Supply Restoration in Power Distribution Systems: a Reference Problem in Diagnosis and Reconfiguration

Marie-Odile Cordier Sylvie Thiébaux
IRISA
Campus de Beaulieu
35042 Rennes Cedex, France
<Firstname>.<Lastname>@irisa.fr

Olivier Jehl Jean-Paul Krivine
EDF-DER
1, Av. du Général De Gaulle
92140 Clamart, France
<Firstname>.<Lastname>@edfgdf.fr

Abstract

This paper proposes to use the problem of supply restoration in power distribution systems as a reference in the area of model-based diagnosis and reconfiguration. This real-world problem is at once simple to understand, easily scalable, and complex to solve. As a representative of a wide range of economically important contexts, it is an ideal testbed for systems claiming to address the necessities of the real world. As a representative of up-to-date research in diagnosis and also in planning, it offers a challenge to both communities.

Motivation for the problem

In real industrial environments, diagnosis is often coupled with repair or reconfiguration in a close loop fashion: observe, diagnose, act (i.e., repair/reconfigure), observe again, revise diagnoses in accordance, and so on. A rationale is that time-consuming observations and prohibitive breakdown costs force us to act urgently, based on early diagnoses. Another reason is that by changing the state of the system, actions bring useful information wrt. discrimination between diagnoses. This information might not be available by other means, or else at much higher cost.

A typical case is supply restoration in a faulty power distribution system. This problem, which is a major concern for electricity distributors such as Electricité de France (EDF), consists in localizing the faulty lines on the distribution network and reconfiguring this latter so as to resupply most of the customers left without power as an effect of the fault. Naturally enough, this has to be done in real-time. Available sensors and actuators are not always reliable, which leads to missing and erroneous information about the state of the components and about the configuration of the system.

In virtue of this accumulation of realistic features, the problem is an ideal testbed for systems claiming to address the necessities of the real-world. Moreover, power distribution systems are representative of a range of other economically important contexts, such as those of water distribution, communication, and motorway networks.

An advantage of this problem compared to as realistic others is that it is relatively simple to understand. Although the structure of the system changes following the execution of a reconfiguration action by the operator, the model of the system does not evolve with time. Only a few classes of components are involved whose functionalities are straightforward. There exist essentially one type of fault and a few behavior modes for sensors and actuators. Further, the topology of power distribution systems makes it easy to scale the problem up or down in order to assess the efficiency of algorithms.

Finally, despite this simplicity, the size of real distribution systems makes them very challenging, not only for tools developed in the diagnosis community, but also for methods existing in related areas such as planning for partially observable environments.

The paper is organized as follows. The first two sections describe the physical aspects of power distribution systems and the problem of supply restoration, respectively. The third section details the features that makes this problem a challenging and representative testbed. The fourth section gives an overview of the scope of the problem wrt. existing work in the diagnosis and planning literature. The last section sketches our own work on this problem. In appendix, we deal with the availability of data and software related to this problem, and give an example of a supply restoration scenario.

Description of the physical system

Topology

A power distribution system, as in Figure 1, can be viewed as a network of electric lines connected via switching devices (SDs – represented by small squares in the figure) and fed via circuit-breakers (CBs – represented by large squares). Switching devices and circuit-breakers have two possible positions: either open or closed (white devices in the figure are open and the others are closed). A circuit-breaker supplies power iff it is closed, and a switching device stops the power propagation iff it is open. Distribution networks have a meshable topology and are exploited radially: the positions of the devices are set so that the paths taken by the power of each circuit-breaker form a tree called feeder. The root of a feeder is a circuit-breaker, its leaves are open switching devices, and each line belongs to a single feeder (in the figure, adjacent feeders are distinguished using different gray levels, and the indexes of the devices located on the same feeder have the same first digit). Consumers may be located on any line, and are then only supplied when this line is fed.

