SHORT: A KNOWLEDGE-BASED SYSTEM FOR FAULT DIAGNOSIS IN POWER TRANSMISSION NETWORKS

Pietro Baroni, Ulrico Canzi, Giovanni Guida

Università di Brescia, Dipartimento di Elettronica per l'Automazione
Via Branze 38, I-25123 Brescia, Italy

ABSTRACT

This paper presents SHORT a knowledge-based system for fault diagnosis in power transmission networks. Differently from previous approaches to this very complex diagnostic problem, SHORT includes a knowledge representation formalism, named plot, capable of taking into account temporal aspects of component behavior, and solves the diagnostic problem as an instance of a more general class of temporal reasoning problems, called history reconstruction. Both plot formalism and history reconstruction method are illustrated in the paper through an application example. SHORT has been tested on both simulated and real cases of faults in an Italian regional transmission network, showing very good diagnostic capabilities.

1. INTRODUCTION

Fault diagnosis in power transmission networks is a very complex problem, of huge practical relevance; for this task many attempts have been made to develop automated support tools. These efforts have been based on a wide variety of approaches, most of which fall in the domain of intelligent systems [8]. In a recent specialized conference [4] many systems for fault diagnosis in power systems have been presented, based on very different technological solutions, such as model-based diagnosis [9], neural networks [5], Petri nets [10], fuzzy systems [1]. Somewhat surprisingly, temporal reasoning aspects, which are intrinsic to this problem, have not been considered in these proposals, thus giving rise to important limitations in their diagnostic capabilities. In this paper we present SHORT, a knowledge-based system, which tries to overcome these limitations using a novel knowledge representation and temporal reasoning technique. SHORT is the result of a research effort aimed at developing diagnostic applications that merge model-based and temporal reasoning features. SHORT exploits a properly designed knowledge representation formalism, called plot, which can capture temporal aspects of behavioral knowledge about the components of a complex dynamic system. The problem of diagnosis of dynamic systems has then been framed into a more general class of problems, named history reconstruction, and a general method for its solution has been devised and implemented. SHORT has been extensively tested on both simulated and real fault cases of a regional power transmission network, where it has shown fully satisfactory diagnostic capabilities. The paper is organized as follows. Section 2 gives an overview of the problem of fault diagnosis in power transmission networks. Section 3 describes the plot knowledge representation formalism. Section 4 introduces the history reconstruction problem and illustrates the reasoning mechanism designed to solve it. Section 5 gives some details about system implementation and experimental results. Conclusions and directions of future work are presented in section 6.

2. FAULT DIAGNOSIS IN POWER TRANSMISSION NETWORKS

2.1 The Protection System

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 mechanisms 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 it begins to operate in a dangerous condition, 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 to each protection. Each protection has to protect mainly one component, but must also operate as a backup to other protections nearby. Currently, in the Italian transmission network, three kinds of protections are used:

- *Differential protections* are used for bus-bars and detect the difference between input and output power. If this difference is greater than a fixed threshold the protection has to intervene, tripping all the breakers connected to the faulty bus as soon as possible. The intervention time is about 50 millisecond from the instant when the faulty condition has been detected.

- *Max-current protections* are used to protect transformers and generation groups from overloads: when other protections do not properly insulate a fault, they act as backup and trip the associated breakers if the current is greater than a fixed threshold. The intervention times range from 1.3 to 1.7 seconds.
- *Distance protections* are used for lines: they are in charge of estimating the direction and the distance of the fault with respect to the protection and, accordingly, of deciding if the protection has to intervene and how long it must wait before intervening. Each line is endowed with two distance protections associated with its terminal bays. Each distance protection continuously receives the measurements of both voltage and current relevant to the associated bay. The intervention time is tuned according to the following principle: the nearer the fault the faster the tripping command. More precisely distance protections are characterized by four intervention steps. When a fault condition is detected (i.e., the impedance seen by the protection is less than a fixed threshold Z_{start}) the protection starts and the relevant breaker is tripped after a delay which depends on the ratio between the nominal impedance of the line to be protected and the impedance actually seen by the protection and on the direction of the energy flow, as shown in the table reported in Figure 1.

Because of the tuning of the first step at the 85% of the nominal impedance of the line to be protected, if a fault occurs near the opposite part of the line with respect to a given protection, it is likely that this protection does not realize that the fault is in the line it has to protect. For this reason, the two protections of a line are connected by means of a telecommunication system: when a protection detects a first step condition, it sends also a telesignal to the other line protection. All the logical signals coming from the protections and the breakers are chronologically recorded by an event recorder and transmitted to a Regional Control Center (RCC) where they are used for faults localization. Records consist of a unique address of the event source, a time stamp, and an event code.

2.2 Network Fault Analysis

As explained above, when a fault occurs the protection system reacts in order to isolate the fault, so that the components involved in the fault are permanently disconnected from the network. In this case, the operators of the RCC have to decide as soon as possible (possibly within one minute) where the fault is located and what recovery actions have to be applied. It should be clear that fault localization is crucial to decide recovery actions and that possible malfunctions of the protection system hugely complicate this task. In fact, for a correct fault localization, it is necessary to evaluate the performance of the whole protection system, checking the behaviour of each protection and of each breaker. In other words, it is necessary to assess, for instance, if a particular protection correctly recognized the actual distance of the fault, if it tripped the relevant breaker within the expected time, if the breaker actually interrupted the current flow within the nominal time, etc. This task is currently performed by the operators of the RCC on the basis of the information coming from the event recorders of the different substations. Very often, the current status of the breakers only is considered, but this is not sufficient for a univocal fault localization. Other information is available to the RCC, such as the protection intervention step or the intervention time, which might help to formulate an unambiguous diagnosis, but this is currently not considered by operators, since they can not manage such an amount of data within the required response time.

