount of pressurized fuel individually for each injection.

The distributor-type injection is an approved system which has been in widespread use for many years. Diagnosis of this system is focused on off-board diagnosis and driver-perceivable effects related to emissions and performance. Often, such effects involve incomplete fuel combustion and increased carbon emissions, and are therefore called "black smoke" problems. The main categories of failures which can lead to this class of symptoms are insufficient airflow to the engine, wrong timing of the fuel injection, an excessive quantity of fuel injected or external influences like poor fuel quality.

Scenarios A list of scenarios, which comprises more than 40 single component faults, has been set up for this system. An additional category consists of 4 external influences such as the quality of diesel fuel. All of the scenarios share the symptom of increased emissions, especially black smoke, and performance deterioration of the engine as their main symptoms.

The failures to be dealt with involve air leakages (intake and exhaust gas pipes or vacuum pipes), malfunctions of valves (e.g. stuck-at-open or stuck-at-closed), increased friction in bearings (resulting in a delay of actuators) and control signal disturbances due to electrical failures. For some of these failures, causal chains have been developed together with experts in the domain, sometimes based on reusing and refining existing failure mode and effects

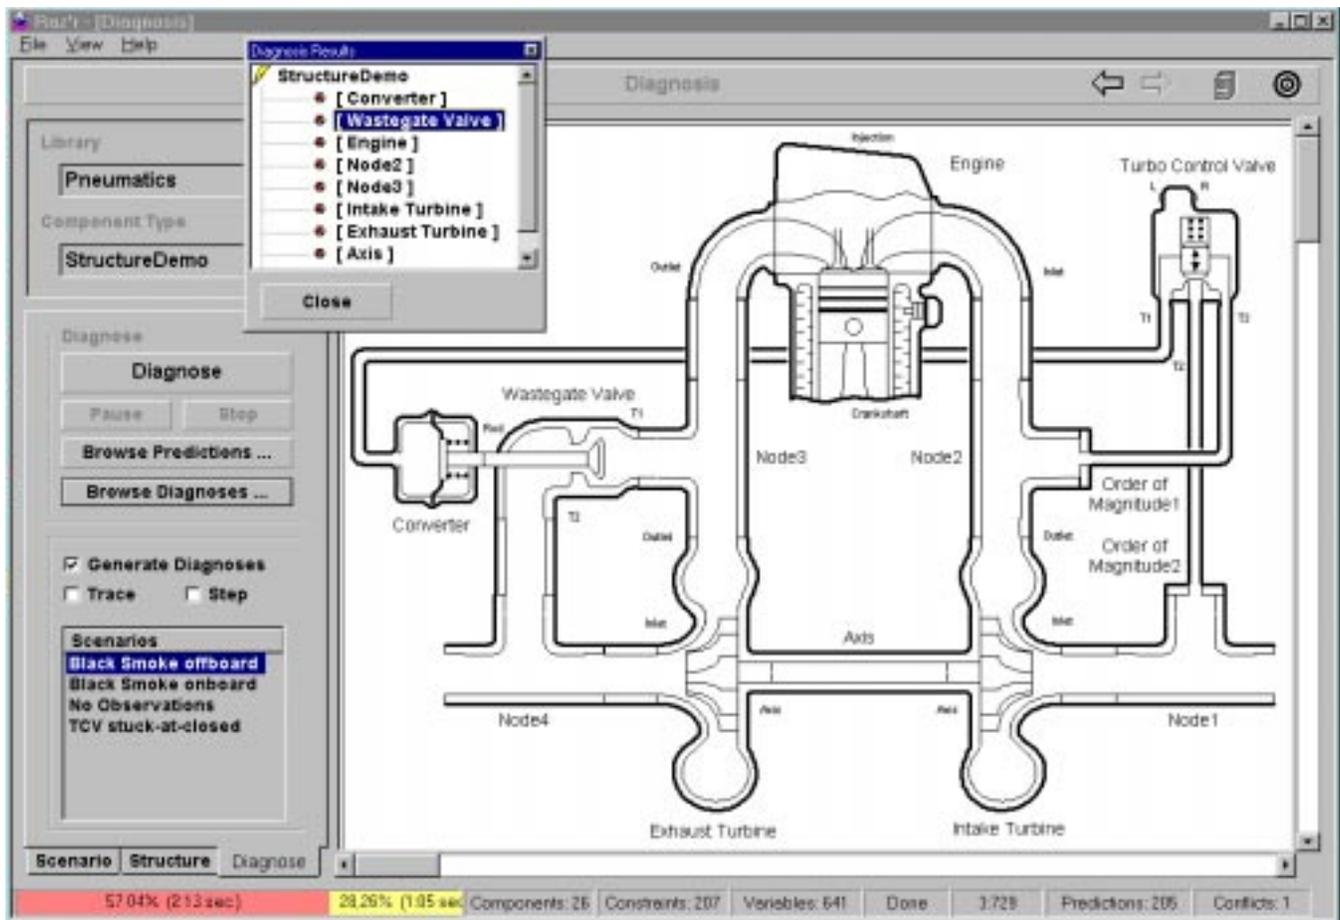


Figure 3: Diagnostic session in RAZ'R using a model of the DTI turbo control subsystem

analysis (FMEA) documents. A large part of categories involves faults in the air supply, which can be decomposed into the exhaust gas recirculation subsystem and the turbo control subsystem (Figure 3). The latter subsystem consists of a turbocharger turbine, which is driven by the engine's exhaust gases, for compressing (and thereby increasing the mass of) the intake air. The ECU controls the boost pressure admitted in a certain driving situation by opening or closing the turbo control valve (component in the upper right corner of Figure 3), which in turn sets the position of a so-called wastegate valve (component on the left). The position of this valve determines how much of the exhaust gas drives the exhaust turbine of the turbocharger.