Faults

There is essentially one type of fault of interest here, namely the permanent short-circuit on a line. Since these short circuits are mainly due to bad weather conditions and lightening, multiple faults are not rare. Upon occurrence of a fault, the circuit-breaker feeding the faulty line opens in order to protect the rest of its feeder from damaging overloads. As a result, all customers located on that feeder are left without power. Using the sensors and actuators described below, a few minutes are then available to locate the faulty line and to reconfigure the network so as to isolate the fault and restore the supply to the non-faulty lines.

Sensors and actuators

As shown in figure 2, devices are equipped with a remote-controlled actuator (AC) used to change their position, a position detector (PD) sensing this position, and a fault detector (FD) sensing the presence of faults.

Fault detectors are the basis on which to localize faults.[1] They permanently indicate whether or not they have "seen a fault pass them", i.e., whether or not a fault is downstream on the feeder. Then ideally, a fault is located on the line between a sequence of switching devices whose fault detectors indicate that it is downstream and a sequence of switching devices whose fault detectors indicate that it is not. Note that in the unfortunate case of multiple faults on the same feeder, only the most downstream faults can be detected that way. More importantly, fault detectors are not always correct and can be in one of the following

[1]Fault localisation may differ from one utility to another, and be based, for instance, on fault distance calculation or on bidirectional fault detectors. Our description refers to EDF distribution systems.

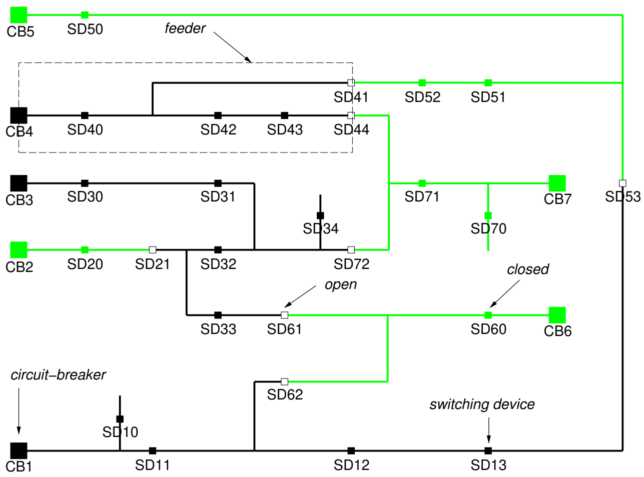

Figure 1: Power Distribution System (for semi-rural areas – representative data supplied by EDF)

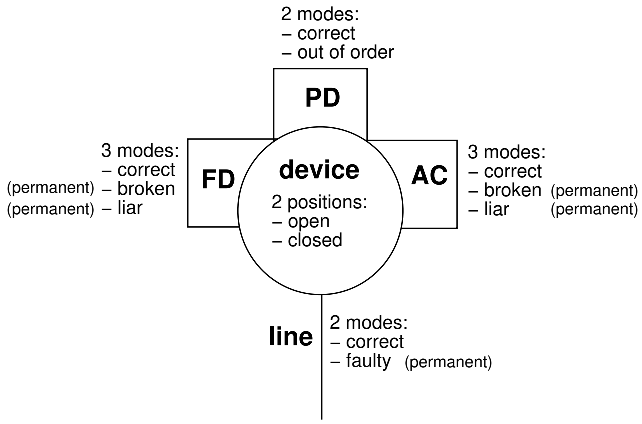

Figure 2: Behavior modes

two permanent abnormal modes: broken (i.e., they do not return any information), or even liar (i.e., they return erroneous information).

Actuators enable us to switch devices so as to isolate suspected lines and to direct the power from other feeders towards the non-faulty lines. Actuators are not always reliable and can be in one of the following permanent abnormal modes: broken (the actuator fails in executing the switching operation and sends a negative notification), or liar (it fails in executing the operation but sends a positive notification).

Position detectors can be consulted in order to reduce uncertainty about the success of switching operations positively notified by actuators. However, position detectors can be out of order (they do not return any information) for an undeterminate time, in which case the configuration of the network remains uncertain.

