scope: *diagnosis*
contribution: *applications*

A model-based approach to fault localisation in power transmission networks

by Anton Beschta, Oskar Dressler, Hartmut Freitag, Michael Montag and Peter Struß

The basic idea of model-based diagnosis is to exploit knowledge about the structure and behaviour of the physical system to be diagnosed in order to identify diagnoses from discrepancies between observed and predicted behaviour. In this paper, we describe the application of model-based diagnosis to the problem of fault localisation in power transmission networks. We also show that the resulting system, DPNet, has an extended competence w.r.t. rule-based approaches (treatment of unknown and multiple faults) while requiring less development effort because of the use of component libraries in the model-based approach.

1 Introduction

Failures in technical systems are often limited to the systems themselves. A faulty copier, for example, will not normally cause further damage; we can simply use another machine. In other cases, system malfunctions may have severe consequences and side-effects on the system environment; examples include energy generation systems and the densely connected power distribution networks. In these cases, for economic reasons, it is necessary to detect and locate faults as quickly as possible in order to restore normal system function, thus minimising downtime.

Work on computer-aided support for these tasks is already underway in many areas. The results, however, are often unsatisfactory. Computer-aided diagnosis therefore is still a research topic. This is particularly true for the work carried out in the area of artificial intelligence, which tries to represent as much knowledge and experience of human experts as possible in so-called *knowledge-based systems*. (Please see References 1 and 2 for the artificial intelligence diagnosis approach in general and References 3–8 for the area of power distribution networks.)

Diagnosis means determining malfunctions. The idea of '*mal*function' presupposes expectations about *function*, i.e. correct behaviour. Therefore, discrepancies between expected and observed behaviour are the primary source of diagnostic information. The expectations necessary under this view must come from a model of what should happen.

Given this simple insight, it is surprising that most of the currently available diagnosis systems do not work on the basis of discrepancies between expected and observed behaviour. Instead, they follow a rule-based approach in the tradition of first-generation expert systems from the mid-1970s. These systems can be understood as a kind of generalised fault dictionary associating faults with symptoms. The approach seemed to be appropriate for *medical* diagnosis because models of the human body are not generally available.

Even before the current widespread application of first-generation expert systems, the scientific community realised that the approach would necessarily fail for the area of *technical* systems. Reasons for this include the following:

- 1 only faults that have been explicitly added to the dictionary can be diagnosed.
- 2 fault symptom associations can, at best, capture the

possible failures of a specific system.

• 3 interactions of faults cannot be adequately described using fault symptom associations only.

From an economic point of view in particular, the first two points make the fault dictionary approach highly undesirable. Shrinking product cycles and permanently changing device structures, such as network topologies, prevent us from even collecting the most probable faults as required by reason 1. The increasing number of variants of the same device/system leads to the expensive development of several expert systems for more or less the same device/system because of reason 2.

In contrast to the fault dictionary approach, models of the correct behaviour of system compoents play a central role in model-based diagnosis. A system is represented by describing the *behaviour* of individual system components and the *structure* of the system, i.e. the interaction paths between components. From this information, a model-based diagnosis system is generated automatically. The aforementioned problems disappear completely.

Continual research over the last decade has brought model-based diagnosis to a point where a second generation of expert systems for technical diagnosis is now being developed for applications. Considering the widespread (and sometimes successful) use of the comparatively simple fault dictionary approach, model-based systems building on libraries of explicit representations of component behaviour, and on an explicit representation of the actual system structure, should have a major impact on all areas with technical diagnosis problems.

As an application of the domain-independent diagnosis system GDE^{+} (Extended General Diagnostic Engine) [9] developed by the Advanced Reasoning Methods Group at Siemens Research, a prototypical application system for fault analysis in power distribution networks, DPNet, was built. Although no domain expert was involved and domain knowledge therefore had to come from textbooks describing the correct behaviour of the system, an initial evaluation on real data was positive; all available message bursts are correctly analysed and faults properly located.

2 The application domain

The energy distribution system is the connecting link between energy suppliers and consumers. According to its size with respect to supply and demand of energy, the network is usually divided into high-, medium- and low-voltage subnets. Subnets transmit energy on more than one voltage level, e.g. the high-voltage network in Germany consists of 380 kV, 220 kV and 110 kV voltage levels. These levels are coupled by means of transformers and generation groups of power plants can be connected to each of them (Fig. 1).

Transformers and generation groups are linked to bus-bars within stations. If the energy distribution network is viewed as a graph, then the bus-bars are the nodes of this graph. Lines connect bus-bars of different stations within, and consumers to, the network. All connections of devices to bus-bars are realised by systems of switches (i.e. insulators and circuit-breakers), which are also called bays. By means of these bays, the topology of the network can be adapted to the current load situation to achieve an efficient distribution of energy resources.

Faulty switch actuations, malfunctioning devices and external influences unfortunately lead to disturbances. Their effects, usually energy breakdown in regions of the transmission network, as well as thermal and mechanical damage of devices are minimised by protection mechanisms [11, 12]. In case of faults, the protection mechanisms detach parts of the network by the automatic actuation of the circuit-breakers of the corresponding bays.

In order to restrict energy breakdowns as much as possible, ideally by detaching only the faulty device, maximal selectivity is required from the protection mechanisms. Exceptions are devices of considerable economic value such as transformers (w.r.t. material) and bus-bars (w.r.t. network connectivity), where selectivity is subordinated to the protection of electrical appliances.

Distance protection is an example of a selective protection mechanism. It is most common for transmission lines and additionally serves as a back-up in case of failure of other protections. The mechanism is equipped with measuring instruments from which the direction of energy flow and resistance can be determined. In the case of a short-circuit in the direction of the connected line, which is reflected by a certain decrease in the resistance, the circuit-breaker of the local bay is actuated after a delay, assuming that the resistance has not returned to the nominal value in the meantime. The delay is determined by a locally adjusted staggered characteristic curve (Fig. 2). The decrease in resistance is a representative estimate for the distance of the short-circuit; the associated delay increases with increasing estimated distance. This behaviour leads directly to the required selectivity as the breakers closer to the fault are guaranteed to trip first. As a positive side-effect, there is always a back-up if one of the distance protections or breakers fails (but with obviously reduced selectivity).

However, in high-voltage networks, selective protection mechanisms are not sufficient, because extremely fast detachment is required to prevent nearby appliances from destruction caused by the large transmitted energy. In such cases, the differential protection is employed for high-speed intervention. This approach exploits Kirchhoff's first law by monitoring the sum of currents of all bays of the bus-bar. If the sum is not approximately zero, all circuit-breakers of all bays are immediately tripped. The intervention time of the differential protection can be an order of magnitude faster than the fastest distance protection level.