Examples For the DTI system, case studies for both behavior prediction and diagnosis were performed. In a first scenario, the task was to predict the effect of the turbo control valve being stuck-at-closed. This fault has the following effects: If the turbo control valve is stuck-at-closed, it fails to reduce the control pressure in the pipe that connects it to the converter. Therefore, the control pressure in the converter remains too high. Due to the too high pressure, the wastegate valve will be opened too wide. Therefore, too much of the exhaust gas bypasses the exhaust turbine and, hence, the rotational speed of the

turbocharger is lower than normal. As a result, the air intake pressure is too low. This leads to a reduced mass of oxygen in the engine's combustion chamber. Provided that the injected fuel quantity remains normal, this leads to incomplete combustion of the diesel fuel which causes increased carbon emissions in the exhaust gas (due to unburned particles) and reduced torque of the engine.

Using the models described above, these effects can be derived automatically. The stuck-at-closed failure of the valve can be described as $[A] = 0$, $[\Delta A] = -$. Furthermore, we have to state that the engine is running, the atmospheric pressure applies at the air inlet, and other operating conditions. Once these observations have been defined, prediction proceeds in two stages: the system applies the models of pneumatic and mechanical components to perform behavior prediction and derives that the air pressure at the engine inlet is too low. The engine model is then used to derive increased carbon emissions ($[\Delta EC] = +$) and decreased torque ($[\Delta \text{Torque}] = -$) at the engine crankshaft as a result of this deviation. These results may further be used as an input to the model of the exhaust subsystem.