2.3 Defining the Problem

The *network fault analysis problem* can be stated as follows:

Given:

- a set of state values of some components (for example, breakers and insulators open or closed, generators and transformers in service or not);
- the set of messages recorded by the event recorder since the occurrence of the fault till the network reaches a stable configuration (note that, whereas the presence of a message is sufficient to guarantee that a certain event happened, the absence of a message does not prove that this message had not been sent: due to transmission problems, some messages might be lost);

identify:

- the component (for example, a line or a bus bar) involved by the short circuit (*short circuit localization problem*);
- the list of faulty protections and breakers -if any, possibly with the indication of the diagnosed failure (such as: missed intervention, improper intervention, intervention delayed or in advance, lost message, etc.) (*protection behaviour evaluation problem*).

While the above two problems are distinct, they are interdependent and can not be solved separately. Short circuit localization is deduced from messages coming from the protections: an evaluation of their behaviour is requested in order to make correct deductions. In some cases messages coming from protections can even lead

Impedance actually seen (Z_s)	**Energy flow direction**	**Tripping time**
$Z_s < 0.85*Z_n$	Toward the line	50 ms
$Z_s < \max[1.2*Z_n, 0.85*(Z_n + 0.85*Z_{ns})]$	Toward the line	400 ms
$Z_s < 1.2*(Z_n + Z_{nl})$	Toward the line	900 ms
Z < Z [illegible]	Any	2 s

to contradictory deductions: in this case the evaluation is necessary to find out which of them have to be considered correct. Moreover, in order to establish what the correct operation of a protection should be, a hypothetical fault localization has to be assumed. Therefore, these two problems are strictly connected and a common solution has to be found.

2.4 An Approach to the Solution

First of all, let us note that a complete solution of the above described problems requires temporal reasoning capabilities. For instance, the evaluation of protection operation can not be reduced to verify if it intervened or not in presence of a certain fault: the intervention, when required, has to be performed not after (and in certain cases also not before) a certain time interval and a violation of these time constraints must be considered a failure.

Note also that the occurrence of a short circuit within the system and the consequent intervention of protections cause a transition of system configuration toward a new stable configuration, where the fault is no more present. In order to find out what fault has occurred, system history has to be reconstructed knowing the initial and final configurations and the, possibly uncomplete, set of messages that track the temporal evolution of the system.

A single protection intervention can not be associated with a specific fault: for instance, an intervention at the second step can be caused either by a fault in the terminal part of the protected line, or on the opposite bus bar or in the initial part of the next line. In each of these cases, a different failure of other parts of the protection system has to be hypothesized. While a single intervention can not be associated to a fault, the consideration of all interventions may lead to univocally identify the occurred fault. However, it is practically impossible to foresee all the combinations of messages which could be associated to a certain fault, because of the complexity of the system and of the huge variety of possible cases. Furthermore, even if all combinations could be considered, lost messages would complicate the problem so that this approach would be impractical. For these reasons a possible method for the solution of the network fault analysis problem is that of reconstructing step by step the system evolution, so that all messages received by RCC are explained within a single coherent history.

3. PLOTS: A NEW APPROACH FOR REPRESENTING KNOWLEDGE ABOUT SYSTEMS AND EVENTS

3.1 Basic Concepts

In our approach, a physical system is characterized by defining its structure and its behavior.

As far as structure is concerned, a physical system can be represented in our approach through a finite set D of *components.* Components are instances of *component types*; the set of component types is denoted by DT. We will distinguish in the following between *abstract components*, i.e. generic components characterized merely by their type, and *concrete components*, i.e. real components having individual properties such as a serial number, a physical location, and so on. Where not differently specified, the term components will be used to denote both abstract and concrete components. It is assumed that the set D includes also a virtual component, the *world*, which represents anything the system can interact with in the external world: users, observers, external measurement equipment, the nature, the fate, etc. Components may be connected to each other trough binary, directed *links.* The finite set of links existing between the components of a system is denoted by L. A link is a pair l = <d, d'>, where d and d' are the connected components (the link is directed from d to d').

As far as behavior is concerned, to each component type dt a finite set S_{dt} of possible *states* is associated. A state represents a possible operational mode of a component of a given type.
A *state value* is a 4-tuple v = <d, s, t', t">, where:

- d is a component, denoted as Comp(v);
- s is one from its possible states;
- t' and t" are two time instants, t' is denoted as Start(v) and t" is denoted as End(v), which represent the limits of a time interval during which d is in the state s, (no assumptions are made about the state of the component out of these limits: it is considered unknown).