Size

As many European power distribution systems, EDF networks are composed of some hundreds of feeders (typically from 100 to 300), each of which contains a few remote-controlled switching devices (the objective is to equip each feeder with 2-3 of them). However, a feeder has only a very few neighbors (typically from 1 to 4), and essentially only those will play a role in supply restoration. Hence reasoning is very local, and the network in figure 1 is a good representative of the complexity of the real problem. The problem can trivially be scaled up or down by modifying the number of switching devices per feeder and the number of neighbors of feeders in this basic network.

The problem

Supply restoration

The problem of supply restoration is that of reconfiguring the network in view of resupplying the customers following the loss of one or more feeders. It amounts to building a restoration plan consisting of switching (opening/closing) operations. This plan must enable the isolation of the faulty lines by prescribing to open the switching devices surrounding them, as well as the restoration of the supply to the non-faulty areas of the lost feeders by prescribing to operate devices so as to direct the power towards these areas.

The following constraint determines which restoration plans are admissible: circuit-breakers and lines can only support a certain maximal power. This might prevent directing the power through certain paths and resupplying all the non-faulty areas. The goal of the restoration is to optimize certain parameters under this constraint, such as minimizing breakdown costs (i.e., resupplying as many consumers as possible, as fast as possible (within a few minutes in any case), giving priority to critical consumers like hospitals), minimizing the number of switching operations so as to stay close to the configuration in which the network is normally exploited (called the normal configuration), and balancing power margins of circuit-breakers in anticipation of the next load peak.

Fault localization

Being able to isolate faulty lines requires diagnosis, i.e., the localization of these lines. This is done by looking at the information transmitted by the fault detectors. It follows from possible fault detector failures that these might disagree, and more generally that several hypotheses of fault location exist, each of which corresponds to an hypothesis concerning the behavior mode of the fault detectors.

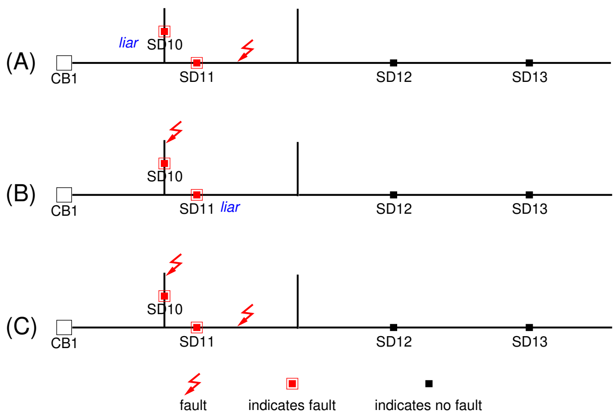

Figure 3: Behavior of fault detectors

There exist preferences between these hypotheses (the probability of multiple faults is much smaller than that of a fault detector lying, and this latter is higher when the fault detector indicates a fault downstream than when it does not because fault detectors do not detect all types of faults). But in fact, only the reconfiguration phase may enable us to discriminate, especially when this one goes wrong. This and the fact that we are trying to minimize breakdown costs, explain why reconfiguration and diagnosis need to be closely intertwined, as illustrated in the example below.

Example

Taking our example network, suppose that CB 1 opens, leaving the bottom-most feeder on the figure unsupplied. Suppose further that the fault detectors of SDs 10 and 11 indicate a fault downstream, while the other fault detectors do not. Among the most probable diagnoses are those shown in Figure 3: (A) there is only one fault, downstream of SD 11, and the fault detector of SD 10 lies, (B) there is only one fault, downstream of SD 10, and the fault detector of SD 11 lies, and (C) there are two faults, downstream of SD 10 and downstream of SD 11, and none of the fault detectors lie.

Assuming that we consider (A) to be the most likely, a possible restoration plan goes as follows: isolate the line between SDs 11 and 12 by opening these devices, resupply the lines upstream of the fault by reclosing CB 1, and have CB 5 resupply the downstream lines by closing SD 53.

