Examples of cooperative situations and their implementation

R.P. Aarnts (Volmac, Netherlands)
J. Corera (Iberdrola, Spain)
J. Pérez (Labein, Spain)
D. Gureghian (Framentec, France)
N. Jennings (Queen Mary and Westfield College, England)

A hot topic in the world of Distributed AI is currently cooperation. Although several researchers are working on this topic, no clear working definition of cooperation exists. Several of them have tried to provide a definition but none of them seems to be workable. The ESPRIT project ARCHON (funded by the CEC) is also suffering from the lack of a working definition for cooperation. Several exhaustive discussions within the project did not result in a usable definition. While none of the researchers working on the project was able to provide such a definition, almost everybody was able to provide examples of situations that can be called cooperative. It is rather easy to recognize cooperation but to describe it is hard.

In order to overcome this problem, members of the Archon project started collecting examples of cooperation (cooperative situations) in order to make clear what they mean by cooperation. Furthermore they wanted to convince the world of the applicability of the cooperation paradigm and therefore they implemented some of the collected cooperative situations.

This article will describe a few of these implemented cooperative situations in order to show the outside world what the Archon team means by cooperation and how it can be applied.

The examples provided in this article are all related to one of the demonstrator applications involved within ARCHON, the Iberdrola (formerly Iberduero) application. After a brief introduction of Iberdrola, an overview of the application and its expert systems will be given. This background information is necessary in order to understand the subsequent section, the cooperative situations. Finally, the implementation of a demonstrator and a few cooperative situations are discussed.

Iberdrola

Iberdrola S.A. is an electricity utility located in northern Spain with an installed generation of about 6,500 MW, 4,200 of which are hydraulic. Distribution takes place in an area of over 100,000 square kilometers with a population of 7,000,000 inhabitants. The generation and transmission is controlled from a central dispatch, equipped with a control system that was put into operation in the year 1981. The large nuclear power units that were put into operation during the last few years have increased the probability of a disturbance in the network. This induced Iberdrola in 1987 to undertake the development of an expert system that could assist the operators in the dispatch center to analyze and recover from disturbances. Right from the beginning it was clear that the problem was too complex for a single expert system and alternatives were necessary. Cooperation between semiautonomous expert systems seems to be a very

©Copyright 1991, by Vleermuis Software Research b.v.
Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the Journal reference and VSR copyright are included on the first page. The title and abstract, but no other portions, of this article may be copied or distributed royalty free without further permission. Permission to publish any other portions of this article must be obtained from the Editor.

Acknowledgement

This article is based on the work performed in the first ARCHON Task Force (November 1990-February 1991). The consortium for the ESPRIT 2256 project ARCHON consists of the following partners: Krupp Atlas Elektronik (Germany), AMBER (Greece), CERN (Switzerland), CNRG-NTUA (Greece), EA Technology (England), Framentec (France), FWI University of Amsterdam (Netherlands), Iberdrola (Spain), IRIDIA Université Libre Bruxelles (Belgium), JRC Ispra (Italy), Labein (Spain), Queen Mary and Westfield College (England), University of Porto (Portugal) and Volmac (Netherlands).

promising alternative for the development of Iberdrola's large applications and so they did not hesitate to join the ARCHON consortium when it was founded in 1989. Iberdrola's main impact on the project is the provision of an application (test-bed) containing several expert systems which should operate in a cooperative manner at the end of this project.

Overview of the Iberdrola application

A high-voltage transmission network is in general controlled at two levels: at a local level and at a global level. At the local level the network contains several automatic safety and recovery possibilities. In the case of a short circuit or a loss of power, breakers will open automatically. When the power comes back again, the breakers will close in order to restore the power supply to the customers.
At the global level control takes place from a central dispatch center situated in Bilbao. This dispatch center receives information from all crucial elements in the network by making use of telemetry.