A state value v is abstract or concrete, depending on the fact that Comp(v) is an abstract or concrete component. Two components, connected by a link, can interact exchanging symbolic, instantaneous *messages.* A message is a 4-tuple msg = <d', d", t, m>, where:

- d' is the *sender component,* denoted as Sender(msg);
- d" is the *receiver component*, denoted as Receiver(msg);
- m is the message content, taking its value from a predefined set MT of message types;
- t is the communication time instant (it is assumed that there is no propagation delay, so that sending and receiving instants are coincident), denoted as Instant(msg).

Figure 2: A sample plot.

A message msg is abstract or concrete, depending on the fact that Sender(msg) and Receiver(msg) are abstract or concrete components (note that both have to be either abstract or concrete). In order to represent time instants, we adopt the temporal model proposed in [2], where time instants are symbols interconnected by a set C of *temporal constraints*. Each temporal constraint is a 4-tuple c = <t', t'', a, b> and means that the time interval starting at time instant t' and ending at time instant t'' has a duration limited by the two real values a and b, i.e., $a \leq duration(t', t'') \leq b$. Given a set T of time instants and a set C of constraints on T, it is possible to map T in R (the set of real numbers) by means of a function τ, where $\tau(t)=x$ transforms a symbolic time instant t into a real value x that can be interpreted as an absolute temporal allocation of the time instant. The function τ is said *consistent* if, interpreting interval durations as differences between real numbers,

$$\forall (t',t'',a,b) \in C,\ a \leq \tau(t'')-\tau(t') \leq b.$$

A given set C of constraints is said *unfeasible*, if it is impossible to define any consistent function τ, *feasible* otherwise (note that when C is *feasible*, in general infinite consistent alternative functions τ exist). For a given set T of time instants and a set C of constraints on T, the set of all consistent functions τ is denoted by W. If, given two instants t_1 and t_2, $\forall\ \tau \in W$, $\tau(t_1) \leq \tau(t_2)$, it is said that t_1 Precedes t_2.

3.2 Representing Behavioral Laws through Plots

In order to represent behavioral laws of a system component, we introduce a novel representation technique, based on a semantic unit called *plot*.

A *plot* is a 4-tuple <I,C,M,V> where:

- I is a non empty set of time instants;
- C is a set of temporal constraints on I;
- M is a non empty set of abstract messages;
- V is a set of abstract state values;

and the following constraints are satisfied:

- $I = \{Start(v): v \in V\} \cup \{End(v): v \in V\} \cup \{Instant(m): m \in M\}$
- $\forall\ v \in V,\ \exists\ m \in M$: Comp(v) = Sender(m) or Comp(v) = Receiver(m)
- $\forall$ v',v'' $\in$ V, Comp(v') = Comp(v''), End(v') Precedes Start(v'') or End(v'') Precedes Start(v').

From an intuitive point of view, a plot is an abstract agglomerate of events involving one or more abstract components: it is intended to be a snapshot of a possible behavior of a part of a system. A plot specifies a fragment of a possible history of the considered system: it states some temporal constraints between certain messages and certain state transitions, which represent typical, partial system behaviors that can occur in a variety of temporal contexts. In this way, we can consider real events occurring within the system as temporal instances of plots, which, in turn, play the role of abstract history fragments, free of absolute temporal

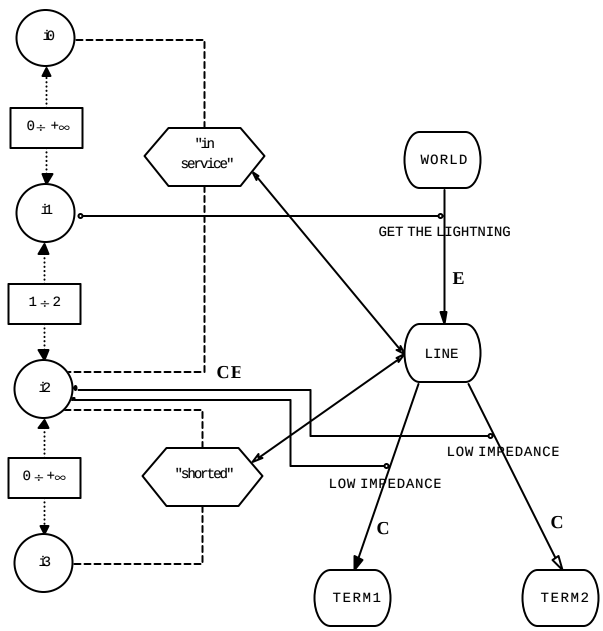

allocation.

For a better understanding, let us consider a very simple plot representing the fall of a lightning on a line in service. After a delay from 1 to 2 ms, the line becomes shorted and, as a consequence, low impedance is seen at line terminals. This is represented by a plot (see Figure 2) involving:

- four time instants i0, i1, i2, i3;
- three temporal constraints c0 = (i0, i1, 0, +∞), c1 = (i1, i2, 1, 2), c2 = (i2, i3, 0, +∞);
- three messages m1=(World,Line,"Get the lightning", i1), m2 = (Line, Terminal1, "Low impedance", i2), m3=(Line, Terminal2, "Low impedance", i2)
- two state values v1 = (Line, InService, i0, i2), v2=(Line, Shorted, i2, i3).