The protection mechanisms and the breakers of each bay send numerous messages to the associated control centre (Fig. 3). In particular, there are

□ warnings (the current in the bay is beyond the admissible threshold), eventually including information about the direction of energy flow, affected phases etc. (these messages are marked ANR in Fig. 3).

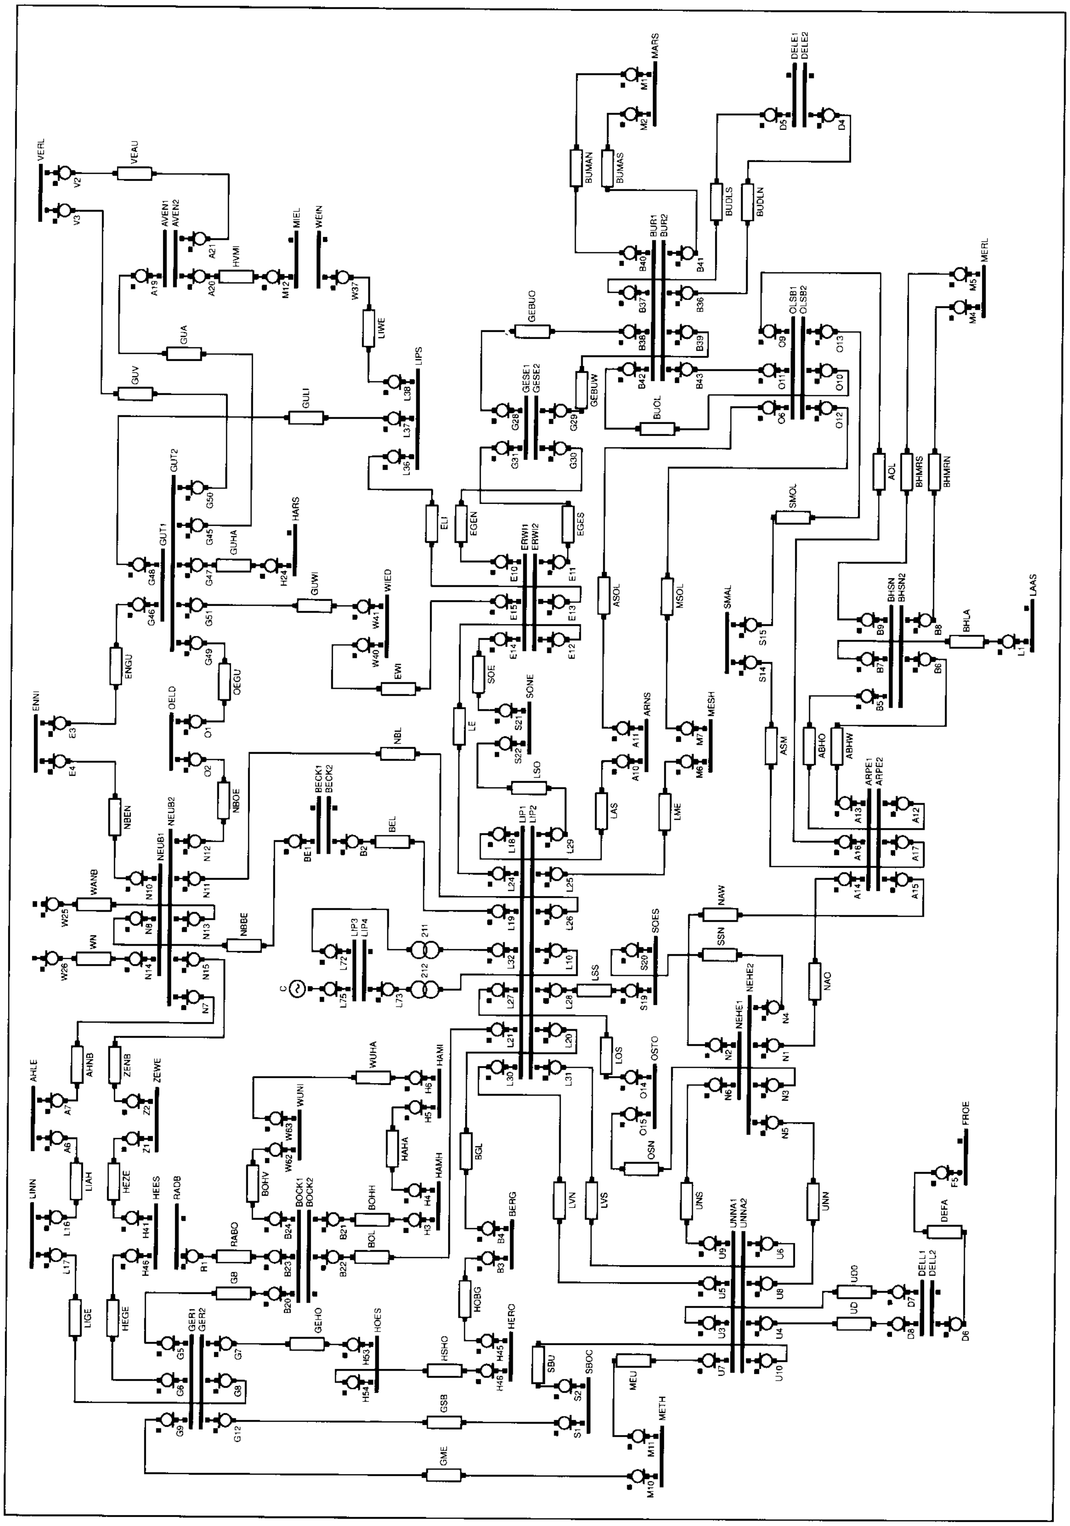

Fig. 1 Part of a power transmission network [10]