Management of high-voltage transmission networks in emergency situations is very complex, due to the large number of constraints to be taken into consideration and the insufficient quality of the information arriving at the control center. An emergency situation typically originates from a short circuit in a line, a bus bar or a transformer, and can be worsened by equipment malfunctioning (that is, a breaker failing to open) or subsequent overloads (for example, a house of cards: triggering of a line overloads other lines, which will also trigger...). Actions to restore service must be taken rapidly and accurately, as in this situation the network is less tolerant to operation errors. Actions will basically consist of breaker operations, topology changes and activation/deactivation of automatisms and protective relays.

Briefly, the sequence of steps to follow in the occurrence of a disturbance are:

- detect the origin of the disturbance: if it is a fault (a short circuit), or a pure overload.
- in case of a fault, identify damaged elements in order to exclude them in the restoration process.
- track the black-out area as the disturbance evolves.
- prepare a restoration plan and consider how the execution of it has to be carried out.

In the past all these tasks had to be carried out by the operator. The main problem is how to maintain an overview of what is happening. The slightest mistake during the restoration process can damage the network. Damage is not only expensive, it can also jeopardize the restoration at the risk of causing large areas (cities) without power for a long period. This can be very critical. For this reason Iberdrola started to develop a Network Restoration Assistant System. This system is capable of performing all the above mentioned tasks except the execution of restoration plans. The restoration is still foreseen to be carried out by the operator. But it is the system that keeps track of what is happening.
Since these tasks are too complex to make use of one single expert system, the application has been subdivided into a few smaller expert systems.

The distribution of tasks within the expert systems of the application cannot be done without considering the characteristics of the information available. Actually this is the big constraint of the application, and the use of redundant information together with cooperation between the agents is required for good overall results. The main types of information used by the application and provided by the control system of the network are the following:

- **Alarms** including operation of breakers and protective devices; their time reference is not precise (a tolerance of ± 5 seconds), but they are produced as soon as the control center is aware of the operation through the SCADA system (the SCADA system is the interface between the application and the network). This can cause a disordering of the "alarms".

- **Chronological alarms** they contain the same information as mentioned above, but their time tag is far more precise; the disadvantage is that they are produced at the substation, and arrive at the control center with a delay.

- **Snapshots** including breakers' and switches' statuses, and selected analog measures (that is, bus bar voltages and power flows); each snapshot is a relatively consistent picture of the whole network; the disadvantage is that information on the operation of breakers which return to the same state (that is, open&close) is lost, if the operation takes place in between snapshots.

Agent description

The expert systems as discussed before were transformed into ARCHON agents in order to achieve a cooperative behavior between them. No major modifications had to be made on the actual expert systems, instead they got an Interworking Layer on top of them. It is the Interworking Layer that provides the cooperation between the agents.

The Interworking Layer has control over the underlying expert system. It therefore sends messages to the expert system. Expert systems can only communicate with other expert systems via their Interworking Layer. The type of control the Interworking Layer has over the expert system is of the kind "strongly request". It is up to the expert system to decide whether to schedule a certain task or not.

For the purpose of this article it suffices to look at an agent as an independent piece of software, capable of performing several tasks and with its own local goals.

The Iberdrola application is composed of several agents. The following section provides a brief overview of the agents and their tasks as far as they are relevant for the examples of cooperative situations.

CSI: Control System Interface

The main goal of this agent is to acquire information about the network from the Control System and to provide the information to the community of agents. The information is available in three kinds: alarms, chronological alarms and snapshots. The CSI processes the information and sends it to interested agents. In order to achieve its goal, the CSI is capable of performing several tasks. The most relevant ones with respect to the examples are:

- ALARM_MESSAGES_ACQUISITION
  This task receives the alarms from the Control Center and combines them into blocks of alarms, probably belonging to one state change of the network. The algorithm used to determine which alarms belong to one state change is rather simple: as long as messages arrive within ten seconds after the last message received it is assumed they belong to the same state change of the network. If no new message arrives within ten seconds, the whole block of alarms will be sent to interested agents.

- CHRONOLOGICAL_MESSAGES_ACQUISITION
  This task collects chronological alarms and combines them into blocks of chronological alarms. These alarms are time-stamped and thus they can be sorted. Every alarm message has also a corresponding chronological alarm which will be prepared, but this arrives with a delay. Due to this fact, the CSI knows which chronological alarms to expect. If, for some reason, the chronological alarm is missing (did not arrive in time), the alarm message without time-stamping will be used.