Plots are a knowledge representation technique typically oriented towards multiple use. Consider, for example, the multiple uses that the simple chunk of knowledge shown in Figure 2 can support:

- It can be used for simulation and prevision: if a lightning strikes a working line, you can deduce that the line will be shorted and low impedance will be seen from line terminals in a time interval from 1 to 2 ms.
- It can be used for past events and states reconstruction: if you see low impedance at line

terminals, you can deduce that, from 1 to 2 ms before, a lightning fell on the line and the line became shorted.

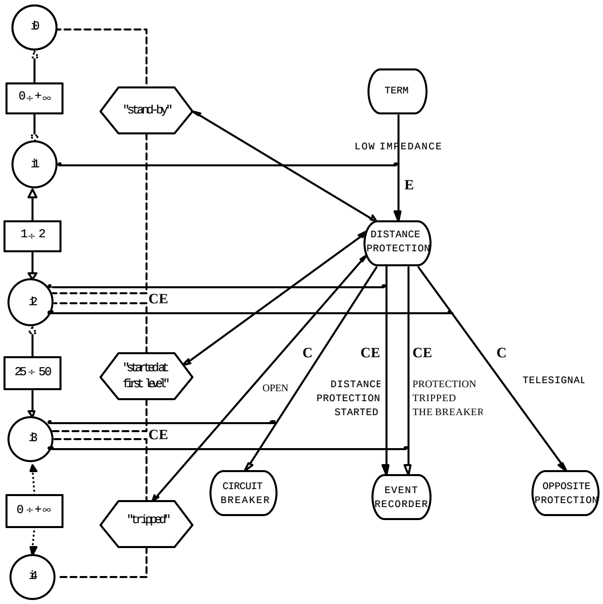

- It can be used for state evaluation: if you know that a lightning struck the line and you did not see low impedance within the foreseen time, you can conclude that the line was not in service.
- Moreover, if a notion of correct behavior is associated to plots, past event, state reconstruction, and state evaluation can be exploited for diagnostic purposes.

The plot formalism is not oriented towards any specific purpose: it intends to represent only a pack of messages and state values you can expect to happen according to a specified set of temporal relations. In other words, a plot is a collection of interrelated events you can expect to bump into, while observing the system.

Consider now a bit more complex example: the first step of a line protection. In this case the protection, being in state "stand-by", detects a low impedance at the line terminal and, within 1-2 ms, enters the state "started at first intervention level" and sends two messages:

- a message to the event recorder, signaling the occurred intervention;
- a message to the other protection of the same line, signaling the detected low impedance.

After a delay between 25 and 50 ms, the protection enters the state "tripped" and sends two other messages:

- a message to the relevant circuit breaker, ordering it to open;
- a message to the event recorder, signaling the opening command given.

The corresponding plot is shown in Figure 3. With respect to the previous example, additional ways of exploiting this chunk of knowledge can be considered. For instance, it is possible to look for lost messages: if the message "PROTECTION TRIPPED THE BREAKER" was received by the event recorder but the circuit breaker did not receive the command, it is possible to deduce that the relevant message was lost. Furthermore, it is possible to verify whether a state transition did actually occur or not. If it is known that low impedance was seen from the terminal, but no message was noticed, it is possible to deduce that something prevented the state transition (or, less probably, that all emitted messages were lost).

Through plots, it is possible to represent behavioral temporal knowledge about a part of a complex system in a form that is entirely independent from any reasoning mechanism that might be applied to it. A detailed analysis of the properties of plots as a general knowledge representation formalism and a comparison with other approaches to representing time and action (events) is however beyond the scope of this paper.

4 THE HISTORY RECONSTRUCTION PROBLEM

4.1 Basic Concepts

Fault analysis in power transmission networks can be considered as an instance of a more general problem, named the *history reconstruction problem*. Let us first define the concept of *history fragment*: an history fragment is a concrete message or a concrete state value. The *history reconstruction problem* can then be stated as follows:

Given:

- a physical system defined through the specification of its structure (namely, a set of components and a set of links among them) and behavior (namely, a set of plots);
- a set of initial history fragments for the system, i.e. a set of concrete messages and state values;

determine all the possible sets of history fragments that:

- are compatible with the structure and behavior of the system at hand;
- include all the initial history fragments provided in input;
- refer to a specified time interval $[t_1, t_2]$, including at least the temporal scope of all the initial history fragments.

History reconstruction does not imply a fixed temporal direction: the fragments initially available can be at the beginning of the chain, at the end, in the middle or arbitrarily distributed along it. This entails that either backward or forward temporal deductions can be carried out: the aim of the reconstruction can be

