

Event-Based Reasoning for Short Circuit Diagnosis in Power Transmission Networks

Gianfranco Lamperti

Dipartimento di Elettronica per l'Automazione
Università degli Studi di Brescia
via Branze 38, I-25123 Brescia, Italy
lamperti@bsing.ing.unibs.it

Paolo Pogliano

CISE s.p.a.
via Reggio Emilia, 39
I-20090 Segrate (Milano), Italy
paolo@sia.cise.it

Abstract

This document introduces an original, model-based technique for short circuit diagnosis in power transmission networks. The diagnostic task is viewed as a bottom-up interpretation process of available observations. Three different interpretation layers have been defined: *local interpretation*, *global interpretation*, and *heuristic interpretation*. This paper presents the first two layers. The local interpretation focuses on the behavior of single protection components that are distributed in the power network and operate when a short circuit occurs. The global interpretation provides a global behavior of the protection apparatus by combining consistent local behaviors so that given *interface constraints* among components are met. Finally, the heuristic interpretation is meant both to shrink the cardinality of the global interpretation on the basis of application-dependent *heuristic constraints*, and to eventually localize the short circuit and possibly faulty protection components within the transmission network. The implementation of the proposed technique is under way. The resulting system will be tested by ENEL, the Italian electricity board.

1 Introduction

To avoid service interruptions caused by short circuits in a power transmission network, it is necessary to isolate faults as soon as possible. For this reason, there are a number of protection components distributed over the network. The protection system is in charge of detecting dangerous conditions, of disconnecting a component (such as a line, a bus, a transformer or a generation group) as soon as its operating starts to be dangerous, and of keeping in operation non-faulty components as much as possible, in order to avoid a black-out. This is

achieved by tripping the circuit breaker associated with each protection. Each protection has to protect mainly one component, but must also operate as a backup to other protections nearby.

When intervening, all the protections and the breakers send logical signals reporting their operation: these are recorded by an event recorder and transmitted to a *Regional Control Center* (RCC) where they are used for fault localization. Records, called *messages*, consist of a unique address of the event source, an event code, and possibly a timestamp. Operators of the RCC have to decide within one minute where the fault is located and what recovery actions have to be applied. Normally, due to time constraint reasons, only the current status of the breakers is considered, but this is not generally sufficient for fault localization. It is recognized that as a result of this an automated tool for fault diagnosis is an important requirement.

To this end, a technique based on the modeling of the protection system and on the interpretation of the reaction of the protection system to the occurrence of a short circuit has been defined. Three different interpretation layers are involved: the *local interpretation*, the *global interpretation*, and the *heuristic interpretation*. The focus of this paper is on the first two layers.

2 Modeling the protection system

Each protection component is modeled by means of a finite state machine (FSM), called *reactive model*, which is driven by a series of *input* events and generates some other *output* events when transitions between (*reactive*) states occur. For example, a distance protection is started when the impedance on the relevant line goes under a certain threshold (input event). After that, the protection is expected to change state at fixed time intervals (clock input events) and, possibly, to trip the associated breaker (tripping output event). Note that the tripping output event for the protection generates an input (tripping) event for the breaker: we say that the output event is *exported* by the protection and is

imported by the breaker. Reactive states can be either *steady* or *unsteady*. This classification improves the semantics of reactive models and indirectly poses a number of constraints on the behavior of protection components which can be conveniently exploited by the interpretation algorithm. A *null event* ϵ is a formal notation to specify the absence of events.

Formally, a reactive model M is a record of four elements: $M = (\Sigma, I, O, \delta)$, where Σ is the set of the reactive states, I is the set of input events, O is the set of output events, and δ is the transition function: $\delta : \Sigma \times (I \cup \epsilon) \times (2^O \cup \epsilon) \rightarrow \Sigma$, (2^O is the power-set of O). This means that a transition T from state S_1 to state S_2 is in general triggered by an input event i and generates, before changing state, the list of output events $\langle o_1, o_2, \dots, o_n \rangle$. This is denoted by $T = S_1 \xrightarrow{i | o_1, o_2, \dots, o_n} S_2$. Furthermore, $\Sigma = \Sigma_s \cup \Sigma_u$, $\Sigma_s \cap \Sigma_u = \emptyset$, where Σ_s and Σ_u denote the set of steady and unsteady states respectively.