- DISTURBANCE_DETECTION
  The purpose of this task is to analyze the blocks of alarms on possible disturbances within the network. This is the first possibility for the system to detect that things go wrong within the network. The result of this task is an indication that something might be wrong. This indication is sent to interested agents.

- PROVISION_OF_SNAPSHOTS
  Snapshots are prepared by the Control System. It is this task that acquires the snapshots from the Control System. The frequency by which the snapshots are acquired depends on the needs of other agents. Under normal conditions, that is, no disturbance detected yet, snapshots are provided every 15 minutes.

- LOAD_FLOW_ANALYSIS
  In the case one of the agents needs to know what the load flow is if certain elements of the network change state, they can request the CSI to perform this task. The CSI does not perform the calculation itself. Instead it sends the problem to the Control System which is equipped with a load-flow calculation program. The LOAD_FLOW_ANALYSIS task is an interface to this program.

AAA: Alarm Analysis Agent

The main goal of the AAA is to analyze the network and to identify possible elements in fault. The AAA makes use of the blocks of alarms with a disturbance detection as supplied by the CSI. This is the earliest available information to be used for an analysis about a possible disturbance. The result is a list of hypotheses about elements in fault.

- HYPOTHESIS_GENERATION
  This task generates a list of elements possibly in fault. It will also calculate certainty factors for these elements. The hypotheses list is sorted based on these certainty factors.

- HYPOTHESIS_VALIDATION
  This task will analyze the hypotheses in depth. The hypotheses of the fault are modelled according to the behavior of the elements and the automatic devices involved. The comparison between the theoretical results for each hypothesis and the actual alarms will allow validation of the hypothesis. This will produce a new list of validated hypotheses with new certainty factors. This task can run for