We let the reader elaborate on the case where CB 1 reopens when we attempt to close it. Either this is due to a wrong diagnosis (e.g., (B) or (C) was the case), or to a fault which could not be detected (here, a fault between CB 1 and SD 11), or to a failure of an operation meant to isolate the fault (the actuator of SD 11 was liar, and this could not be detected if the position detector of this device was out of order). Since faults can in principle occur at anytime, it could even be the case that a new fault has appeared while we were attempting to restore the supply. However, for simplicity, we will assume throughout this paper that no new fault can occur during power restoration.

Characteristics of the problem

Two main characteristics of supply restoration in power distribution systems make it a challenging and representative reference problem.

Uncertainty

Firstly, it features numerous sources of uncertainty due to *missing* and *erroneous observations*:

- As usual in diagnosis, information gained from existing sensors (fault and position detectors) is not sufficient to get a precise idea of the state of the network and consequently to decide unambiguously where the faults are.
- As usual in real applications, sensors are unreliable and may stop providing any information or provide erroneous information.
- Furthermore, actuators are unreliable, which implies that the actual configuration of the network is not precisely known.

Consequently, the state space and the set of candidate diagnoses to be dealt with are very challenging. For instance, the network in Figure 1 has about 10^{56} possible states.

Intertwining diagnosis & reconfiguration

Secondly, diagnosis and reconfiguration are deeply intertwined:

- An accurate diagnosis requires reconfiguration. Indeed it is impossible to discriminate between candidate diagnoses by asking for new measurements, because these are unavailable. The only possible solution is to perform reconfiguration actions (which at the same time do modify the structure of the network) and to confront the resulting sensing information with the expected one.
- A good reconfiguration requires an accurate diagnosis. Indeed actions are selected according to their expected utility, and thus each candidate diagnosis, even highly unlikely, should ideally be taken into account for this evaluation. Pruning candidates is problematic since it may spoil the evaluation by forgetting unlikely but risky (costly) states.[2] However, the candidate set is too large to be integrally considered, without further discrimination.

Consequently, the challenge here is that we cannot be sure to choose an adequate reconfiguration action before discrimination, but also cannot be sure to discriminate adequately before starting reconfiguration.

Related work

As we just explained, the application we describe is very representative of the issues we have to address when dealing with real-world applications. It is also representative of up-to-date research challenges in the diagnosis and planning communities. As a matter of fact, these communities have examined related problems, and proposed solutions which are worth mentioning here even if they have some limitations when faced to this application.

Integrating diagnosis & repair

In the model-based diagnosis community, systems have been built that interleave diagnosis and repair so as to minimize breakdown, observations, and repair costs.

These systems often rely on a two-level architecture. At the lower level, a *classical planner* is responsible for returning an action sequence achieving a given repair goal for a given candidate diagnosis. At the top level, a *diagnostic reasoner* maintains a set of candidates (or a probability distribution on candidates) and chooses, at each step, whether it is preferable to discriminate between candidates by taking a new measurement or to undertake some repair activities. The decision results from a look-ahead search based on the plans returned by the planner and on a utility function taking into account breakdown, observations, and repair costs. Upon executing an action and obtaining new observations, the diagnostic reasoner updates the candidate set.

(Friedrich and Nejdl 1992; Sun and Weld 1992) are representative of this approach and generalize (de Kleer and Williams 1987) dedicated to the choice of the best next measurement. An important limit of these works is their dependence on the assumption that every relevant observation can be made reliably when needed, and on very basic reliable repair actions. Furthermore, the examples used for their illustration are restricted: the sets of candidates and available actions are small enough to be entirely considered, which is impossible in the case of power distribution systems.

[2] In our supply restoration example, a fault between SDs 12 and 13 is very unlikely given the observations, but is very costly wrt. applying the proposed restoration plan because we will loose CB 5 when we close SD 53.

Integrating diagnosis & reconfiguration

In the context of dynamic systems or in contexts where uncertainty is present, approaches integrating model-based diagnosis and reconfiguration have adopted architectures similar to their diagnosis-and-repair homologues. The sense-response loop is composed of two main phases. A diagnosis phase identifies the most probable states from knowledge of the previous state, observations, and events that have occurred or actions which have been performed. Based on this identification, a reconfiguration phase determines the lowest cost actions achieving a set of reconfiguration goals.