The specification of the reactive model of each type of protection component and the instantiation of such classes of components into a given network topology yields the whole model of the diagnosed system.

Each reactive model describes both the correct and faulty behavior of a class of protection components. Furthermore, the model allows for uncertainty due to possible loss of messages. Specifically, if during a transition $T_1 = S_1 \xrightarrow{i | m} S_2$, message m may be lost, then the δ function will include an additional transition $T_2 = S_1 \xrightarrow{i | \epsilon} S_2$.

In the sequel, we assume the completeness of each reactive model, that is to say, the reactive model is assumed to describe each possible reaction.

3 Misbehaviors, observations, histories

When a short circuit occurs, we say there is a *misbehavior* μ of the transmission network. The set of protective components reacting to a misbehavior μ is called the *misbehavior extent* of μ , and is denoted by $extent(\mu)$. After the specification of each reactive model, it is possible to interpret a given sequence of observations, the messages, in order to eventually find out the sequence of transitions covered by the involved protection components during μ . When a short circuit occurs, each actual protection component is expected to react in a way that corresponds to an instantiation of the model, called the *local history* of the component. As such, a local history is a sequence of transitions within the reactive model. The initial and final states of the history are required to be steady. The local history is derived on the basis of the *local observation* of the component P , namely a list of messages relevant to P .

Definition 1 (*local history*) Let $M = (\Sigma, I, O, \delta)$ be a reactive model of a protective component. A *lo-*

cal history of M is a (possibly empty) sequence $h = \langle T_1, T_2, \dots, T_n \rangle$ of transitions in M so that h conforms to the following *morphology constraints*:

1. *Determinism*. Each transition $T_i \in h$ is adorned with at most one allowed input event.
2. *Contiguity*. For each pair of contiguous transitions T_i, T_{i+1} in h , the final state of T_i coincides with the initial state of T_{i+1} .
3. *Stability*. Both the initial state of T_1 and the final state of T_n are in Σ_s .

If the sequence of transitions is empty, h is called a *null history*, and is denoted by ϵ .

Definition 2 (*local observation*) Let M be a reactive model of a protection component P . A *local observation* obs of P , $obs(P) = \langle o_1, o_2, \dots, o_n \rangle$, is a sequence of temporally ordered observable output events, generated by a local history of M . A local observation which is empty is called a *null observation* and is denoted by ϵ .

4 Local Interpretation

The first step towards the detection of the short circuit is represented by the local interpretation of every local observation. A local interpretation algorithm (*lia*) has been defined. The algorithm takes as input a local observation $obs(P)$ and generates a set h^* of consistent local histories on the basis of the relevant reactive model. In general, a local interpretation gives rise to several consistent local histories, namely: $h^* = lia(obs(P)) = [h_1, h_2, \dots, h_n]$.

Interestingly, a reactive model M of P can be formally interpreted as the grammar of a regular language L where an observation $obs(P)$ represents a phrase of L [Hopcroft and Ullman, 1979]. Consequently, the problem of determining the local interpretation of $obs(P)$ corresponds to the recognition of a string of tokens (the messages) as a phrase f of L . What is essential in the local interpretation is to keep track of the sequence h of transitions which are relevant to the recognition of f .

However, a model M may possibly be non-deterministic, involving transitions which are not associated to any observable event. These are called *silent transitions*. Moreover, silent transitions may possibly form *silent cycles* so that $h^* = lia(obs(P))$ may potentially include an unlimited number of local histories, all of them being consistent with $obs(P)$. For example, if $h_1 = S_1 \xrightarrow{i_1 | m_1} S_2 \xrightarrow{i_2 | m_2} S_1$ is consistent with $obs(P)$, and the model of P includes the silent transition $S_2 \xrightarrow{i | \epsilon} S_2$, then $h_2 = S_1 \xrightarrow{i_1 | m_1} S_2 \xrightarrow{i | \epsilon} S_2 \xrightarrow{i_2 | m_2} S_1$ is consistent with $obs(P)$ as well, since the silent transition is immaterial for the observation.