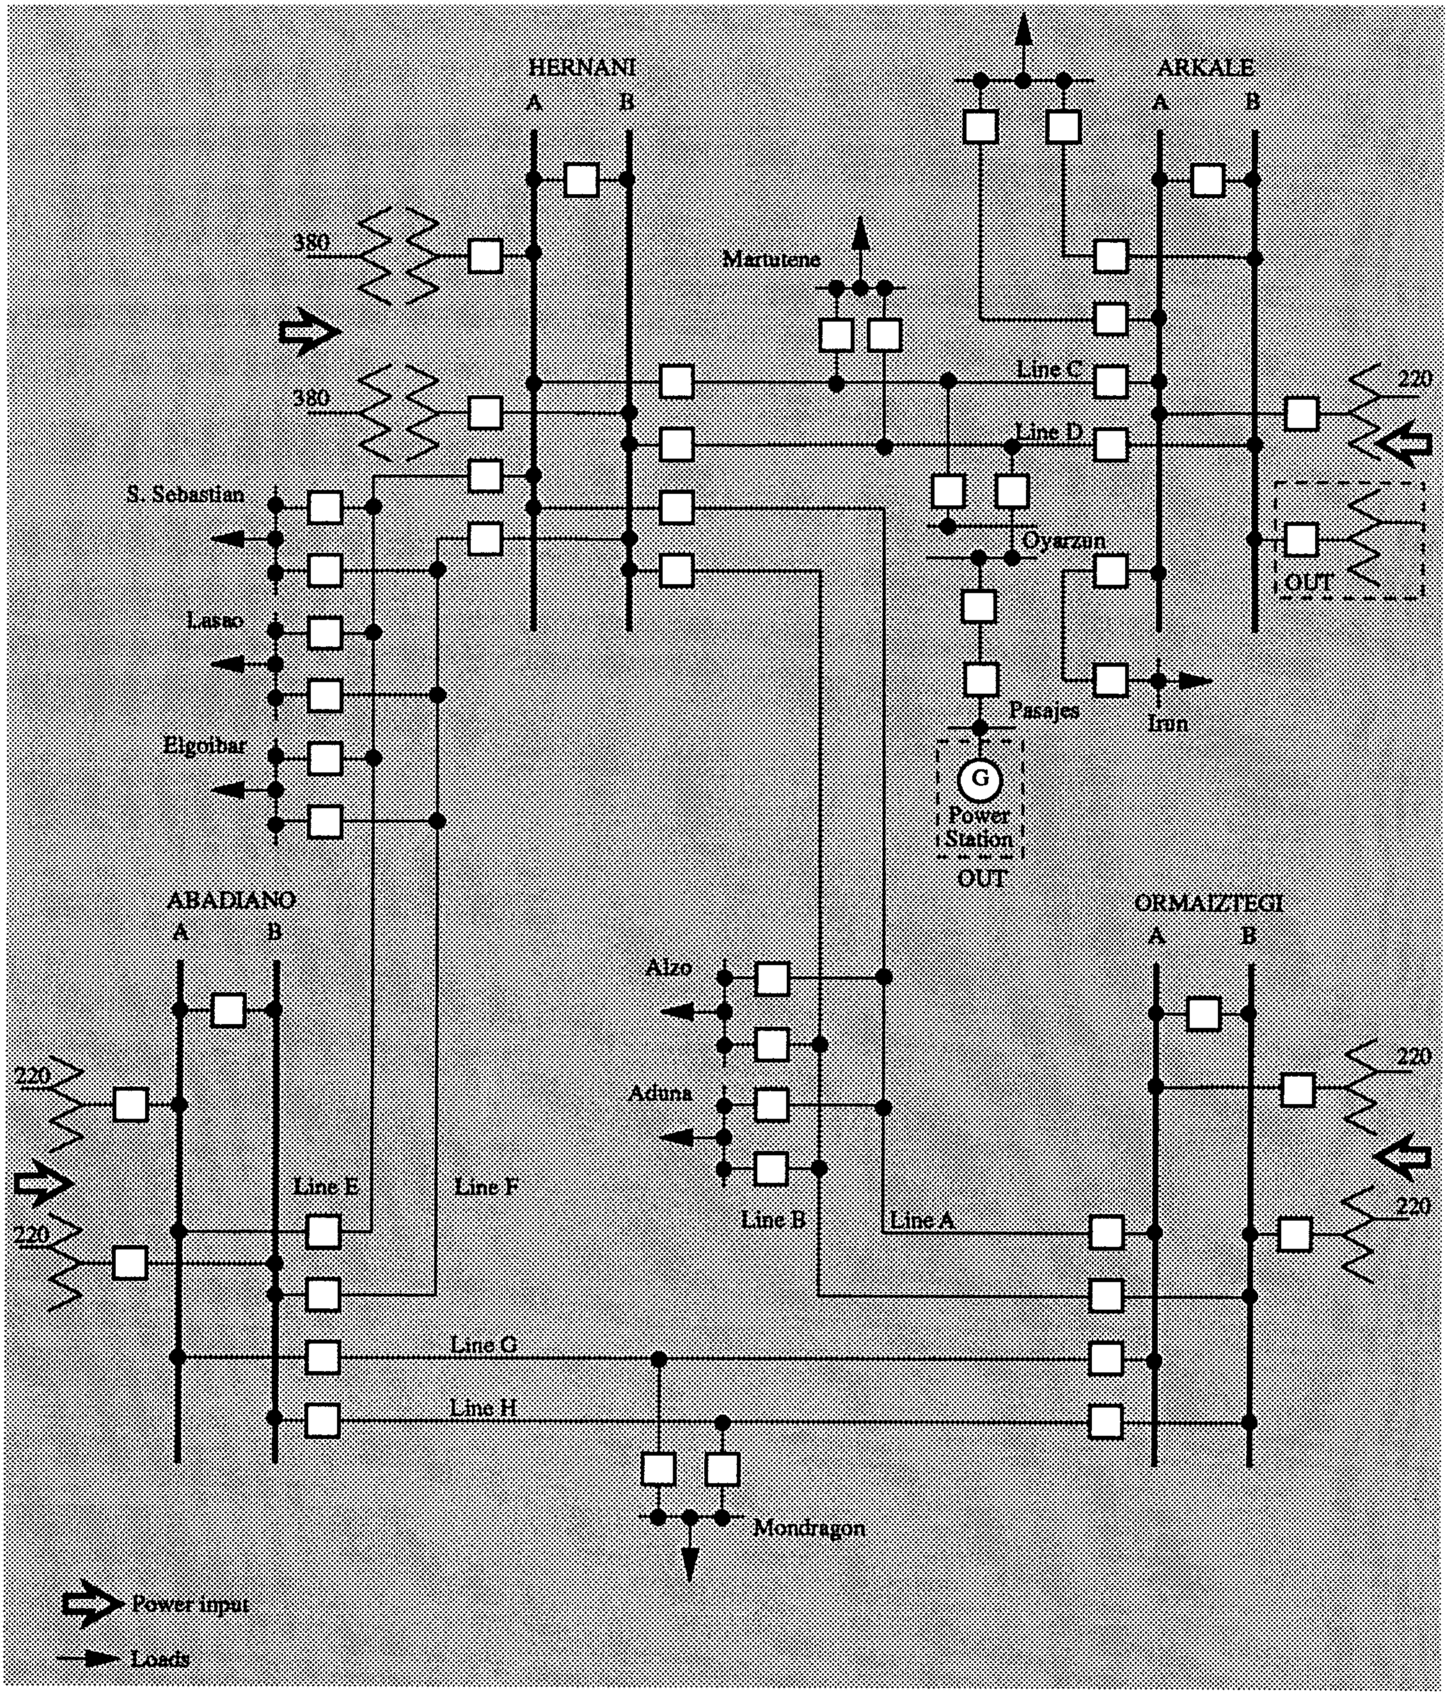

Figure 1 Part of the network controlled by Iberdrola

one hypothesis at a time. The typical duration of this task is approximately 20 seconds.

- HYPOTHESIS_REFINEMENT
The list of hypotheses can be reduced before the validation task is performed by making use of the list of elements of the initial black-out area. Elements in the initial black-out area cannot be in fault, so they can be removed from the list.

- RECLOSING_TRIALS_ANALYSIS
If a fault occurs in the network, the protective relays will isolate it by tripping the breakers. This will cause several alarm messages, which are combined to the first block of alarms after the disturbance occurred. However after a short delay the automatism of the network will try to restore the lost power. Breakers will be closed and if the fault still exists they trip again. These reclosing trials will cause successive blocks of alarms. This task is responsible for the analysis of these reclosing trials in order to determine how the fault evolves trough the network. This task is also capable of detecting alarms not belonging to the current disturbance.

- DISTURBANCE_DETECTION
If the task RECLOSING_TRIALS_ANALYSIS detects a subset of alarms, not belonging to the current disturbance, the subset has to be analyzed. These alarms can be the result of ordinary maintenance operations or of an isolated second disturbance. It is necessary to determine whether this subset indicates a second disturbance or not. Disturbance detection is the purpose of this task.

BRS: Breaker and Relays Supervisor
The goal of this agent is similar to that of the AAA but it uses blocks of chronological alarms instead of blocks of alarms. The diagnosis process can be more accurate but will also start later because chronological alarms will arrive with a delay at the Control System. The tasks the BRS can perform are similar to the tasks of the AAA.

BAI: Black-out Area Identifier
This agent's main goal is to identify the elements in black-out, that is, without power. The BAI can perform the following tasks:

- INITIAL_AREA_OUT_OF_SERVICE
This task determines which elements are in black-out during the first notification of a disturbance. The analysis is done using alarm messages.

- BLACK_OUT_AREA_FOLLOW_UP
A distinction is made between elements in black-out during the first notification of a disturbance and the elements in black-out during the phase after the disturbance took place. This task identifies which elements are in black-out during the phase after a disturbance took place. The analysis is done using snapshots. In order to be as up-to-date as possible, snapshots have to be provided with the shortest possible interval.