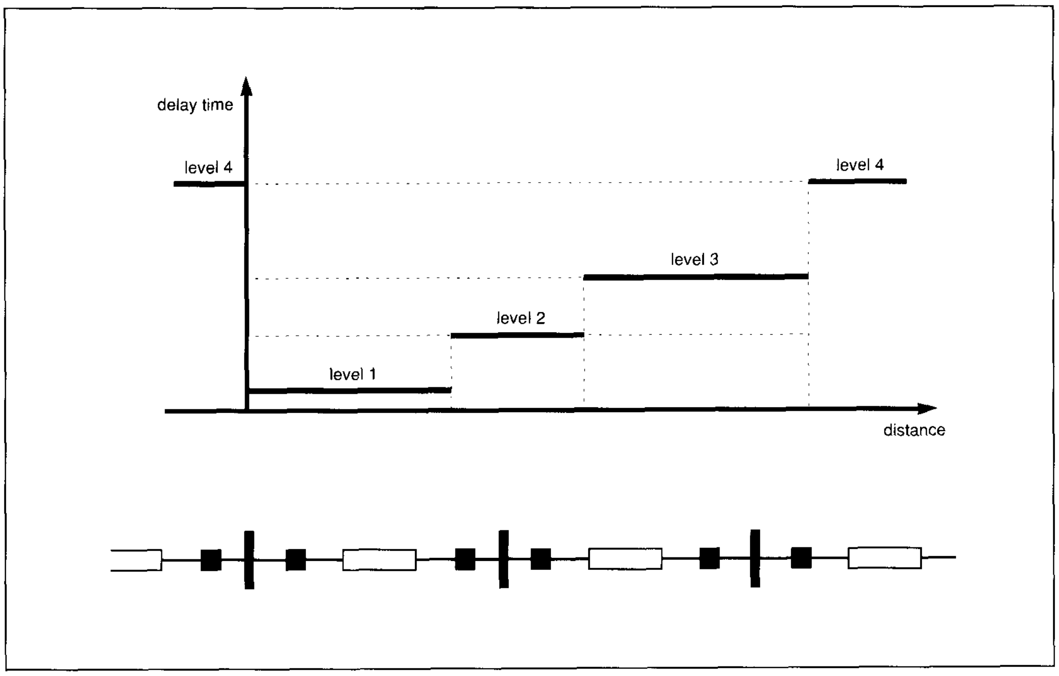

Fig. 2 Tuning of a distance protection system [11]

Level 4 triggers the breaker independently of the direction of the short-circuit current.

□ messages about interventions of the protection mechanism (marked AUS in Fig. 3).
□ requests and acknowledgments of actuations of circuit-breakers, including information about location and time.

During fault situations, the control centre personnel and a computer-aided diagnostic system are confronted with the following problems:

- depending on the kind of fault, hundreds of messages arrive within seconds.
- the messages do not arrive in chronological order.
- messages might be delayed significantly or even be missing.
- multiple temporally overlapping faults are difficult to separate.
- the reaction time to a fault should be as short as possible (real-time constraints).

The task of the engineer, and hence the replacement diagnostic system, is to locate and reconstruct as precisely as possible the cause of the fault from the incoming messages, eventually resulting in a sequence of breaker commands. These include the attachment of unnecessarily detached devices, switching to available back-up appliances in order to re-establish energy supply and operations for achieving an optimal load distribution.

Diagnostic systems which allege to be effective in this field must satisfy the following additional requirements:

□ fast adaptation to changes of the network topology through actuations of breakers.
□ flexibility in handling many different existing device generations and the introduction of new devices.

3 Knowledge-based diagnosis

3.1 The first generation: experience-based diagnosis

Diagnosis is one of the main application areas of so-called first-generation knowledge-based systems. These systems from the mid-1970s were quite successful in *limited* domains, e.g. certain areas in medical diagnosis. They operate on a set of fault-symptom associations typically encoded in the form of rules ('IF ⟨condition⟩ THEN ⟨consequence⟩'). The limitations of these purely experience-based systems became apparent when they were applied to technical domains. Here physical laws and general principles governing the selected domain play a crucial role in performance at the expert level. These laws and principles are available and should be used by the system. However, this type of knowledge can hardly be encoded in fault-symptom associations and, if so, only in an implicit way.

The main restrictions imposed on experienced-based diagnosis systems are the following:

- they are limited to known symptoms and their causes.
- they are specialised for one or a small set of device type.
- they fail to degrade gracefully because they do not know the boundaries of their competence.

07/15.08.28	FAV	110		LS-		FALL	
07/15.08.28	FAV	110	G	LS	AUS		
07/15.08.28	FAV	110	EQQU-Q	LS	AUS		
07/15.08.28	EQQU	110		LS-		FALL	
07/15.08.28	EQQU	110	FAV-Q	LS	AUS		
07/15.08.36	KQPR	110		S-AMR			
07/15.08.36	GISN	110		LS-		FALL	
07/15.08.36	GISN	110		S-ANR			
07/15.08.36	GISN	110	FAV	LS	AUS		
07/15.08.36	QOEG	110		S-ANR			
07/15.08.36	GOSN	110		LS-		FALL	
07/15.08.36	GOSN	110		S-ANR			
07/15.08.36	OQQA	110		S-ANR			
07/15.08.36	GOSN	110	NEV	LS	AUS		
07/15.08.36	IOFT	110		S/ANR			
07/15.08.36	FAV	110	PDU-J	LS	AUS		
07/15.08.26.23	FAV	110	PDU-J	DIST	ANR	G	.05
07/15.08.26.29	GED	220	532	DIST	ANR	T Q	.24
07/15.08.36	ODJA	110	FAV	LS	AUS		
07/15.08.26.23	EQQU	110	KOPR	DIST	ANR	RST Q	.77
07/15.08.26.64	UDVO	110	QORO-I	DIST	ANR	RST	.13
07/15.08.26.66	MOD	110	RIOC	DIST	ANR	S Q	.06
07/15.08.26.64	GED	110	533	DIST	ANR	RST	.11
07/15.08.26.80	EQQU	110	QORO-C	DIST	ANR	RSTMQ	.19
07/15.08.26.28	FAV	220	533	DIST	ANR	RSTM	.53
07/15.08.26.64	UDVO	110	133	DIST	ANR	ST	.06
07/15.08.26.23	MOD	110	GISN	DIST	ANR	RST	.48
07/15.08.26.25	FAV	110	TEVI-C	DIST	ANR	R Q	.43
07/15.08.26.65	GED	110	532	DIST	ANR	RST	.19
07/15.08.26.80	EQQU	110	QORO-Q	DIST	ANR	RSTMQ	.19
07/15.08.26.66	MEP	110	IOFT	DIST	ANR	RST	.11
07/15.08.26.64	UDVO	110	QORO-J	DIST	ANR	RST	.14
07/15.08.26.23	MOD	110	133	DIST	ANR	RST	.26
07/15.08.26.65	GED	110	MOCO-J	DIST	ANR	G	.21
07/15.08.26.22	EQQU	110	FAV-C	DIST	ANR	RSTMQ	.77
07/15.08.26.66	MEP	110	OQQA	DIST	ANR	RST	.11
07/15.08.36	FAV	110	G	GEG	AUS		
07/15.08.26.23	MOD	110	133	DIST	ANR	RST	.26
07/15.08.26.65	GED	110	MOCO-J	DIST	ANR	G	.21
07/15.08.26.22	EQQU	110	FAV-C	DIST	ANR	RSTMQ	.77
07/15.08.26.64	MEP	110	OQQA	DIST	ANR	RST	.13
07/15.08.26.71	FAV	110	G	GEG	AUS		
07/15.08.26.70	MOD	110	KOPR	DIST	ANR	RST	.18
07/15.08.26.64	GED	110	MOCO-I	DIST	ANR	G	.22
07/15.08.26.20	EQQU	110	FAV-Q	DIST	ANR	RSTM	.80
07/15.08.26.28	MOD	110	CGIS	DIST	ANR	RST	.74
07/15.08.26.22	FAV	110	GISN	DIST	ANR	RSTMQ	.53
07/15.08.26.98	EQQU	110	FAV-Q	DIST	AUS		.77
07/15.08.26.23	FAV	110	533	DIST	ANR	RST	.58
07/15.08.53	FAV	110	533	LS	AUS		
07/15.08.26.22	FAV	110	ODJA	DIST	ANR	RST Q	.59
07/15.08.26.63	FAV	110	QOEG	DIST	ANR	RST	.21
07/15.08.26.23	FAV	110	GOSN	DIST	ANR	RSTMQ	.61
07/15.08.26.68	FAV	110	533	DIST	AUS		.45
07/15.08.26.65	FAV	110	535	DIST	ANR	RST	.19
07/15.08.26.22	FAV	110	ICPI	DIST	ANR	RST Q	.64
07/15.08.26.25	FAV	110	EQQU-C	DIST	ANR	RSTM	.67
07/15.09.05	FAV	110	TEVI-C	DIST	ANR	RST Q	
07/15.08.26.85	FAV	110	TEVI-C	DIST	ANR	RST Q	.16
07/15.08.26.21	FAV	110	EQQU-Q	DIST	ANR	RSTMQ	.91
07/15.08.26.93	FAV	110	EQQU-Q	DIST	AUS		.72
07/15.12.36	FAV	220		LS-		FALL	
07/15.08.12.36	FAV	220	S	LS	AUS		
07/15.12.34.09	FAV	220	S	GEG	AUS		