In the *lia* algorithm, a non-deterministic reactive model M_n is first translated into a deterministic model

M_d , in which a connected subgraph of M involving only silent transitions (called a *silent subgraph*) is collapsed into a new reactive state. The interpretation of $obs(P)$ is therefore performed on the basis of the new deterministic model M_d , instead of the original non-deterministic model M_n . However, once a local history h is determined, for example $h = S_1 \xrightarrow{i_1 | m_1} S \xrightarrow{i_2 | m_2} S_2$, if S corresponds to a collapsed silent subgraph, then the actual local history relevant to the original model M_n will include a so-called *hypertransition*, namely the totality of the paths in S which start at an initial state S_i and end at a final state S_f such that $S_1 \xrightarrow{i_1 | m_1} S_i$ and $S_f \xrightarrow{i_2 | m_2} S_2$ are legal transitions in M_n . A hypertransition from S_i to S_f is denoted by $S_i \Rightarrow S_f$. A local history involving a hypertransition is called a *hyperhistory*. In the above example, the hyperhistory will be: $S_1 \xrightarrow{i_1 | m_1} S_i \Rightarrow S_f \xrightarrow{i_2 | m_2} S_2$. A *hyperhistory graph* is the graph representation of the hyperhistory and is obtained by connecting the graph representation of the involved (hyper) transitions to each other.

5 Global Interpretation

Once observations of protective components reacting to a misbehavior μ are interpreted by *lia*, the set of consistent global histories must be generated.

Definition 3 (*global history*) Let $[P_1, P_2, \dots, P_m] = extent(\mu)$ be the set of protective components involved in a misbehavior μ . A *global history* H of μ , $H(\mu) = (h_1, h_2, \dots, h_m)$, is the aggregation of the local histories $h_1, h_2, \dots, h_m$ relevant to $P_1, P_2, \dots, P_m$ respectively.

Definition 4 (*global history domain*) Let $[obs(P_1), obs(P_2), \dots, obs(P_n)]$ be the set of local observations relevant to $extent(\mu)$. The *global history domain* of μ , $H_d(\mu)$, is the cartesian product of the local interpretations of $P_1, P_2, \dots, P_n$, namely: $H_d(\mu) = lia(obs(P_1)) \times lia(obs(P_2)) \times \dots \times lia(obs(P_n))$.

Definition 5 (*global interpretation*) The *global interpretation* of μ , $H^*(\mu)$, is a relation among the local interpretations $[lia(obs(P_1)), lia(obs(P_2)), \dots, lia(obs(P_n))]$, namely a subset of the global history domain $H_d(\mu)$: $H^*(\mu) \subseteq H_d(\mu) = lia(obs(P_1)) \times lia(obs(P_2)) \times \dots \times lia(obs(P_n))$ so that the following conditions hold:

1. $\forall H = (h_1, h_2, \dots, h_n) \in H^*(\mu)$, H is globally consistent with respect to the interface constraints, and
2. $\neg \exists$ a globally consistent $H' \in H_d(\mu)$ such that $H' \notin H^*(\mu)$.

A systematic approach for selecting non spurious global histories is to apply a *balance* function to every pair of local histories (h_1, h_2) , where $h_1 \in lia(obs(P_1))$, $h_2 \in lia(obs(P_2))$, and P_1 is connected to P_2 . In the simplest case in which both h_1 and h_2 are plain local histories (not

hyperhistories), *balance* is a boolean function checking the consistency of h_1 and h_2 with respect to the interface constraints relevant to h_1 and h_2 only. Informally, each event of h_1 exported by P_1 to P_2 is required to be imported by h_2 and vice versa.

To this end, a *consistency matrix* is used, where the result of each application of *balance* is recorded. A consistency matrix Ξ_μ is relevant to a misbehavior μ and is composed of rows and columns associated with local interpretations relevant to $extent(\mu)$. Thus, $\Xi_\mu(h_1, h_2)$ denotes the element of the matrix which corresponds to histories h_1 and h_2 . The *inconsistency set* of Ξ_μ , denoted by $\mathfrak{I}(\Xi_\mu)$ is the set composed of those local history pairs (h_1, h_2) for which $balance(h_1, h_2) = false$. On the basis of the inconsistency set we are allowed to remove the inconsistent global histories from the global history domain $H_d(\mu)$. This operation corresponds to the deletion of the global histories which include a pair of the inconsistency set. Therefore, $H \in H_d(\mu)$ is a consistent global history if and only if $\neg \exists (h_1, h_2) \in \mathfrak{I}(\Xi_\mu)$ so that $(h_1, h_2) \subseteq H$.