- CHECK_RESTORATION_PREMISES
The purpose of this task is to track elements needed for a restoration plan on availability. As soon as one of the elements needed is in black-out, it is no longer available. In order to perform this task, it is necessary to understand restoration plans.

SRA: System Restoration Agent
This agent is responsible for the provision of a network restoration plan. This plan will be presented to the operator (via the UIA) who is responsible for the execution of the plan.
The SRA will monitor the execution of the plan by analyzing state changes of the network. In case of an exception or a deviation from the restoration plan, the SRA will stop the execution of the plan and start planning again.
Before the SRA can create a restoration plan, an analysis of the network has to be done with respect to external network security, criticality of the situation and the availability of possible sources of energy.

UIA: User Interface Agent
The main purpose of this agent is to provide the interface between end users and the multiagent system. The view of this agent is domain-related. In the particular case of the Iberdrola application tis concerns the network, the alarms, the hypotheses, restoration plans, etcetera.

Cooperative Situations

In the following paragraph we will discuss several examples of cooperative situations.

Situation 1: AAA-BAI
The first example shown is between the AAA and the BAI. Taking a better look at the AAA learns that the agent is responsible for the analysis of the network and the identification of elements in fault. The AAA starts this work as soon as it receives a block of alarms with a disturbance indication. Based on the block of alarms and the already prepared topology of the network, the AAA is able to determine which elements are in fault. Additional information or assistance from other agents is not necessary. The AAA can perform its task autonomously.

The process of identification of elements in fault is carried out in two stages. First of all the AAA will generate a list of hypotheses about elements possibly in fault (HYPOTHESIS_GENERATION). For each hypothesis a certainty factor will be calculated. The list of hypotheses is sorted using these certainty factors. During the second stage the hypotheses are analyzed in depth (HYPOTHESIS_VALIDATION). This second stage is a very time-consuming process (typically 20 seconds per hypothesis) so it is worth skipping hypotheses of elements which are unlikely to be in fault.
Faulty elements must be isolated from energy sources at least during the time they are in fault. It is the BAI that is capable of providing a list of elements which have been isolated from the energy sources during the first set of alarms, called the Initial Black-out Area. For this reason, any element which is not in the Initial Black-out Area cannot be in fault.

If the list of hypotheses and the initial black-out area as prepared by the BAI are available, the AAA can make use of them by deleting all hypotheses of which the elements are not contained in the initial black-out area. For this purpose an additional task for the AAA is defined called HYPOTHESIS_REFINEMENT. The typical duration of this task is approximately 2 to 3 seconds, so it pays off to reduce the hypotheses list using this task (when compared to the duration of the validation task).

The task HYPOTHESIS_REFINEMENT is scheduled when the task HYPOTHESIS_GENERATION has finished and the initial black-out area is available. If the initial black-out area is not available the task HYPOTHESIS-VALIDATION will be scheduled for each element in the list of hypothesis. If the initial black-out area arrives during the execution of this task, the scheduling of the task HYPOTHESIS_VALIDATION will be stopped and the task HYPOTHESIS_REFINEMENT will be started in order to reduce the remaining list of hypotheses. After the refinement task has been finished, the validation task will be started again.

Conclusion: if the initial black-out area is available, the AAA will make use of it in order to improve both performance and quality.

Situation 2: BRS-CSI

A similar cooperative situation as between the AAA and BAI can be found between the BRS and the BAI because the diagnosis processes are very similar. Similar to the AAA, the BRS goes through the process of first generating hypotheses and then validating them. The main difference is the kind of alarm message used for the analysis. The BRS uses chronological alarms.

If an Initial Black-out Area is available, the BRS can make use of it by reducing the list of hypotheses to be validated. However, in this case the Initial Black-out Area covers several blocks of alarms analyzed by the BRS. For this reason, the same Initial Black-out Area can be used for several blocks of alarms.