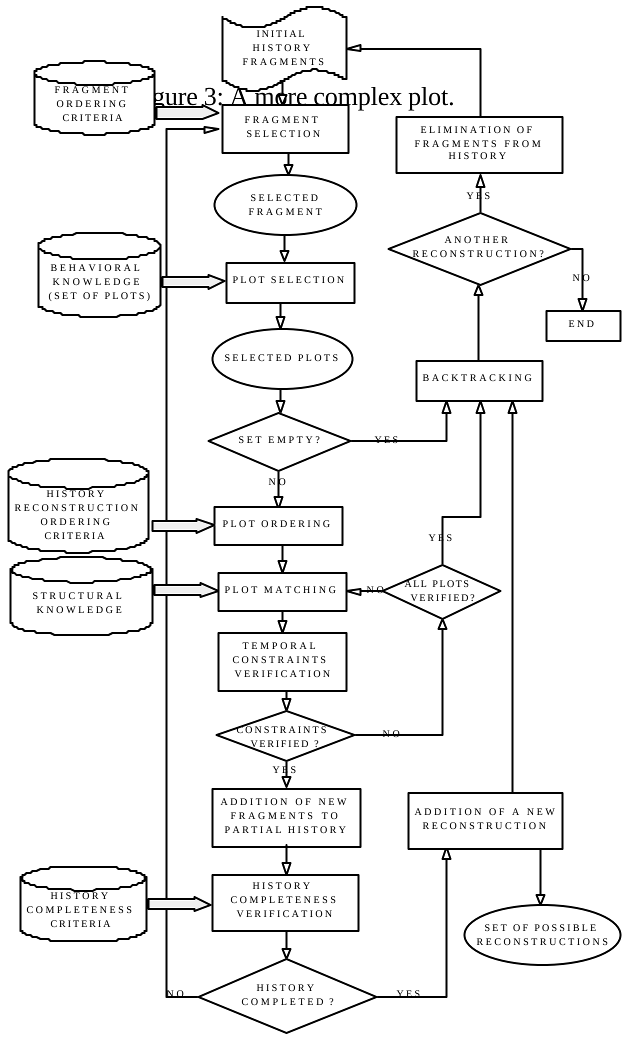

gure 3: A more complex plot.

either the detection of past events, or the prevision of future ones, or both.

In order to better define (and delimit) the history reconstruction problem, we make some additional assumptions:

- all the available fragments are reliable, though we are not sure to have all possible fragments; in other words, we admit lost data, but not false or counterfeit ones;
- it is impossible (or impractical) to make additional observations on the system before or during the solution of the history reconstruction problem; therefore, there is no possibility of increasing the initially available information.

As it is evident, fault analysis in a power transmission network is an instance of the history reconstruction problem. In this case, the initial history fragments are the messages received by the event recorder and the information available about the state of system components. In particular, it is possible to assume that:

- before the occurrence of a fault the state of (almost) all the components is known;
- after the occurrence of the fault, the system evolves toward a new stable state in which the fault has been isolated: messages received by the event recorder refer to this evolution;
- after a stable state has been reached again, the state of (almost) all the components is known.

4.2 Exploiting Plots for History Reconstruction

As described above, plot are abstract prototypes of fragments of real histories. The basic idea behind our approach is that of referring to these abstract prototypes in order to explain and, possibly, complete the fragments that have been collected through observation. An overview of the approach is given in Figure 4, where rectangular boxes represent the phases of the reconstruction process and the elliptical boxes represent intermediate data. Initially, one of the fragments is selected for plot matching, then a set of plots that can be matched with the selected fragment is identified: each selected plot represents a branching alternative in the reconstruction process. If the set is empty, a backtracking procedure is activated, otherwise the alternatives are ordered according to heuristic criteria (see below). Plot matching is then applied to the first plot in the ordering: if matching gives rise to unfeasible temporal constraints, the following plot is tried. If none of the plots satisfies temporal constraints, eventually backtracking procedure is activated, otherwise, new fragments are inserted in the history, according to the matched plot. The completeness of the reconstructed history is then checked. If the reconstructed history is complete, a new possible reconstruction is produced and backtracking is activated, otherwise the reconstruction process proceeds including the new fragments in the history which is being reconstructed and a new fragment for plot matching is selected. Backtracking procedure is aimed to check if it is possible to try to build an alternative reconstruction and, if the case, it is in charge of removing the history fragments that have been added during the previous reconstruction.

For the sake of clarity, let us introduce a simplified example of network fault diagnosis problem, before giving more details about each step of the reconstruction process. The example refers to a portion

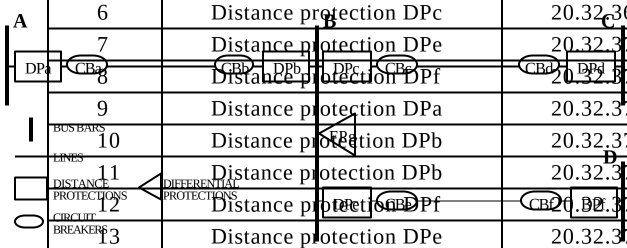

Figure 5: A fragment of power transmission network

Id	Sender Component	Time stamp	Message
1	Distance protection DPe	20.32.36.740	Distance protection started
2	Distance protection DPf	20.32.36.745	Distance protection started
3	Distance protection DPd	20.32.36.750	Distance protection started
4	Distance protection DPa	20.32.36.755	Distance protection started
5	Distance protection DPb	20.32.36.760	Distance protection started
6	Distance protection DPc	20.32.36.780	Received telesignal from opposite protection
7	Distance protection DPe	20.32.37.145	Distance protection switched to 2nd step
8	Distance protection DPf	20.32.37.145	Distance protection switched to 2nd step
9	Distance protection DPa	20.32.37.155	Distance protection switched to 2nd step
10	Distance protection DPb	20.32.37.160	Distance protection switched to 2nd step
11	Distance protection DPb	20.32.37.665	Distance protection switched to 3rd step
12	Distance protection DPf	20.32.37.695	Distance protection switched to 3rd step
13	Distance protection DPe	20.32.37.700	Distance protection switched to 3rd step
14	Distance protection DPa	20.32.37.705	Distance protection switched to 3rd step
15	Distance protection DPf	20.32.37.730	Distance protection tripped the breaker
16	Distance protection DPa	20.32.37.740	Distance protection tripped the breaker
17	Distance protection DPd	20.32.37.780	Distance protection tripped the breaker
18	Circuit breaker CBf	20.32.37.785	Breaker opened
19	Circuit breaker CBa	20.32.37.800	Breaker opened
20	Circuit breaker CBd	20.32.37.820	Breaker opened