However, the trouble arises when the approach based on the consistency matrix is extended to deal with hyperhistories: in this case, *balance* is expected to return more information than a single boolean value. For example, if h and h^* denote a local history and a hyperhistory respectively, three different cases are possible for $balance(h, h^*)$, namely:

1. $balance(h, h^*) = \emptyset$. This is the easiest case since it establishes that every candidate global history involving both h and h^* is to be discarded. $\Xi_\mu(h, h^*)$ is marked with the empty set symbol.
2. $balance(h, h^*) = [h_1, h_2, \dots, h_n]$. In that case the hyperhistory is *resolved* and a finite number of local histories is determined (possibly a singleton). $\Xi_\mu(h, h^*)$ is marked with a reference to this set.
3. $balance(h, h^*) = h'^* = [h_1, h_2, \dots, h_n, h_{n+1}, \dots]$, $h'^* \subseteq h^*$. The hyperhistory is *reduced* but not resolved. A reference to the hyperhistory graph of h'^* is put in $\Xi_\mu(h, h^*)$.

But the real problem is how to extend the global history generation process once the Ξ_μ has been filled with this information. When all the local interpretations are composed of a limited number of local histories, we are able to enumerate all the candidate global histories $H \in H_d(\mu)$. In that case we say that we have an *extensional* representation. By contrast, the possible inclusion of hyperhistories in the local interpretation forces us to maintain a so called *intensional* representation, by means of the hyperhistory graph.

Definition 6 (*hyper global history*) Let $H \in H_d(\mu)$ be a candidate global history. If there exists a hyperhistory $h^* \in H$, then H is called a *hyper global history*.

Extending the scope of *balance* to cope with hyperhistories, requires Ξ_μ to record a more complex set of information. Specifically, the codomain of *balance* must be represented by a set ∇ including elements which can be classified as follows (h^* denotes either a hyperhistory or a finite set of local histories): $\nabla = [\text{true}, \text{false}, \emptyset, h^*, (\emptyset, \emptyset), (h^*, \emptyset), (\emptyset, h^*), (h_1^*, h_2^*)]$ which in turn can be seen as the union of three parts: $\nabla = \nabla_{h,h'} \cup \nabla_{h,h^*} \cup \nabla_{h^*,h'^*}$, whereby: $\nabla_{h,h'} = [\text{true}, \text{false}]$, $\nabla_{h,h^*} = [\emptyset, h^*]$, and $\nabla_{h^*,h'^*} = [(\emptyset, \emptyset), (h^*, \emptyset), (\emptyset, h^*), (h_1^*, h_2^*)]$, corresponding respectively to the domain of two local histories, a local history and a hyperhistory, and two hyperhistories.

Definition 7 (*resolvable hyper global history*) Let H^* be a candidate hyper global history involving a number of infinite hyperhistories: $H^* = (h_1, h_2, \dots, h_k, h_{k+1}^*, h_{k+2}^*, \dots, h_n^*) = (h_1, h_2, \dots, h_k) \times h_{k+1}^* \times h_{k+2}^* \times \dots \times h_n^* = [H_1, H_2, \dots, H_i, H_{i+1}, \dots]$. If it is possible to enforce the interface constraints so that H^* is reduced to a finite set $H^+ = [H'_1, H'_2, \dots, H'_p] \subset H^*$, then H^* is called a *resolvable hyper global history*.

Definition 8 (*symbolic intersection*) Let h_1^* and h_2^* be two hyperhistories relevant to the same protective component P . The hyperhistory graph corresponding to $h^* = h_1^* \cap h_2^*$ is called the *intersection graph* of h_1^* and h_2^* . The process of deriving the intersection graph is called *symbolic intersection*.

Definition 9 (*residue*) Let H^* be a candidate hyper global history. Let $h^* \in H^*$ be a hyperhistory embraced by H^* . Let $h_1^*, h_2^*, \dots, h_n^*$ be the list of reduced hyperhistories relevant to *balance*(h_i^*, h^*), $h_i^* \in H^*$, $h_i^* \neq h^*$. The *residue* of h^* in the context of H^* , denoted by $\mathcal{R}(h^*, H^*)$, is the symbolic intersection of $h_1^*, h_2^*, \dots, h_n^*$, namely $\mathcal{R}(h^*, H^*) = h_1^* \cap h_2^* \cap \dots \cap h_n^*$.