Situation 3: CSI-AAA

One of the tasks of the AAA is to analyze the reclosing trials (RECLOSING_TRIALS_ANALYSIS). Due to the automatism in the network, breakers will be closed automatically. If there is still a disturbance in the network, those breakers that feed the element in fault will be opened again by the protective relays. This information can be used to determine if a disturbance is permanent or transient (in which case no restoration plan has to be prepared).
This task is also able to determine whether or not alarms belong to the disturbance currently under analysis. Alarms not belonging to this disturbance can be the result of ordinary maintenance operations or a new (second) disturbance. The selection of alarm messages not related to the disturbance currently under analysis should be done by the AAA and not by the CSI because it is necessary to have knowledge about the topology of the network and the CSI lacks this knowledge.
During this task of reclosing trials analysis, the AAA may detect that some of the alarms are not related to the disturbance being analyzed. It is at this point that the AAA would need to activate a task for disturbance detection on the identified subset of alarms. This rather important task can be performed by the AAA. However, when the AAA detects a subset of alarms, it is usually busy performing other tasks (for example, RECLOSING_TRIALS_ANALYSIS).

Under normal conditions, that is, no disturbance has been detected yet, the CSI will check each block of alarms on a possible disturbance. When a disturbance is detected, the CSI will no longer perform this task until the end of the disturbance has been signalled. The reason for this is that the CSI does not have the ability to make the distinction between two separate disturbances. The CSI can only determine if one or more alarm messages of a block of alarms are the result of a disturbance.
Since the AAA is usually busy, it makes sense that in the case of a subset of alarms not belonging to the current disturbance, the CSI performs its task DISTURBANCE_DETECTION using this subset of alarms. If the CSI has enough resources and performs its task, the AAA will be relieved.

Situation 4: BAI-SRA

The SRA prepares plans for the restoration of the network. It also sends these plans to the operator and monitors the execution of the plan which is performed by the operator. In case of a failure of a certain restoration action, the SRA has to stop the execution of the plan and start preparing a new one.

The BAI is capable of detecting whether a restoration plan is no longer feasible (for example, if an element necessary for the restoration process drops out). The BAI is capable of detecting this long before the SRA does, because the SRA only reasons about the current restoration step. The SRA only finds out when the step using the unavailable element has to be executed.

If the BAI and the SRA cooperate, the restoration process can be improved. If the BAI receives the plan from the SRA, it is capable of determining the feasibility of the plan. If one or more of the elements needed in the restoration process become unavailable during the execution, the BAI can inform the SRA, which will then stop the execution of the plan and restart planning again.

As an alternative, the SRA can send a request to the BAI to monitor the availability of certain elements instead of sending complete plans. In this case it is not necessary for the BAI to have any knowledge about restoration plans.

Situation 5: BRS-AAA

Both the BRS and the AAA diagnose faults in a similar way. They both go trough the process of hypothesis generation, refinement and validation. The main difference lies in the fact that they use different information. Due to this one could imagine that the BRS and AAA can provide conflicting hypotheses. The quality of the hypothesis can be improved if the BRS and AAA decide to cooperate.

Assume both agents have completed their hypothesis generation tasks at more or less the same time and they have exchanged their hypotheses sets. If the first, let's say four, hypotheses of both sets are the same, then there is no conflict and the confidence in these hypotheses is increased. Each agent can then increase the certainty factors of these hypotheses.

If they are different, various actions can take place, depending on the alarms received by the BRS. The BRS receives more accurate alarms because they contain a time stamping, but when a chronological alarm gets lost it is replaced by the corresponding alarm without time stamping (done by the CSI). This decreases the reliability of the alarms the BRS receives.

If the BRS receives only chronological alarms, its diagnosis is based on better information and therefore its diagnosis is more reliable. The hypotheses set of the AAA should then be reordered by increasing the certainty factor of the elements, using the BRS set. Elements not contained in the set of the BRS will not be removed if they have enough certainty factors.

When the BRS has received some alarms without time stamping, the quality of its diagnosis is decreased. In that case the certainty factors of the BRS set will be adapted using the set of AAA.

The described situation occurs under the assumption that both the BRS and the AAA complete the generation task at about the same time. However one could imagine that the BRS completes this task later because it makes use of time-stamped information that arrives later. In that case it may be possible to let the interaction take place between the refined hypothesis of the AAA and the generated hypothesis of the BRS.
The weighting for the refined hypothesis will be greater than for the generated hypothesis produced by the AAA.