Fig. 3 A sample message burst [10]

STATUS	LEVEL	FDIST
OPEN	1	1
OPEN	2	{1, 2, 3}
OPEN	3	{3, 4, 5}
OPEN	4	{5, 6, 7, . . .}
OPEN	4	{−1, −2, −3, . . .}
CLOSED	0	Z\{0}
CLOSED	1	{1, 2, 3, . . .}
CLOSED	2	{3, 4, 5, . . .}
CLOSED	3	{5, 6, 7, . . .}
CLOSED	{1, 2, 3}	{−1, −2, −3, . . .}

Fig. 4 Model of a breaker

- they are not constructive as there is no methodology
  - for adapting a changed or similar device.
  - for generating diagnosis systems from given design data.

Therefore, developing and maintaining experience-based diagnosis systems requires a substantial amount of resources. From a financial point of view, this approach cannot be a basis for widespread industrial applications [13].

3.2 The second generation: model-based diagnosis

From the beginning of the 1980s, this experience of limitations of first-generation systems caused the development of so-called *model-based* diagnosis systems. Based on previous work on model-based reasoning and the 'General Diagnostic Engine' [14], the diagnosis framework GDE$^+$ has been developed [9, 13] and evaluated by a number of realistic studies. In the following, we use this system (which is also the basis for the DPNet prototype) to illustrate the basic ideas of model-based diagnosis.

A major prerequisite of this approach is a description of the device to be diagnosed in terms of its components, their normal behaviour and connections. Based on this description, the diagnosis system predicts the behaviour of the whole system, assuming that each component behaves correctly. Additionally, dependency records are maintained, which relate behavioural aspects to the correctness of those components responsible for producing this behaviour. When *discrepancies* are detected between observed and predicted behaviour, these dependency records can be analysed to retrieve the underlying sets of correctness assumptions, so-called *conflicts*. As a *diagnosis* has to explain each conflict, it has to retract at least one correctness assumption of each conflict. This normally yields several diagnoses, which may be further discriminated by making additional measurements.

The principal advantages of this approach are thus

- generality: a domain-independent diagnosis framework.
- a domain-specific component library.
- extendability to new device structures and new components (due to the explicit representation of the device structure which may be derived from CAD data).
- treatment of unknown symptoms and faults, in particular multiple faults.

These features allow us to overcome the limitations of first-generation systems, and provide a basis for developing and maintaining knowledge-based diagnosis systems in an economically feasible way.

A major requirement of this approach is a domain-specific component library containing a local context-free model for each component of an application domain. A model describes the normal behaviour of a component by its characteristic variables and their relations. In the following, we use a small example from the domain of power transmission networks to illustrate this approach.

3.3 Component models

The component library of DPNet specifies the normal behaviour of components during a disturbance of the network by means of the following parameters:

- the *distance protection level* (LEVEL) at which a breaker has been tripped (corresponding to the different levels in Fig. 2).
- the *state* (STATUS) of the breaker.
- the *direction* of the short-circuit (+ means the location of the fault in the direction of the line and − in the direction of the bus-bar).
- the *distance* to the location of the fault (in terms of the number of lines and bus-bars between the current position and the fault), which may be computed using the distance protection level.

The last two parameters are combined to produce a directed distance FDIST (whose sign is the direction and the amount is the distance to the location of the fault).

In the following, we describe models as (mathematical) relations which are simply written as tables listing the tuples of the relation.

The model of a breaker (Fig. 4) describes the following.

- If the breaker has been tripped (STATUS = OPEN), we can compute the direction and distance of the short-circuit using the current distance protection level (LEVEL). If the breaker has been tripped at distance protection level 1, the short-circuit current is directed to the line and the distance is 1 (i.e. the line is possibly faulty). For distance protection levels 2 and 3, there is no unique fault distance as one step in the distance protection's characteristic curve does not exactly correspond to the length of a line. If the breaker has been tripped on level 4, the short-circuit current may also be directed towards the bus-bar.
- If the breaker's associated switch is closed, we can derive similar constraints for the values of LEVEL and FDIST.
- Clearly, using this relation we can also derive the expected state of the breaker (STATUS), if we know the distance protection level and the fault distance, or we

may infer the actual distance protection level from the state of the breaker and the fault distance.

The behaviour of a line is modelled (Fig. 5) as follows:

□ a line operating in normal mode has no short-circuit, i.e. the location of the fault is not on this line. Thus, the fault distance must be either ≥ 2 (i.e. at least two components in the direction of the line) or ≤ -1 (at least one component in the direction of the bus-bar).
□ moving closer to the fault reduces the fault distance by 1, and moving in the opposite direction increases the fault distance by 1. In both cases, the sign of the fault distance has to be inverted to indicate the direction of the short-circuit current.

(Clearly, in the case of a disturbance, it is wrong to assume that all components behave correctly. However, the basic idea of this approach is to exploit discrepancies between observed and predicted behaviour to infer diagnoses.)