5.1 Global Interpretation Algorithm

The *gia* global interpretation algorithm can be concisely described as a function having as input a global history domain $H_d(\mu)$ and returning a global interpretation $H^*(\mu) \subseteq H_d(\mu)$ by means of the following steps:

1. $H^*(\mu) \leftarrow H_d(\mu)$;
2. create the consistency matrix Ξ_μ ;
3. by applying *balance*, associate with each element of Ξ_μ a relevant instantiation of a symbol in ∇ ;
4. build the inconsistency set $\mathfrak{Z}(\Xi_\mu)$ as composed of the pairs of (possibly hyper) histories (h_i, h_j) for which: *balance*(h_i, h_j) = *false* $\vee$ *balance*(h_i, h_j) = $\emptyset \vee$ *balance*(h_i, h_j) = $(\emptyset, h_{j'}) \vee$ *balance*(h_i, h_j) = $(h_{i'}, \emptyset) \vee$ *balance*(h_i, h_j) = $(\emptyset, \emptyset)$;
5. remove from the global interpretation those global histories H including a pair $(h_i, h_j) \in \mathfrak{Z}(\Xi_\mu)$, namely:

$$H^*(\mu) \leftarrow H^*(\mu) - [H \mid \exists (h_i, h_j) \subseteq H, (h_i, h_j) \in \mathfrak{Z}(\Xi_\mu)];$$

6. for each hyperhistory h^* of every hyper global history $H^* \in H^*(\mu)$, replace h^* with the corresponding residue, namely: $\forall H^* \in H^*(\mu), \forall h^* \in H^*, h^* \leftarrow \mathcal{R}(h^*, H^*)$.

6 Comparison with related work

Several attempts to develop automated support tools for fault diagnosis in power transmission networks are reported in the literature. Somewhat surprisingly, they are based on a variety of different technological approaches including neural networks [Kezunovic *et al.*, 1994], fuzzy expert systems, [Cho *et al.*, 1994], Petri nets [Wang and Trecat, 1994], model-based diagnosis [Tornielli *et al.*, 1994; Beschta *et al.*, 1993], and temporal reasoning techniques [Baroni *et al.*, 1995].

In [Kezunovic *et al.*, 1994] a neural network is trained to recognize typical fault patterns, by processing in input the values of voltage and current measured in a given substation. Therefore this system aims mainly to provide a detailed characterization of a local fault, within the context of a specific substation and of its electrical and structural configuration, rather than to perform a diagnosis at the level of the overall network, considering the behavior of the protection system. Moreover, even within this limited context, it seems that a major difficulty of this approach is related to the huge amount of preliminary work required, since the neural network must be trained on a wide number of different cases each time it has to be applied to a different substation configuration.

In [Cho *et al.*, 1994], explicit relations between the location of a fault in the networks and the operation of protections and breakers are represented through sagittal diagrams. Temporal behavior of the components is not considered, however in order to encompass the fact that some protections should operate before others when a fault is located in a given position, a numerical label is associated with each relation, representing the possibility that each protection operates when the fault is in the considered position. Different diagnostic hypotheses can be generated starting from observations and are ranked according to their possibility value. This approach, not including an explicit model of the diagnosed system, suffers from all the well-known limitations of first-generation, rule-based, diagnostic systems (see for instance a discussion in [Beschta *et al.*, 1993]).

An alternative approach [Wang and Trecat, 1994] resorts to Petri nets in order to build a model of the protection system behavior. However this model is oversimplified: temporal aspects are not considered and the overall protection system behavior is modeled as a simple two steps activity. More detailed models are used

in [Tornielli *et al.*, 1994; Beschta *et al.*, 1993], where a classical model-based diagnosis approach is adopted, using the GDE+ diagnostic engine [Struss and Dressler, 1989]. The approaches differ in the nature of the model adopted. In [Tornielli *et al.*, 1994], the behavior of network components is defined in terms of equations between the admittance values seen from different points, whereas the behavior of the components of the protection system is related to the impedance value seen by the component itself. The approach of [Beschta *et al.*, 1993] is quite similar but a qualitative evaluation of the fault distance is used instead of the admittance value in the component models. A major problem with these approaches is that the component models adopted, being quite abstract and not very detailed, do not take into account the rather complex behavior of the protection system and are not able to produce detailed diagnoses about its malfunctions.