of the system composed by three lines converging on the bus bar of a substation (see Figure 5). Let us suppose that, before the fault happens, it is known that all lines are in service, all circuit breakers are closed, and all protection are in stand-by. Then, a fault occur and, after the intervention of the protection system, the breakers CBa, CBd and CBf are open and all the lines and the bus bar are out of service. Using only this information, four diagnostic hypotheses can be done:

- a fault in the bus bar B with faulty behavior of the relevant differential protection FPg;
- a fault in line AB with faulty behavior of the protection DPb or of the breaker CBb;
- a fault in line BC with faulty behavior of the protection DPc or of the breaker CBc (actually the correct one);
- a fault in line BD with faulty behavior of the protection DPe or of the breaker CBe.

In the meantime, the event recorder has received some messages that can be profitably exploited to solve the diagnostic problem through the reconstruction process, as it will be explained below. Received messages are shown, in chronological order, in the table of Figure 6.

Fragment Selection

History reconstruction starts by selecting one of the available history fragments (a message or a state value).

Domain dependent heuristics can profitably direct this selection in order to improve the efficiency of the reconstruction process. For instance, since, in the domain at hand, a message gives a more specific account of what happened than a state value, messages are preferred to state values in this selection. Moreover, among the received messages, one has to be selected for starting reconstruction: in this case, a heuristic which proved useful is that of choosing the messages which appear less frequently among the fragments. Therefore message 7 is selected.

Another heuristics, which has been profitably applied, states that if a fragment has been selected in the previous step, but it has been only partially explained (see below for the definition of the concept of explanation), it has to be selected again in order to reach a complete explanation.

Plot Selection and Ordering

In general, many plots may be available that match with a selected history fragment. Therefore, it is necessary first to identify all plots that match the selected history fragment and then to establish the order in which plots have to be effectively exploited for history reconstruction. The identification of all plots that match a selected fragment, (in our example the message "Received telesignal from opposite protection"), is a simple task that can be performed efficiently through an indexing mechanism on the plot database. The ordering of the plots strongly influences the efficiency of the history reconstruction process, and should be driven by both general and domain-dependent heuristics. A general heuristic we adopted for the ordering is that precedence is given to plots that require a lesser number of new fragments to be added to the partially reconstructed history (see below). This heuristic corresponds to the intuitive idea that "compact" reconstructions, requiring few additional hypotheses, should be preferred and explored first.

Plot Matching

Let us suppose now that the plot shown in figure 7, has been selected first for matching with message 7. The concrete message 7 is then matched with the abstract message "TELESIGNAL RECEIVED" within the plot. Since sender and receiver of message 7 are known (the distance protection DPc and the event recorder of the considered region), they are matched with the abstract components "DISTANCE PROTECTION" and "EVENT RECORDER" specified in the plot.

Then other abstract components present in the plot can be matched with concrete components, according to knowledge about the structure of the system, defined by links between components: OPPOSITE PROTECTION is matched with DPd and CIRCUIT BREAKER is matched with CBc. Matching procedure goes on considering messages. One of the abstract messages (TELESIGNAL RECEIVED) has been already matched with message 7; for each other abstract message, the plot matching procedure looks among the available history fragments in order to find any real message that can be matched with it, i.e. a message that has the same content, the same sender, the same receiver and which respects the temporal constraints specified in the plot about the sending instant.

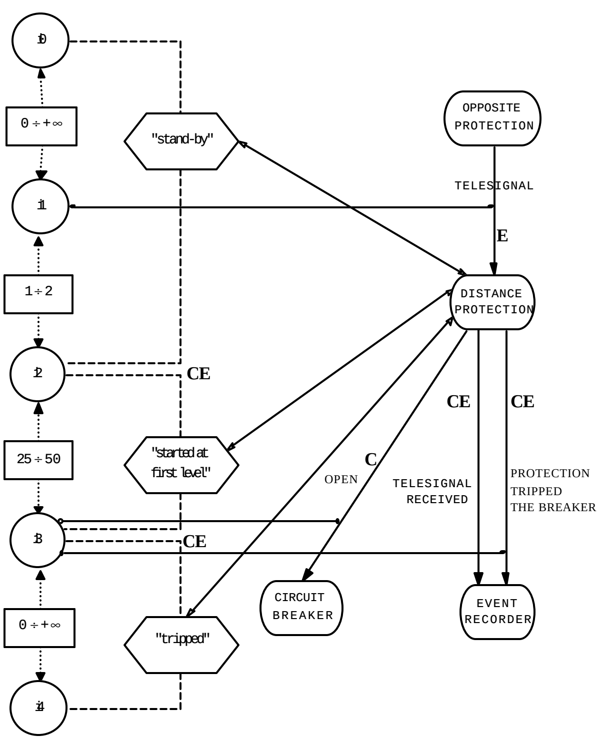

Figure 7: A plot representing the reaction of a distance protection to a telesignal.