The component library contains such models for each component type. Although these qualitative models (no precise quantitative descriptions are used) require a number of simplifying assumptions, they are sufficient to diagnose a large number of possible faults. Furthermore, they provide a basis for the use of more detailed and exact models (see Section 5).

3.4 System modelling

A particular network is modelled by connecting instances of such component models. As a simple example, we use a network consisting of a line L1 connected to two breakers A1 and A2 (Fig. 6).

Using this system description and initial observations, it is possible to determine the state of this system by combining local component models. During this prediction step, the system records which components and thus which correctness assumptions justify the inference steps. This is done by means of an assumption-based truth maintenance system (ATMS) [15], which also provides a basis for dealing with incomplete and prototypical information (see Section 5).

For example, if we observe that breaker A1 has been tripped on distance protection level 1 (STATUS(A1) = OPEN, LEVEL(A1) = 1, we may use the breaker model to derive that

$$\text{FDIST(A1)} = 1 \qquad \{\text{A1}\}$$

$\text{FDIST}_{\text{LEFT}}$	$\text{FDIST}_{\text{RIGHT}}$
{2, 3, 4, . . .}	$1\text{-FDIST}_{\text{LEFT}}$
{−1, −2, −3, . . .}	$1\text{-FDIST}_{\text{LEFT}}$

Fig. 5 Model of a line with LEFT and RIGHT terminals

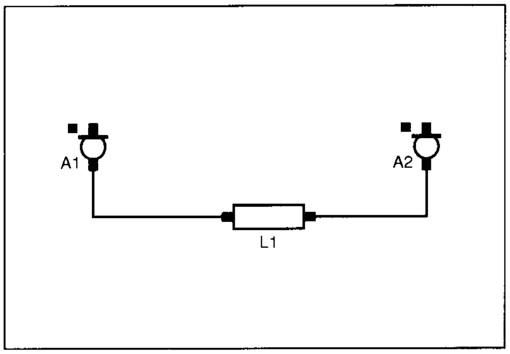

Fig. 6 Part of a network

holds. In addition, we record (in curly brackets) that this inference depends on the correctness of A1 (and the correctness of the observations which we believe unconditionally). Inferences are only performed by a combination of local models, so-called *local propagation*.

3.5 Generating diagnosis candidates

As A1 is connected to the left terminal of line L1

$$\text{FDIST}_{\text{LEFT}}\text{(L1)} = \text{FDIST(A1)} = 1 \qquad \{\text{A1}\}$$

holds. This is inconsistent with the model of a correct line, which requires that the fault distance is either ≥ 2 or ≤ -1. Using the initial observations (i.e. parts of the message burst) and the models for correct behaviour, it is possible to derive contradictory values for some variables; the fault distance is inferred to be 1 and ≥ 2 or ≤ -1 justified by the correctness assumptions for A1 and L1. This conflict implies that at least one of these correctness assumptions is wrong, i.e. diagnosis candidates are [A1] and [L1]. Clearly, it is possible that there are additional faulty components, but in general, we are primarily interested in diagnoses suspecting only a small number of components, and thus we consider only minimal (w.r.t. set inclusion) candidates.

Observing that breaker A2 has been tripped on level 1, we can also derive in a similar way the conflict {A2, L1}. Both conflicts can be resolved by retracting the correctness assumption for L1, or by retracting the correctness assumptions for A1 and A2, i.e. possible diagnoses are [L1] and [A1, A2].

Two faulty breakers are less probable than a short-circuited line, and thus GDE^+ assigns a lower rank to this diagnosis. Ranking of diagnoses in GDE^+ is according to cardinality and *a priori* fault probabilities of components. Using this approach, we can also derive implausible and physically impossible diagnoses. This is because GDE^+ does not require any knowledge about faulty behaviour (which is a principal advantage compared to experience-based systems). Fortunately, by using such fault models, these candidates may be eliminated. The underlying inference pattern is to exonerate a component if none of its known fault models explains the observed behaviour [9, 16].

$FDIST_{S1}$	$FDIST_{S2}$	$FDIST_{S3}$
{−2, −3, −4, . . .}	{−2, −3, −4, . . .}	−1-min($FDIST_{S1}$, $FDIST_{S2}$)
{1, 2, 3, . . .}	{1, 2, 3, . . .}	−1-min($FDIST_{S1}$, $FDIST_{S2}$)

Fig. 7 Model of a bus-bar with three terminals (S1, S2, S3)

Discrimination among candidates can be carried out by additional observations (messages). In order to minimise the number of additional measurements for obtaining a unique diagnosis, GDE^+ uses an information-theoretic method to determine the next mesurement point that provides maximal information gain [17].

4 The DPNet system

In this Section we describe the DPNet prototype developed using GDE^+. DPNet has been tested using real message bursts and parts of real networks (Fig. 1). As GDE^+ already provides the general framework for diagnosis, development work for DPNet focused mainly on the development of a component library for power transmission networks, and thus required only a fraction of the efforts necessary to develop a conventional system.

4.1 The DPNet component library

As previously mentioned, component models describe the normal behaviour of components during a disturbance of the network. The models of the protection mechanisms and their actions allow the localisation of a short-circuit as well as the diagnosis of faulty protection mechanisms (e.g. the unsuitable intervention of a breaker). The current prototype considers only distance protections (because they reveal the most interesting behaviour pattern), but may be easily extended to other protection mechanisms, e.g. differential protections. In the following, we describe these component models. Models for breakers and lines are illustrated in Section 3.

4.1.1 Breaker: a breaker with distance protection level 1 and a short-circuit current in the direction of the line trips the associated breaker (if the breaker works correctly). On the other hand, in normal operating mode, it is possible to derive the fault distance from the tripping command and the distance protection level.

4.1.2 Line (with LEFT and RIGHT terminals): in normal operating mode, there is no short-circuit on the line, i.e. the short-circuit currents observed at both sides of the line must have opposite signs. The fault distance describes the distance between the current component and the short-circuit by the number of components. Thus, by crossing the line, we have to add 1 to the fault distance.

4.1.3 Bus-bar (with three terminals S1, S2, S3): a correctly working bus-bar has no short-circuit, i.e. the fault distance of at least one terminal must be positive (the short-circuit current must be flowing to a line). Furthermore, the fault distance of at least one terminal must be negative; otherwise, the bus-bar would produce current. Moving across the bus-bar subtracts 1 from the negated minimum of the fault distances.

4.1.4 Transformer (with LEFT and RIGHT terminals): with respect to a disturbance of the normal operation of the network, a transformer qualitatively behaves like a line.

4.1.5 Generator: in the case of a disturbance, the short-circuit current is usually provided by a generator, i.e. the fault distance must be negative.

$FDIST_{LEFT}$
{−1, −2, −3, . . .}

This component library can be easily extended by new component types, e.g. connections between parallel bus-bars.