Detailed models of component behaviors, including temporal features, are used in [Baroni *et al.*, 1995], where an original diagnostic algorithm for time-varying systems is proposed and applied to power networks. This algorithm is able to produce a detailed temporal reconstruction of the events occurred in the network, but it relies on the assumption that the timestamp of each message received is available, that is not always the case in practice. By contrast, the approach presented in this paper, while retaining the advantages of using detailed component models, does not rely on the availability of temporal information.

An approach whose rationale has some significant similarities with the work presented here has been developed in the frame of diagnosis of communication protocols [Riese, 1993]. In this context the problem consists of verifying, starting from a set of observations, if an implementation of a communication protocol is compliant with its formal definition. The protocol is modeled through a FSM and observations are matched with the possible evolution of the FSM, simulated on a discrete time scale. If none of the possible evolutions matches with the observations, the algorithm searches for modifications of the FSM that can explain the observations: each modification corresponds to a lack of compliance with respect to the standard. A detailed comparison with this work is beyond the scope of the present paper. However it can be noted that, even though the nature of the diagnostic problem is rather different, we share the use of FSM-based techniques, involving observation matching, propagation, and interpretation, in order to deal with complex diagnostic problems of time-varying systems.

This choice differs with respect to other approaches proposed in literature for diagnosis of dynamic systems, a topic that has been addressed by many researchers in

recent years [Dvorak and Kuipers, 1989; Gluckenbiehl and Schäfer-Richter, 1990; Hamscher, 1991; Ng, 1991; Lackinger and Nejd, 1991; Dressler and Freitag, 1994].

As a matter of fact, even though the importance of dynamic system diagnosis and the need of specific techniques for this problem was recognized early [Pan, 1984; Hamscher and Davis, 1984], the *classical* theory of diagnosis [Reiter, 1987; de Kleer and Williams, 1987] has been conceived for static systems only. Subsequent attempts to overcome this limitation have been, however, strongly influenced by the approach initially adopted for static systems. In fact, most of these approaches [Dvorak and Kuipers, 1989; Hamscher, 1991; Ng, 1991; Lackinger and Nejd, 1991; Friedrich and Lackinger, 1991; Dressler and Freitag, 1994] share a common rationale: Reiter's algorithm is simply applied iteratively to subsequent instantaneous snapshots of the behavior of the system, generated through a suitable dynamic model. This method presents however some significant drawbacks, since it is unable to deal with faults whose manifestation involves several time instants and does not exploit efficiently knowledge about system dynamic behavior. An approach closer to our ideas is presented in [Gluckenbiehl and Schäfer-Richter, 1990] where the diagnostic activity for a dynamic system is seen as the task of reconstructing the history of the system starting from some temporally located information about system attributes and using a model of system behavior. However the behavioral representation adopted is made up of *if-then* rules including temporal information and suffers therefore from limited expressiveness.

7 Discussion and conclusions

Although featuring an original approach to model-based diagnosis of event-based dynamic systems, the proposed method relies on the availability of correct observations. More precisely, besides the intrinsic unavailability of messages, the method can cope with lost messages, but it assumes that those messages which are available are correct. The knowledge about possible missing messages must be specified in the model, as illustrated in Section 2. Intuitively, the richer is the set of observations, the smaller is the number of diagnostic alternatives produced. Furthermore, the algorithm is clearly designed for domains where discrete (and complete) models of correct and faulty behaviors are available.

Another important point concerns the resolvability of hyper global histories, which can be guaranteed by enforcing some *safety* constraints on a reactive model M . Intuitively, safety may be achieved when either no silent cycles are incorporated in M or when, under additional assumptions, interface events are associated with silent transitions within silent cycles.

Finally, a formal analysis of the complexity of the

method and real benchmarks of the implementation are expected for the future. However it is already evident that the performance of the algorithm will strongly depend on the complexity of the reactive models exploited, and the extent to which they allow an accurate interpretation of the available observations. It is also worth pointing out that local interpretations might in principle be generated in parallel, either by a single parallel machine or by several machines installed in different substations.

Acknowledgments

The present work is partially supported by the Commission of the European Union under the ongoing Esprit III Project 8491 *Timely* (Time-Constrained Integrated Management of Large-Scale Systems).