Examples of such approaches includes our own work on diagnosis and reconfiguration of power distribution systems (Thiébaux *et al.* 1996b) which we sketch below, as well as the approach to reactive self-configuring autonomous systems implemented in Livingstone (Williams and Nayak 1996). The main difference between these two works is that Livingstone is a general purpose kernel which uses concurrent transition systems to represent the dynamics of the environment (normal behavior, and following failures and repairs), while our prototype is domain-specific and takes advantage of the structure of power distribution systems in order to remain efficient in spite of uncertainty.

Planning and POMDPS

In the terminology used in the planning community, a power distribution system is a partially observable (stochastic) environment in which we must act optimally. Coping with such environments is a topic that has recently emerged as highly motivating in the planning literature. A well-known approach consists in modeling the problem as a partially observable Markov decision process (POMDP) and solving the decision process using algorithms inspired from traditional operation research methods (Cassandra *et al.* 1994).

The algorithms for solving POMDPs that have been experimented so far are based on extensional representations of the system. They potentially explore the whole belief state space, which clearly makes them too complex for power distribution systems. However, algorithms based on more compact representations are starting to be developed (Boutilier and Poole 1996), and one can hope that they will lead to significant improvements in the field. We believe that power distribution systems are an excellent testbed for the assessment of the benefits of these algorithms.

Power systems

Other works of interest here are those concerned with similar applications in power systems. Power systems being economically important, many such applications exist (see e.g., (ISAP 1996)). We only mention a few of them here.

(Friedrich and Nejdl 1992; Beschta *et al.* 1993) study diagnosis and supply restoration in power *transmission* systems. A crucial difference with our proposed testbed is that observations and actions are assumed to be reliable. This may be reasonable when considering transmission systems, but far too restrictive for power distribution systems.

(Moinard 1994) presents the DIAGRAL system, which tackles exactly the problem we describe by adopting an architecture integrating multiple models representing different types of knowledge. This system has been coupled with a power distribution system simulator and has been experimented in a training environment.

(Gonzalez *et al.* 1996) proposes a model-based approach to diagnosing faults in power distribution systems, identifying sensor failures and carrying out appropriate corrective action. An Intelligent Power Controller has been developed whose capabilities have been tested on a power distribution system breadboard representative of space station power systems. The diagnostic reasoning is a variation of Reiter's conflict-based approach. A brief mention is made of the fault recovery task, which includes fault isolation, and reconfiguration of the network as in (Crow and Rushby 1991) (in particular, the way diagnosis and recovery phases are interleaved is not specified).

Our current work and results

At present, human operators are in charge of supply restoration in EDF power distribution systems. They rely on their expertise and on pre-established restoration plans, which are both specific to the normal configuration of the network. For some years, EDF has put an important work into studying the automatization of supply restoration, in order to improve its speed, the quality of the restoration plans, and the treatment of large extent incidents (e.g., multiple faults) which force moving away from the normal configuration.

One the major steps of this study was the development of the AUSTRAL prototype (Bredillet *et al.* 1994). This prototype integrates a special-purpose model-based reasoner (for localization of the fault based on initial discrepancies, update of the state of the network upon occurrence of an action, and verification of the admissibility of plans), and an expert system (for plan selection-revision and further hypothetic reasoning). A network simulator is at the heart of the model-based reasoner, and returns the state resulting from performing a given switching operation on

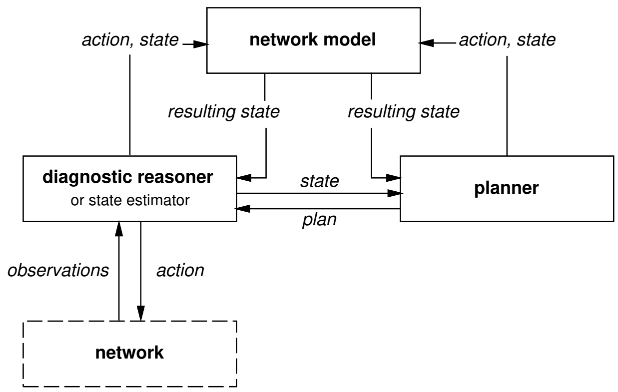

Figure 4: Architecture of SyDRe

a given state of the network. One of the limits of the AUSTRAL prototype is that the hypothetic reasoning performed by the expert system is not systematic: a failure of the current plan does not always lead to a revision of the state hypothesis and may result in an abortion of the restoration. Also, the AUSTRAL prototype is limited to a single fault, and plan evaluation does not account for the consequences of possible failures.