4.2 Fault localisation

In the following, we demonstrate by a sample session of DPNet the localisation of a fault from a message burst (Fig. 3) in a part of a real network (shown in Fig. 8). The message burst is caused by a multiple fault:

- the bus-bar FAV3 is faulty.
- the generator S does not work properly (note that this generator is on a different voltage level than the bus-bar).
- there is an unsuitable intervention of breaker F3.

These faults cause an isolation of bus-bar FAV3 by tripping all breakers on the 'opposite' side (on distance protection level 2). The faulty generator is disconnected from the network by the directly associated breaker.

Usually, only a fraction of the available messages is required to determine the fault location(s). Therefore, the system initially considers only messages concerning the tripping of breakers and neglects the rest, which provides a useful focus for the analysis. With correctly working protection mechanisms, the faults described above cause the opening of breakers E2, G36, IC5, F55, GO5, O1, UC3, F53 and F3. The substations PCPI and UDQC do not have any direct communication lines to

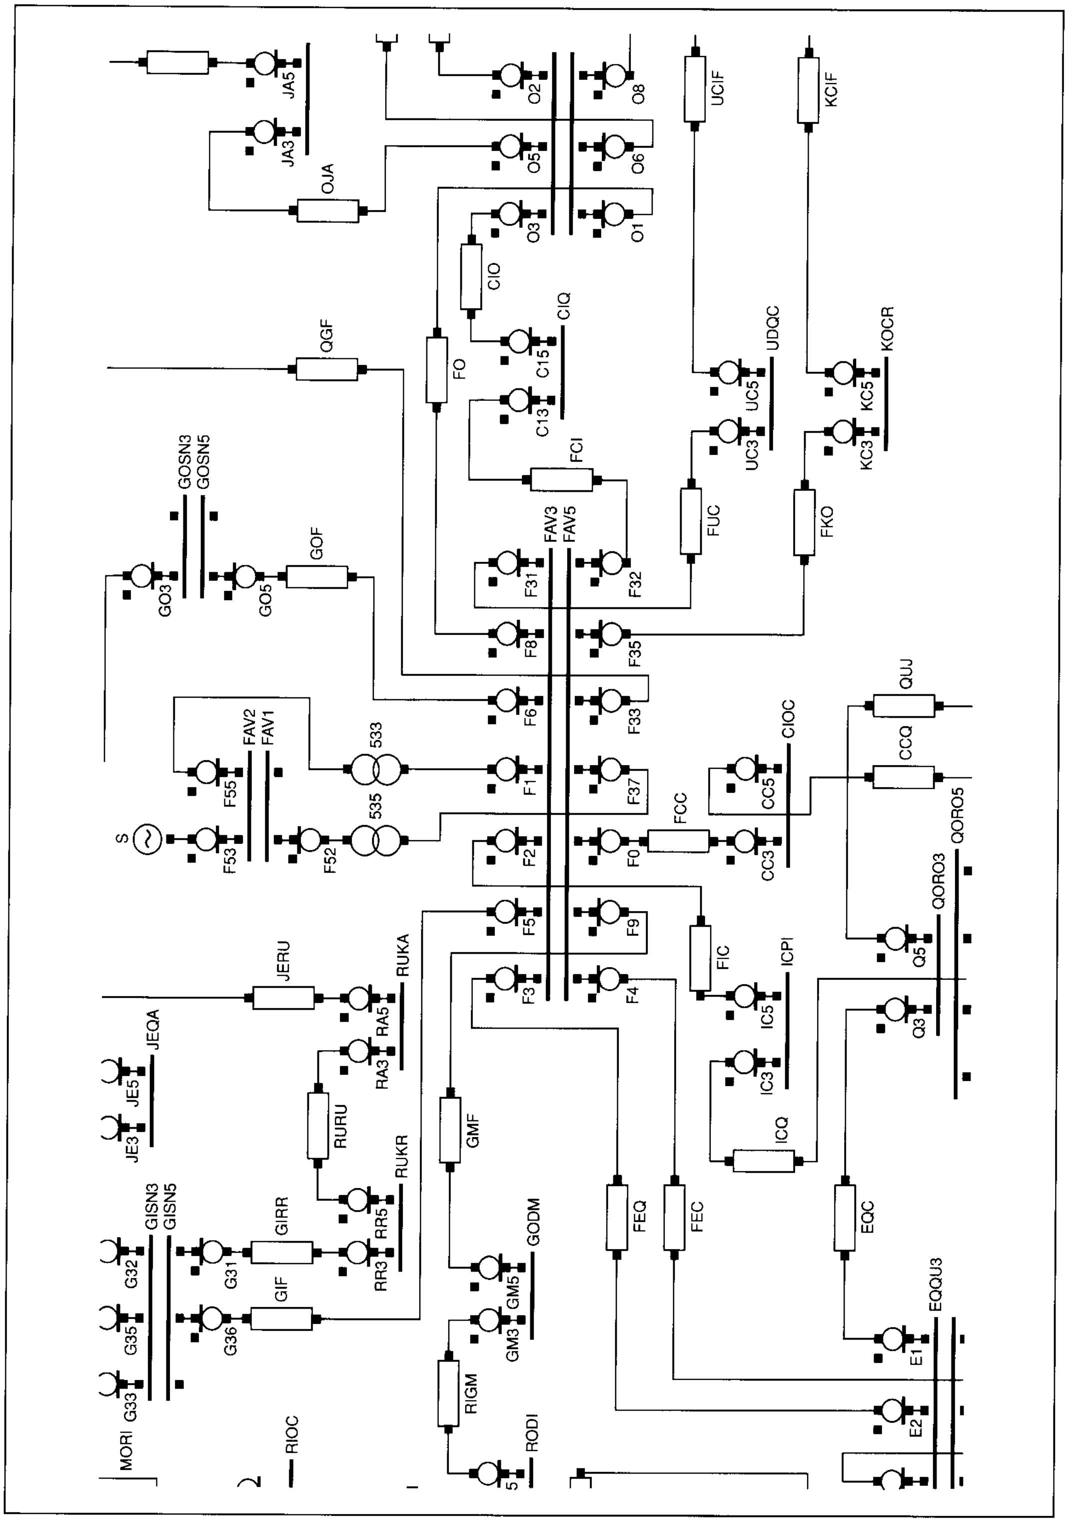

Fig. 8 The faulty network (part of the network in Fig. 1)

the control centre, and thus the message burst contains only the tripping commands for E2, G36, F55, GO5, O1, F35 and F3. Starting from these messages, the system generates as initial diagnoses:

[FAV3, S, F3]
[FAV3, S, FEQ]
[FAV3, F53, F3]
[FAV3, F53, FEQ]

All diagnoses suspect bus-bar FAV3 of being faulty, and they also include generator S or breaker F53 and breaker F3 or line FEQ. There are also several diagnoses which suspect at least seven components and represent different improbable combinations of line and breaker faults. This is not very surprising as the exclusive consideration of tripping commands does not determine a unique diagnosis.