References

- [Baroni *et al.*, 1995] P. Baroni, U. Canzi, and G. Guida. SHORT: a knowledge-based system for fault diagnosis in power transmission networks. In *Proc. of the Eighth Int. Symp. on Artificial Intelligence ISAI 95*, pages 199-208, Monterrey, MX, 1995.
- [Beschta *et al.*, 1993] A. Beschta, O. Dressler, H. Freitag, M. Montag, and P. Struss. A model-based approach to fault localisation in power transmission networks. *Intelligent Systems Engineering*, pages 3-14, Spring 1993.
- [Cho *et al.*, 1994] H.-J. Cho, J.-K. Park, and H.-J. Lee. A fuzzy expert system for fault diagnosis of power systems. In *Proc. of the Int. Conf. on Intelligent System Application to Power Systems ISAP 94*, pages 217-226, Montpellier, F, 1994.
- [Dressler and Freitag, 1994] O. Dressler and H. Freitag. Prediction sharing across time and contexts. In *Proc. of the Twelfth Nat. Conf. on Artificial Intelligence AAAI-94*, pages 1136-1141, Seattle, WA, 1994.
- [Dvorak and Kuipers, 1989] D. Dvorak and B. Kuipers. Model-based monitoring of dynamic systems. In *Proc. of IJCAI-89*, pages 1238-1243, Detroit, MI, 1989.
- [Friedrich and Lackinger, 1991] G. Friedrich and F. Lackinger. Diagnosing temporal misbehavior. In *Proc. of IJCAI-91*, pages 1116-1122, Sydney, Australia, 1991.
- [Gluckenbiehl and Schäfer-Richter, 1990] T. Gluckenbiehl and G. Schäfer-Richter. SIDIA: Extending prediction based diagnosis to dynamic models. *Working Notes of the First Int. Workshop on Principles of Diagnosis*, pages 74-82, Stanford, CA, 1990.
- [Hamscher, 1991] W. Hamscher. Modeling digital circuits for troubleshooting. *Artificial Intelligence*, 51 (1-3):223-271, 1991.
- [Hamscher and Davis, 1984] W. Hamscher and R. Davis. Diagnosing circuits with state: an inherently underconstrained problem. In *Proc. of the Fourth Nat. Conf. on Artificial Intelligence AAAI-84*, pages 142-147, Austin, TX, 1984.
- [Hopcroft and Ullman, 1979] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley, Reading, MA, 1979.
- [de Kleer and Williams, 1987] J. de Kleer and B.C. Williams. Diagnosing multiple faults, *Artificial Intelligence*, 32(1):97-130, 1987.
- [Kezunovic *et al.*, 1994] M. Kezunovic, I. Rikalo, and D.J. Sobajic. Neural network applications to real-time and off-line fault analysis. In *Proc. of the Int. Conf. on Intelligent System Application to Power Systems ISAP 94*, pages 29-36, Montpellier, F, 1994.
- [Lackinger and Nejd, 1991] F. Lackinger and W. Nejd. Integrating model-based monitoring and diagnosis of complex dynamic systems. In *Proc. of IJCAI-91*, pages 1123-1128, Sydney, Australia, 1991.
- [Ng, 1991] H.T. Ng. Model-based, multiple-fault diagnosis of dynamic, continuous physical devices. *IEEE Expert*, 6(6):38-43, 1991.
- [Pan, 1984] J. Y.-C. Pan. Qualitative reasoning with deep-level mechanism models for diagnosis of mechanism failures. In *Proc. of the First IEEE Conf. on AI Applications*, pages 295-301, Denver, CO, 1984.
- [Reiter, 1987] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57-95, 1987.
- [Riese, 1993] M. Riese. Diagnosis of communicating systems: dealing with incompleteness and uncertainty. In *Proc. of IJCAI-93*, pages 1480-1485, Chambéry, F, 1993.
- [Struss and Dressler, 1989] P. Struss and O. Dressler. Physical negation-integrating fault models into the general diagnostic engine. In *Proc. of IJCAI-89*, pages 1318-1323, Detroit, Michigan, 1989.
- [Tornielli *et al.*, 1994] G. Tornielli, L. Capetta, S. Cermignani, A.S. Fabiano, and R. Schinco. An interval-based approach to fault diagnosis of power systems. In *Proc. of the Int. Conf. on Intelligent System Application to Power Systems ISAP 94*, pages 809-816, Montpellier, F, 1994.
- [Wang and Trecat, 1994] J.P. Wang and J. Trecat. A parallel fault diagnosis expert system for one dispatch center. In *Proc. of the Int. Conf. on Intelligent System Application to Power Systems ISAP 94*, pages 201-208, Montpellier, F, 1994.