In the framework of a collaboration between IRISA and EDF, another step investigated ways of overcoming some of these limits (Thiébaux *et al.* 1996b). In particular, we built SyDRe, an entirely model-based prototype which handles hypotheses about the state of the network in a systematic way, including multiple fault hypotheses. As shown in Figure 4, SyDRe's architecture relies on a domain-specific model, which, close to the network simulator found in AUSTRAL, accounts for both the logical and quantitative aspects of power distribution systems.

A diagnostic reasoner is in charge of maintaining hypotheses about the state of the system throughout the sense-response loop. Its main steps are the following: (1) compute the possible states and their initial probability distribution from the initial observations, (2) ask the planner for a restoration plan for the most probable state, (3) if all actions in the current plan have been executed then exit the restoration process successfully, (4) otherwise execute an action in this plan, gather the resulting sensing information and update the probability distribution accordingly, (5) decide whether the current plan is still appropriate for the new most probable state of the distribution; if so go on with (3) otherwise go on with (2). In order to cope with complexity, the initial probability distribution only considers a single fault per lost feeder; states involving n+1 multiple faults are only considered after all state hypotheses involving n faults are proven inconsistent. Naturally enough, the probability distribution at any time considers that the switching devices that have not yet been operated are in their normal mode. This is not restrictive.

A domain-specific planner takes as input the most probable state supplied by the diagnosis reasoner. It returns a plan (sequence of switching operations) isolating the faults and restoring the supply to the non-faulty lines in that state. In order to cope with complexity, the search space is restricted to plans that only extend existing feeders, i.e., do not discharge any circuit-breaker of part of its load after the incident. In fact, other types of plans are rarely used in the reality, and the space of such admissible plans for the given state is small enough to be entirely explored. For our network example in Figure 1, it contains most of the time less than a hundred plans. These are all evaluated using a utility function that captures the criteria mentioned above, and the best one is returned to the diagnostic reasoner. As another concession to complexity, risks due to partial observability and actuator failures are not taken into account in the evaluation. Therefore, the returned plan is optimal iff the state hypothesis provided by the diagnostic reasoner is correct and none of the actuators changes its behavior mode during the process.

Further details on the principles underlying SyDRe, including the actual construction of plans, are out of the scope of this presentation. They can be found in (Thiébaux *et al.* 1996b). SyDRe has been tested on a number of simulated scenarios involving multiple faults on multiple feeders of semi-rural and urban networks, and could restore the supply in real-time. The restoration sessions have been judged satisfactory by experts. SyDRe is currently being integrated to the AUSTRAL software platform for power distribution systems (Krivine and Jehl 1996), a platform which will be operational in two French distribution centers by the beginning of 1998. However, much remains to be done with power distribution systems in the directions of a better assessment of risks and of a more general system.

Acknowledgments

This work has been supported by contract R34/1H9487 between IRISA and EDF. We thank the anonymous reviewers for useful comments.

References

A. Beschta, O. Dressler, H. Freitag, M. Montag, and P. Struß. A model-based approach to fault localisation in power transmission networks. *Intel. Syst. Eng.*, 1-2:190–201, 1993.

C. Boutilier and D. Poole. Computing optimal policies for partially observable decision processes using compact representations. In *Proc. AAAI-96*, pages 1168–1173, 1996.

P. Bredillet, I. Delouis, P. Eyrolles, O. Jehl, J.-P. Krivine, and P. Thiault. The AUSTRAL expert system for power restoration on distribution systems. In *Proc. ISAP-94*, pages 295–302. EC2, 1994.