Discrimination between these initial diagnoses can be carried out by exploiting additional information about the state of the network. Therefore, the measurement proposal routine of GDE$^+$ is used to select additional messages from the message burst. The only measurable quantity that can be obtained from messages is the direction of the short-circuit current. In our example, the system selects the direction of the short-circuit between line FEQ and breaker F3 as the best measurement. The corresponding message provides the − direction (i.e. towards the bus-bar), which eliminates all diagnoses including FEQ. The remaining diagnoses are

[FAV3, S, F3]
[FAV3, F53, F3]

Additionally, observing the direction between S and F53, which yields +, leaves

[FAV3, S, F3]

as the single best diagnosis, as well as several other implausible and improbable diagnoses, i.e. using only tripping commands and a few additional messages is sufficient to obtain a fairly good estimation of the fault location.

GDE$^+$ does not require any assumptions about the faulty behaviour of components, which allows the treatment of unknown symptoms but may also cause physically implausible diagnoses. This is the reason for some of the additional diagnoses in our example. Fortunately, GDE$^+$ exploits knowledge about faulty behaviour in order to eliminate such diagnoses. Applying these additional inferences, we have the diagnoses

[FAV3, S, F3]
[E2, G36, F55, GO5, O1, S, F3]

The second diagnosis represents the situation that all open breakers are unsuitable interventions, a very improbable but physically possible situation.

Evaluating DPNet against 20 real message bursts revealed that the system produces a fast localisation of faults (response time is about one minute). As only the relevant part of the network (focusing by tripped breakers) and only a fraction of the message burst are considered, this response time is more or less independent of the size of the network. It is important to note that the basis for this implementation, GDE$^+$, is only a research prototype. As the main part of the work was the development of a domain-specific component library, DPNet required only about six working months of development effort (including domain analysis and improvements of the GDE$^+$ implementation). GDE$^+$ (and DPNet) is implemented in Common Lisp and available on SUN workstations and SYMBOLICS Lisp Machines.

5 Discussion

We discuss the potential, the current limitations and possible improvements of the current prototype, based on this case study. This discussion considers the following aspects:

- processing of available information:
  - selective message processing.
  - dealing with incomplete information.
  - revision of intermediate results.
  - incremental message processing.
- diagnostic results:
  - dealing with unknown, unusual and multiple faults.
  - ranking of alternative diagnoses.
  - refinement of results.
- adaptability, extensibility and maintainability:
  - changes of topology.
  - new technology and new types of components.

5.1 Selective information processing

DPNet automatically filters the incoming message burst by selecting only messages which potentially refine the current diagnostic results (by discriminating between different, alternative diagnoses). This filtering depends on the *current situation*, and thus does not require an *a priori* ranking of single messages. Selecting messages according to the expected information gain only affects the performance of the system and not the final result. The heuristic used in DPNet to favour messages about interventions by the protection mechanism seems to be appropriate for this purpose.

5.2 Reasoning with incomplete information

The fact that for some substations no messages are available has been already mentioned. This, and the possibility of delayed messages, requires the computation of diagnoses based on the available, possibly incomplete information. The example in the previous Section, however, has already demonstrated that the system can cope with this additional difficulty. In contrast to rule-based systems, DPNet does not need to match complete patterns in order to apply a rule, but is able to infer the unknown state of a breaker based on the currently available information and use this new information for further inferences.

This causes the problem that some of these conclusions may be invalidated by additional information (e.g. delayed messages) and therefore have to be retracted.

5.3 *Retractable inferences*

The task of retracting inferences based on assumptions poses no additional difficulty for GDE^+. In Section 3, we illustrated how the system keeps track of the assumption sets underlying a fact by means of an assumption-based truth maintenance system. In the case of discrepancies between a derived fact and an observation (which is considered as true), the underlying sets of assumptions are marked as inconsistent, and thereby all other facts which depend on these assumption sets are also retracted.

The ability to deal with incomplete information and to revise intermediate conclusions opens up the following perspective (which is not yet realised in DPNet).

5.4 *Incremental message processing*

As the system does not rely on the completeness of a message burst (when is a message burst complete?) and is able to revise preliminary conclusions, it is possible to process messages as they are coming in, and thus reduce the space for possible diagnoses. Further messages then may provide additional discrimination w.r.t. these diagnoses. This property seems to be very important for satisfying real-time requirements while maintaining completeness w.r.t. diagnoses.

5.5 *Unknown, unusual or multiple faults*

As already mentioned, DPNet uses only models of correct behaviour for computing diagnoses. It does not rely on knowledge about previously observed fault situations, and therefore is also able to deal with unknown and unusual faults and fault combinations, which makes this approach superior to purely experience-based approaches. The example in Section 4 illustrates that multiple faults can be diagnosed, which enables the system to identify short-circuits *and* faults in the protection system.

On the other hand, this may result in a huge space for possible diagnoses, allowing either technically implausible or theoretically possible but improbable faults. It has been shown that the first class of diagnosis can be excluded by additionally exploiting knowledge about faulty behaviour of components (fault models).

5.6 *Ranking of diagnoses*

As this approach also provides theoretically possible but practically improbable diagnoses, an additional ranking according to the corresponding likelihood (not necessarily in a mathematical sense) is required to focus the system on a small set of interesting diagnoses. (Note that, for example, a possible diagnosis is to always suspect all breakers which have been opened.) GDE^+ provides the technical basis for such focusing, which, however, may only be controlled by very general criteria, such as probabilities.

One of the major goals of future work therefore is to investigate how to exploit domain-specific knowledge, *e.g.* dependent faults. Furthermore, the *purpose* of diagnosis should also be considered; if the goal is to perform a quick and reliable reconfiguration of the network, a possible discrimination between different diagnoses may be irrelevant. This consequently affects the ranking of diagnoses and the termination criteria of the system.

5.7 *Refinement of analysis*

Different fault situations may require models of different granularity. The majority of the cases considered can be solved by using simplified models, as sketched in Sections 3 and 4. The concept of fault distance used by the model of the breaker is based on the hypothesis that the impedance measured by the distance protection actually corresponds to the distance of the fault. This assumption may be invalid in the case of a short-circuit with a high impedance, and in that case this model may lead to wrong or even no diagnoses. Clearly, it is possible to describe a model which does not rely on this simplifying assumption, and thus also allows the diagnosis of high-impedance faults [18].