If no such message is found, a new message, with the specified characteristics is added to current history fragments, since it is necessary in order to carry on reconstruction. In the considered example, three new messages are added to current history fragments, along with the corresponding time instants and the temporal constraints between them. Finally, state values, along with the relevant time instants and constraints specified in the plot, are matched with state values already present or are added to the history.

Temporal Constraints Verification

When history fragments are matched with a plot, it is necessary to verify that constraints already present in the history together with those entailed by the plot are feasible. The same check has to be performed when new time instants and new temporal constraints are added. Furthermore, the constraint that a component can not be in two different states at the same time must be respected too. Constraint feasibility is checked through the constraint propagation technique proposed in [3].

Completeness Verification

Let us suppose that constraints are satisfied and the partial history has therefore been added with new fragments: then it has to be decided if the reconstruction carried out so far is satisfactorily complete or if it must be continued. In order to make this decision, a completeness criterion has to be defined, which is, in general task- and domain-dependent. For instance if history reconstruction is applied to a continuous process, either for forecast or for diagnosis, an history could be considered satisfactorily complete when it spans over a sufficiently large time interval.

In our diagnostic application, an history is considered complete when all fragments it includes have been *explained*. In order to define the concept of explanation, we introduce a (rough) concept of causality: we state that all events in the history, i.e. state transitions and message emissions, must be causally explained, i.e. a cause for the event must be identified. Moreover events must also be eventually explained, i.e. the consequences they produce have to be explored.

Knowledge about causal relations between events might be deduced from plot structure according to a general scheme of causal relation. A detailed discussion about this point is however outside the scope of this paper and is not strictly necessary for the comprehension of the following concepts.

Therefore, let us simply assume that causal knowledge is included in the definition of plots by associating to each event in the plot an explanation label L ∈ {C, E, CE}. The labels C, E, CE, mean respectively that the event is causally explained, eventually explained, or both, within the plot. For example, turning back to figure 3, the message LOW IMPEDANCE is eventually explained, in the sense that its effects (not its causes) are included in the plot, messages TELESIGNAL and OPEN are causally explained since their cause is included in the plot but their effects have to be explored, messages PROTECTION STARTED and BREAKER TRIPPED are both causally and eventually explained since their causes are included in the plot and they have not further effects. Finally the two state transitions included in the plot are both causally and eventually explained.

When plot matching is applied, explanation labels are transfered to real history fragments so that real messages and state transitions in the history reconstruction become causally explained, eventually explained, or both.

History reconstruction can therefore be considered completed when all messages and state transitions (either observed or introduced by plot matching) are both causally and eventually explained, i.e. what has been reconstructed is self-explaining. Initial causes in any history, are messages coming from the external world, which do not need further causal explanation, such as lightning strokes or component faults.

When history reconstruction is completed control is passed to the backtracking procedure, which is in charge of deciding if an alternative reconstruction has to be attempted or the process has to be definitively terminated.

Backtracking Procedure

In the history reconstruction process, the backtracking procedure is invoked when a failure occurs, i.e. when no plots are found that can match with the fragments of a partially reconstructed history, or when a reconstruction is completed, i.e. when all fragments are completely explained. In both these cases a choice has to be made between attempting an alternative reconstruction and terminating definitively the process. Let us consider first the failure case: for the sake of clarity we will still refer to our concrete example.

After first plot matching has been applied to message 7, the messages TELESIGNAL RECEIVED and PROTECTION TRIPPED THE BREAKER are completely (CE) explained, whereas the message OPEN included in the plot is just partially (C) explained. According to heuristics described above, the message OPEN is selected for plot matching in order to complete its explanation. Let us suppose that a plot representing the correct operation of the breaker is selected: among other facts, a transition of breaker DPc to the state OPENED is added to the history. However, it is known that the final state of breaker DPc is CLOSED: therefore an explanation for this state transition has to be found. Since there is no plot including such a transition, a failure is generated and the backtracking procedure is invoked.

The backtracking procedure is in charge of finding an alternative reconstruction, if any one exists, for the initial fragments. In order to do this, some of the choices made in the failed reconstruction have to be revised. In particular, since, at any step, a plot has been selected for matching and additional fragments have been added to the history, one of these previous steps has to be retracted and the relevant additional fragments eliminated from the history. After that, reconstruction can continue exploring the choice of a different plot. In the current version of SHORT, chronological backtracking is applied, i.e. the last plot matched is retracted first.

Referring to the example, the plot representing the correct operation of the breaker is retracted and another plot is applied, where the breaker is in a faulty state and does not open on reception of the opening command.

The reconstruction can then be completed in a consistent way: reconstruction includes a short circuit in line BC, a fault of breaker CBc, and the loss of 1 message addressed to event recorder.

When this reconstruction is completed, the backtracking procedure is invoked in order to search for other reconstructions. Referring to the example, a different plot can be matched with message 7: in this plot the protection is in a faulty state and, on reception of telesignal, does not activate the tripping mechanism. Also in this case a consistent reconstruction can be completed, which includes a short circuit in line BC and a fault of protection DPc.