A. Cassandra, L. Kaelbling, and M. Littman. Acting optimally in partially observable stochastic domains. In *Proc. AAAI-94*, pages 1023–1028, 1994.

J. Crow and J. Rushby. Model-based reconfiguration: Towards and integration with diagnosis. In *Proc. AAAI-91*, pages 836–841, 1991.

J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.

G. Friedrich and W. Nejdl. Choosing observations and actions in model-based diagnosis-repair systems. In *Proc. DX-92*, pages 76–85, 1992.

A. Gonzalez, R. Morris, F. McKenzie, D. Carreira, and B. Gann. Model-based, real-time control of electrical power systems. *IEEE Trans. on Systems, Man, and Cybernetics*, 26(4), 1996.

Proc. International Conference on Intelligent System Application to Power Systems (ISAP-96), Orlando, FL, 1996.

J.-P. Krivine and O. Jehl. The AUSTRAL system for diagnosis and power restoration. In *Proc. ISAP-96*, 1996.

C. Moinard. Diagral : a diagnosis and repair system. In *Proc. DX-94*, pages 203–211, 1994.

Y. Sun and D. Weld. Beyond simple observation: Planning to diagnose. In *Proc. DX-92*, pages 67–75, 1992.

S. Thiébaux, M.-O. Cordier, O. Jehl, and J.-P. Krivine. Génération de plans de reprise de service pour les réseaux de distribution d'EDF. Research Report 3006, INRIA, October 1996. In French.

S. Thiébaux, M.-O. Cordier, O. Jehl, and J.-P. Krivine. Supply restoration in power distribution systems — a case study in integrating model-based diagnosis and repair planning. In *Proc. UAI-96*, pages 525–532, 1996.

B. Williams and P. Nayak. A model-based approach to reactive self-configuring systems. In *Proc. DX-96*, pages 267–274, 1996.

Appendix

Available data and illustrative example

The data available on this reference problem include the specifications of representative sections of several EDF power distribution systems (for semi-rural or urban environments), including normal configuration, maximal power constraints and load curves for circuit-breakers and lines. Functions estimating the utility of network configurations exist. They take into account, for instance, the satisfaction of the power demand, critical customers, power margins of circuit-breakers, and the distance to the normal configuration. Subjective probabilities on the behavior modes of sensors and actuators for a given network could be gathered.

Available software includes a network simulator, which can be used to produce realistic data for any kind of faults on any network having the topology described in this paper.[3] It can also simulate the performance of a switching operation on a given state of the network, and serve as a predictive model. Finally, SyDRe (see Figure 5) could be used as a reference basis for comparative tests. We now give an example of a supply restoration session with SyDRe. Other examples can be found in (Thiébaux *et al.* 1996a).

Taking our example network in Figure 1, two faults cause CB 1 and CB 5 to open: one on the feeder fed by CB 1 between SD 13 and SD 53, and the other on the feeder fed by CB 5 between SD 41 and SD 52. Furthermore, the fault detector of SD 13 and the actuator of SD 52 lie, and the position detector of this latter is broken. All this is unknown to SyDRe, which can solely observe that the two circuit breakers are open, and that the fault detectors of SDs 11, 12, 50, 51, and 52 are the only ones to indicate a fault downstream.

Given these observations, the most probable hypothesis is a fault between SD 12 and SD 13 (which is incorrect) and another between SD 41 and SD 52 (which is correct). The plan is then to open SDs 12 and 13 to isolate the first fault, to open SD 52 to isolate the second, to re-close CB 1 to resupply the lines upstream of the fault on the first feeder, to re-close CB 5 and to close SD 53 to resupply both the downstream lines on the first feeder and the upstream lines on the second one. Note that after operating SD 52, it is unknown whether this one is really open because its position detector is broken, and in fact it is still closed because