In a scenario for diagnosis, the task was to identify the set of component failures in the turbo control subsystem which

can lead to increased carbon emissions in the exhaust gas. In this case, we state that increased carbon emissions ($[\Delta EC] = +$) are observed at the exhaust side of the engine. Due to the unspecific nature of this measurement, and since no fault models are exploited presently, almost all components occur as diagnostic candidates. However, additional observations are available in the off-board situation that can help discriminating and lead us to more specific results. Under workshop conditions, pressure can be measured relatively easily at various points in the system. For instance, refining the scenario by an additional measurement of the converter pressure, stating that it is normal ($[\Delta p] = 0$), cuts down the set of candidates to nine possible single faults. Conversely, if the boost pressure is observed to be normal, diagnosis in this subsystem ends up with the engine as a single fault candidate. Even if the engine itself is not considered to be faulty, its model (by dropping the assumption that the amount of fuel injected in the engine is normal) can serve as a hook to further resume diagnosis in the DTI high-pressure hydraulic subsystem.

Current State and Future Work

So far, the project has been mainly concerned with the analysis of the vehicle subsystems and the application requirements and with feasibility studies based on some initial set of models. Meanwhile, the project enters the phase in which the diagnosis system is evaluated with real data obtained from the vehicles. Driven by the experiments, we will have to

- extend the model library to include more components and, in particular, fault models,
- revise and extend component models (e.g. splitting component behavior, especially the engine model, into operating regions, enhancing the ECU model, adding features, such as time delays and propagation of material properties like the oxygen rate),
- implement a data acquisition tool transforming the signals to a representation which can be fed to the qualitative models,
- explore possible limitations of the state-based diagnosis approach for the given scenarios and, perhaps, develop extensions and revisions.

At this time, we are confident that the project will achieve its goals and produce successful demonstrators. Proving the utility of model-based systems on the vehicles will be an important step in the industrial exploitation of our technology not only in the car industry. Besides these technical issues, we would like to convey two important messages based on our experience:

- to managers and technical staff in industry: the technology matures and becomes relevant to real applications,
- to the research community: addressing industrial applications and pursuing practical goals is a major driving force for research in the field rather than a distraction or road block.

Acknowledgments

We would like to thank the partners in VMDB, the members of the model-based systems and qualitative reasoning group Munich, and OCC'M Software GmbH for their valuable contributions. This work was supported in part by the Commission of the European Union (#BE 95/2128) and by the German Ministry of Education and Research (#01 IN 509 41).

References

- [Bosch 96] Robert Bosch GmbH (ed.) 1996. *Automotive Handbook, 4th Edition*. Warrendale, PA: Society of Automotive Engineers (SAE)
- [de Kleer and Williams 87] de Kleer, J. and Williams, B. C. 1987. Diagnosing multiple faults. *Artificial Intelligence* 32(1): 97-130
- [Dressler 96] Dressler, O. 1996. On-line Diagnosis and Monitoring of Dynamic Systems based on Qualitative Models and Dependency-based Diagnosis Engines. In Proceedings of the European Conference on Artificial Intelligence (ECAI-96), John Wiley & Sons
- [Loeser et al. 98] Loeser, T., Iwasaki, Y., Fikes, R. 1998. Safety Verification Proofs for Physical Systems. Technical Report, KSL-98-14, Knowledge Systems Laboratory, Stanford University
- [Malik and Struss 96] Malik, A. and Struss, P. 1996. Diagnosis of Dynamic Systems Does Not Necessarily Require Simulation. In Workshop Notes of the 10th International Workshop on Qualitative Reasoning (QR-96), AAAI Press
- [Sachenbacher and Struss 97] Sachenbacher, M. and Struss, P. 1997. 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
- [Struss 97] Struss, P. 1997. Fundamentals of Model-Based Diagnosis of Dynamic Systems. In Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI-97), Nagoya, Japan
- [Struss et al. 95] Struss, P., Malik, A., Sachenbacher, M. 1995. Qualitative Modeling is the Key. In Workshop Notes of the 6th International Workshop on Principles of Diagnosis (DX-95), Goslar, Germany
- [Struss et al. 97] Struss, P., Sachenbacher, M., Dummert, F. 1997. 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
- [Tatar 97] Tatar, M. 1997. *Presentation of the VMDB project at the 8th International Workshop on Principles of Diagnosis (DX-97)*, Mont-Saint-Michel, France