Different reconstructions can be evaluated in order to propose them to the user as diagnostic hypotheses in a preference order. Parsimony [6] is the (rough) preference criterion adopted in the current version of SHORT: diagnoses are preferred that include the lesser number of anomalous events (message losses and component faults). According to this criterion, the latter diagnosis, involving a protection fault, is preferred to the former diagnosis, involving a circuit breaker fault and one message lost.

The reconstruction process terminates when all alternative reconstructions have been explored: in our example, no other reconstructions are possible.

5. IMPLEMENTATION AND EXPERIMENTAL RESULTS

The plot matching approach for the solution of the history reconstruction problem illustrated in the previous sections has been implemented in SHORT, a knowledge-based system for fault diagnosis in power transmission networks. SHORT has been developed in CLOS (Common Lisp Object System) on an IBM RISC-6000 workstation. SHORT has been extensively tested on simulated and real cases. The regional power transmission network of Lombardia (Italy) has been used as a testbed for the SHORT prototype: it encompasses 12 primary stations and 9 secondary stations, for a total of about 500 components of 14 different types.

The diagnostic approach adopted has proved to be sound: a correct diagnosis has been identified in all cases considered and, thanks to the heuristics adopted to direct the reconstruction process, has been proposed as the preferred diagnosis in all real test cases. In simulated cases, where initial information about component states has been reduced and the number of lost messages has been increased with respect to real cases, the correct diagnosis has been in any case among the first two preferred diagnoses.

The analysis of the performance of SHORT on the testbed has provided the following results:

- The time required to produce the first diagnosis goes from 10 seconds (in simplest real cases) to 3-4 minutes (in the worst simulated cases). It heavily depends on the complexity of the fault case considered, but is only little influenced by the completeness of the initially available information.
- The time necessary to produce all possible alternative diagnoses and terminate the search, goes from 10 minutes to several hours. Due to the intrinsic combinatorial nature of the problem, it is largely influenced by the completeness of the initially available information.

6. CONCLUSIONS

The contribution of this work is twofold:

- from the theoretical point of view, a general class of temporal reasoning problems, called history reconstruction, has been identified, and a method for its solution has been proposed, exploiting the properties of a suitably defined knowledge representation formalism, called plots;
- from the application point of view, an important real problem, fault diagnosis in power transmission networks has been faced and solved as an instance of the history reconstruction problem and a knowledge-based system, called SHORT, for fault diagnosis in power transmission networks has been successfully implemented and tested.

SHORT is the result of a research work still in progress. The main directions of future activity will include the following issues:

- the study of an original theoretical approach to diagnosis of time-varying system, overcoming the limitations of classical theory of diagnosis [7], which considers the behavior of the system to be diagnosed characterised by static, persistent states and the structure of the system itself as time-invariant;
- the study of intelligent backtracking methods, in order to improve the efficiency of the reconstruction algorithm;
- the association of a probability value to each (partial) reconstruction, in order to allow the system to generate the most probable reconstructions first (search focusing) and, in the cases the completeness of the search is not requested, to stop the generation of reconstructions which are below a fixed probability bound (pruning).

Acknowledgements. The research reported in this paper has been partially supported by CISE s.p.a. (Milano, Italy), who has also provided the testbed concerning the regional power transmission network of Lombardia (Italy).

REFERENCES

[1] H.-J. Cho , J.-K. Park , and H.-J Lee, A fuzzy expert system for fault diagnosis of power systems, *Proceedings ISAP 94*, Montpellier, F, 1994, 217-226

[2] T.L. Dean and D.V. Mc Dermott, Temporal data base management, *Artificial Intelligence* 32, 1987, 1-55.

[3] R. Dechter, I. Meiri, and J. Pearl, Temporal constraint networks, *Artificial Intelligence* 49,1991,61-95.

[4] International Conference on Intelligent System Application to Power Systems, Montpellier, F, 1994.

[5] M. Kezunovic, I. Rikalo, D.J. Sobajic, Neural network applications to real-time and off-line fault analysis, *Proceedings ISAP 94*, Montpellier, F, 1994, 29-36

[6] Y. Peng and J.A. Reggia, Plausibility of diagnostic hypotheses: the nature of simplicity, *Proceedings AAAI-86*, *6th National Conference on Artificial Intelligence*, Philadelphia, PA, 1986, 140-145

[7] R. Reiter, A theory of diagnosis from first principles, *Artificial Intelligence* 32(1), 1987, 57-95.

[8] Y.Sekine, Y.Akimoto, M.Kunugi, C.Fukui, and S.Fukui, Fault diagnosis of power systems, *Proc. of the IEEE, Special Issue on Knowledge-Based Systems in Electric Power Systems*, Vol. 80, No. 5, 673-683, 1992

[9] G. Tornielli, L. Capetta, S. Cermignani , A.S. Fabiano, and R. Schinco, An interval-based approach to fault diagnosis of power systems, *Proceedings ISAP 94*, Montpellier, F, 1994,809-816

[10] J.P. Wang and J. Trecat, A parallel fault diagnosis expert system for one dispatch center, *Proceedings ISAP 94*, .Montpellier, F, 1994, 201-208