[3] The simulator is implemented in Standard ML. We are in the process of building another simulator in Ilog Talk. This latter accounts for a number of subtilities that do not add much to the problem complexity, and which we omitted in order to make our presentation suitable for a reference problem. Among those are for instance hardware-dependent delays of between one millisecond and one minute in the transmission of the sensing information.

```
--------------------------------------
Current hypothesis:
  [SD12, SD13] and [SD41, SD52] faulty

Current plan:
  open SD12, SD13, and SD52
  close CB1, CB5, and SD53

I open SD12
  AC notification: positive
  PD information: open

I open SD13
  AC notification: positive
  PD information: open

I open SD52
  AC notification: positive
  PD information: unavailable

I close CB1
  AC notification: positive
  PD information: closed

I close CB5
  AC notification: positive
  PD information: open

--------------------------------------
Current hypothesis:
  [SD12, SD13] and [SD41, SD52] faulty
  AC of SD52 liar

Current plan:
  open SD51
  close CB5 and SD53

I open SD51
  AC notification: positive
  PD information: open

I close CB5
  AC notification: positive
  PD information: closed

I close SD53
  AC notification: positive
  PD information: open

--------------------------------------
Current hypothesis:
  [SD13, SD53] and [SD41, SD52] faulty
  AC of SD52 liar
  FD of SD13 liar

Current plan:
  open SD53
  close SD12 and CB5

I open SD53
  AC notification: positive
  PD information: open

I close SD12
  AC notification: positive
  PD information: closed

I close CB5
  AC notification: positive
  PD information: closed

Restoration process terminated
```

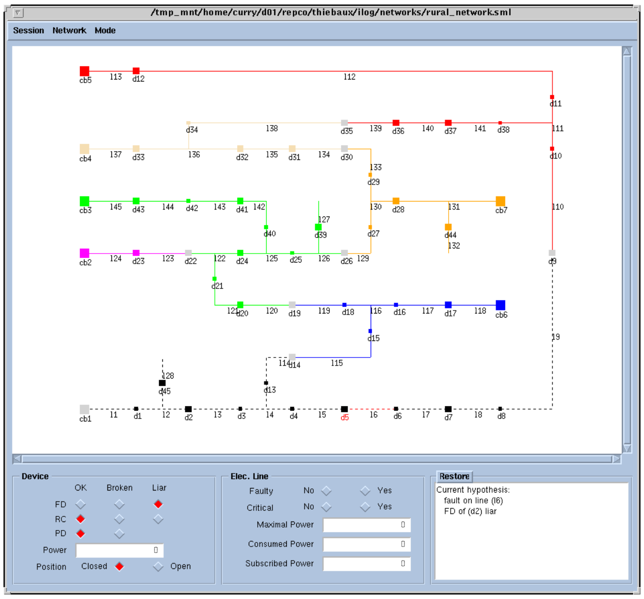

Figure 5: SyDRe

its actuator is lying. Thus, the second fault is not correctly isolated. Therefore, when we attempt to close CB 5, this one opens because it is still feeding the fault. This leads to a change of most probable state hypothesis and to the revision of the current plan, which is materialized by a dash line in the session trace.

From the new observations (SDs 50, 51, and 52 still indicate a fault downstream) and the knowledge of the previously executed actions, it becomes clear that the actuator of SD 52 is liar. The most probable fault locations, however, remain the same. The plan is to open SD 51 instead of SD 52 and to go on as was planned before, that is to close CB 5 and SD 53. This latter switching operation causes CB 5 to open again, since it is now feeding the wrongly localized fault between SD 13 and SD 53. This leads to another change of most probable hypothesis and to a revision of the plan.

From the new observations (SDs 50 and 53 indicate a fault downstream) and the success of the previous closing operation on CB 5, it is immediate that the most probable fault location hypothesis on the first feeder is now the right one (between SDs 13 and 53), and that the fault detector of SD 13 is liar. The plan is then to isolate this fault by re-opening SD 53 (the other side is already isolated because SD 13 has been opened), to resupply the part between SD 12 and SD 13 (which had been isolated) by closing SD 12, and to resupply the second feeder by re-closing CB 5. The completion of these operations ends the restoration process.