The important thing, however, is that in GDE^+ this more detailed model does not automatically replace the coarser, but surprisingly useful, model. Actually, it is possible to define multiple alternative models and select a model that is appropriate for the current purpose. This allows us to start with the simplest model and, in case this is not successful or the purpose of analysis has been changed, to switch to a more detailed model. The underlying ATMS avoids inconsistencies between different models and preserves intermediate results, which are independent of this change (Struss [18] illustrates this model switching).

In addition to such a refinement of behavioural descriptions, GDE^+ allows the structural refinement of models. Subsystems may be described as black-box models; in case such a subsystem is suspected, it may be expanded and diagnosis focused on its substructure. This recursive applicaton of GDE^+ (and the refinement of behavioural descriptions) allows the system to cope with complex systems and has been successfully used in another case study to diagnose systems with a large number of components (approximately 1000) [13]. The necessity and utility of this technique for power transmission networks is a topic for further investigation.

5.8 *Adaptability to a changed topology*

As the topology of a power transmission network is changed according to the current load situation, a system for fault localisation must be adapable to the actual structure of the network. As in GDE^+ the behavioural models of components and the structural description of the system to be diagnosed are represented separately, this requirement is already satisfied by the architecture of the diagnosis system.

5.9 Adaptability to new components and technologies

An even more important requirement is an adaptability to changes of a technological nature, such as the presence of special components or the introduction of new technology. This is one of the major limitations of first-generation systems; the rule base has to be checked to see whether the changed situation affects the applicability of rules (which requires that this question can be decided in case there is no experience with the new technology). This is absolutely necessary, as the new technology cannot simply be anticipated as an exception to a rule.

The principal advantage of the model-based approach is based on the use of a domain-specific, but system-independent, component library as part of the knowledge base. The introduction of a new type of component is thus simply the definition of a corresponding model in this library (and the creation of corresponding instances in the structural description).

Whereas the development and maintenance of complex rule-based systems quickly results in an unreasonable, complex, expensive, and unreliable procedure (see the experience with XCON [19]), model-based systems allow incremental, modular and efficient changes and improvements to a knowledge base. The development of DPNet has exploited and demonstrated this potential.

6 Conclusions

We have shown that a model-based approach is appropriate for the task of fault localisation in power transmission networks by demonstrating how this approach satisfies the particular requirements of the problem. The model-based approach allows the development of diagnostic systems with extended competence (treatment of unknown and multiple faults), while significantly changing the development process. In many cases, the reduced development overheads and the improved quality of the results are necessary preconditions to allowing an industrial application of knowledge-based technology. The work on DPNet up to now has impressively confirmed this claim.

7 Acknowledgments

The authors would like to thank Stefano Cermignani, Mike Chantler, Gerhard Friedrich, Roy Leitch, Ulrich Löwen, Monika Pfau, Alberto Stefanini and Giorgio Tornielli for their helpful discussions.

This work was supported in part by BMFT (ITW 9001 A9) and CEC (P5143, 3249).

8 References

[1] DAVIS, R., and HAMSCHER, W. C.: 'Model-based reasoning: troubleshooting' *in* SHROBE, H. E. (Ed.): 'Exploring artificial intelligence: survey talks from the national conferences on artificial intelligence' (Morgan Kaufmann, San Mateo, California, 1988) pp. 297–346

[2] GERO, J. S. (Ed.): 'Artificial intelligence in engineering: diagnosis and learning' (Amsterdam and Boston, 1988)

[3] CIGRE Working Group 38.02: 'An international survey of the present status and the perspective of expert systems on power system and analysis and techniques', *Electra*, 1989, **123**, pp. 70–94

[4] Conf. Proc. on Expert Systems Applications in the Electric Power Industry, Orlando, Florida, 1989

[5] ETZ, 1990, **111**, (5)

[6] BRAMER, M. A., PEARCE, J. S., MUIRDEN, D., PLATTS, J. C., and VIPOND, D. L.: 'FAUST—an expert system for diagnosing faults in an electricity supply system' *in* KELLY, B., and RECTOR, A. L. (Eds.): 'Research and development in expert systems V' (Cambridge University Press, 1989)

[7] GALLASTEGUI, I., LARESGOITI, I, PÉREZ, J., AMANTEGU, J., and ECHÁVARRI, J.: 'Operating experience of an expert system for fault analysis in electrical networks'. Proc. Int. Working Conf. on Expert Systems in Electrical & Power Systems, Avignon, France, May 1989

[8] RUMPEL, D., and SUN, J. R.: 'Netzleittechnik' (Springer-Verlag, 1989)

[9] STRUSS, P., and DRESSLER, O.: 'Physical negation—integrating fault models into the general diagnostic engine', Proc. IJCAI 89, Detroit, Michigan, 1989

[10] DIETZE, P., LÖWEN, U., and STELZNER, P.: 'Untersuchungen von Netzstörungen mit einem Expertensystem,' ETZ, 1990, **111**, (5)

[11] PHILIPPOW, E.: 'Taschenbuch der Elektrotechnik, Band 6: Systeme der Elektroenergietechnik' (Carl Hanser, München & Wien, 1982)

[12] VON RZIHA, E.: 'Starkstromtechnik' (Wilhelm Ernst & Sohn, Berlin, 1955) Vol. 1, 8th edn.

[13] STRUSS, P.: 'Model-based diagnosis—introduction and survey of recent work by ARM'. Technical Report INF 2 ARM-16-90, Siemens, München, 1990

[14] DE KLEER, J., and WILLIAMS, B. C.: 'Diagnosing multiple faults', *Art. Intell.*, 1987, **32**, (1), pp. 97–130

[15] DE KLEER, J.: 'An assumption-based TMS', *Art. Intell.*, 1986, **28**, (2), pp. 127–162

[16] STRUSS, P.: 'A framework for model-based diagnosis.' Technischer Bericht INF 2 ARM-10-88, Siemens, München, 1988

[17] FREITAG, H.: 'A generic measurement proposer' *in* GOTTLOB, G., and NEJDL, W. (Eds.): 'Expert systems in engineering' (Heidelberg, 1990)

[18] STRUSS, P.: 'What's in SD? towards a theory of modeling for diagnosis' *in* HAMSCHER, W. C., CONSOLE, L., and DE KLEER, J. (Eds.): 'Readings in model-based diagnosis' (Morgan Kaufmann, 1992)

[19] BACHANT, J., and SOLOWAY, E.: 'The engineering of XCON', *Commun. ACM*, 1989, **32**, (3)

The paper was first received on 17 June and in revised form on 1 July 1992.

Anton Beschta, Oskar Dressler, Hartmut Freitag and Michael Montag are with Siemens AG, Otto-Hahn-Ring 6, W-8000 Munich 83, Germany; Peter Struß is with the Computer Science Department, Technical University of Munich, Orleansstraße 34, W-8000 Munich 80, Germany.