Demonstrator Implementation

In order to prove the applicability of the developed concepts for cooperation, the ARCHON project decided to develop a demonstrator. This demonstrator is based on the Iberdrola application. Since this demonstrator had to be used at several sites it was impossible to make use of the real expert systems and thus simulated versions were used.
Almost all agents mentioned earlier in this article were implemented. Only the BRS was left out in order not to exhaust the budgeted time for this experiment.

Although the expert systems were simulated (predefined responses on certain triggers known in advance), the ARCHON Layer, which is composed of several modules, was fully implemented for all agents.

The demonstrator was then used to realize some of the cooperative situations mentioned.

One of the lessons learned during this experiment was that it is quite easy to change the behavior of an agent without changing the underlying expert system. This behavior is changed by changing the mechanism by which agents send/receive their message. If situation 1 is taken as example, there are three principal mechanisms by which an initial black-out area can reach the AAA:

- The AAA has knowledge about the BAI and knows that this agent is able to determine the initial elements in black-out. As soon as the AAA

receives a block of alarm messages with a disturbance indication, it will send a request to the BAI to provide the initial black-out area.
- The BAI has knowledge about the AAA and knows that this agent can make use of the initial black-out area. Therefore the BAI will try to provide the AAA with this information as soon as it is available.
- The BAI has no knowledge about a specific agent in its community that is capable of making use of the initial black-out area. The BAI determines the initial black-out area anyhow, due to its own local goals. The BAI now acts as a "chatty" agent by sending this information to all other agents in the belief that there must be an agent interested in it.

The behavior of agents is determined by the agent models, which are part of the ARCHON Layer.

Conclusion

With this demonstrator implemented, the ARCHON project was able to demonstrate some of the cooperative situations mentioned. The main conclusion of this demonstrator exercise is that it has been proven that the cooperative schemes mentioned are feasible and that the ARCHON architecture proposed is suitable for the implementation of these cooperative schemes.

Another conclusion is that without changing a task of the underlying expert system, it is very well possible to change the behavior of the agent by making changes at the ARCHON Layer. This conclusion is very important for the use of already existing expert systems in a multiagent system based on the ARCHON architecture.
This positive experience proved that ARCHON is on the right track. For this reason it is intended to continue the work with the implementation of the ARCHON architecture in an environment where real expert systems are used. This will not only be done for the Iberdrola application but also for all the other applications involved in the ARCHON project.

Literature

Aarnts R.P. and Groote O.C. de, "ARCHON, (ARchitecture for Cooperating Heterogeneous ON-line systems)". Journal of Software Research, Special Issue, pp. 59-68, December 1989.

Aarnts R.P. and Groote O.C. de, "ARCHON: a matter of cooperation". Journal of Software Research, Volume 3 no. 1, pp. 96-103, February 1991.

Corera J.M., Echavarri J., Laresgoiti I. and Lazaro J.M., ID/ARCHON/TN-004/09-90, "Structural design of the Service Restoration Agent".

Corera J.M. and Laresgoiti I., ID/ARCHON/TN-003/06-90, "Iberduero Application Scenario".

Corera J.M. and Laresgoiti I., ID/ARCHON/TN-005/01-90, "Iberduero Application Scenario II".

Echavarri J., Amantegui J. and Corera J.M., ID/ARCHON/TN-002/10-89, "Specification of Iberduero Application".

Echavarri J. and Laresgoiti I., LAB/ARCHON/TN-001/06-89, "Preliminary Structural Design of the Black-out Area Identifier".

Jennings, N.R., "Cooperation in Industrial Systems". In: "Proceedings of the Esprit Conference '91". North Holland, 1991.

Laresgoiti I. and Barandiaran J., LAB/ARCHON/TN-003/12-89, "Structural Design of the Breakers and Relays Supervisor".

Perez J. and Laresgoiti I., "User guide of the BAI Intelligent System". ■