Sophisticated Cooperation in FA/C Distributed Problem Solving Systems*

Norman Carver
Zarko Cvetanovic
Victor Lesser

COINS Technical Report 91-23
February 1991

Abstract

In the functionally accurate, cooperative (FA/C) distributed problem solving paradigm, agents exchange tentative and partial results in order to converge on correct solutions. The key questions for FA/C problem solving are: how should cooperation among agents be structured and what capabilities are required in the agents to support the desired cooperation. To date, the FA/C paradigm has been explored with agents that did not have sophisticated evidential reasoning capabilities. We have implemented a new framework in which agents maintain explicit representations of the reasons why their hypotheses are uncertain and explicit representations of the state of the actions being taken to meet their goals. In this paper, we will show that agents with more sophisticated models of their evidence and their problem solving states can support the complex, dynamic interactions between agents that are necessary to fully implement the FA/C paradigm. Our framework makes it possible for agents to have directed dialogues among agents for distributed differential diagnosis, make use of a variety of problem solving methods in response to changing situations, transmit information at different levels of detail, and drive local and global problem solving using the notion of the global consistency of local solutions. These capabilities have not been part of previous implementations of the FA/C paradigm.

*This work was supported by the Office of Naval Research under DARPA contract N00014-89-J-1877 and University Research Initiative grant number N00014-86-K-0764.

1 Introduction

In the functionally accurate, cooperative (FA/C) systems paradigm for distributed problem solving [Lesser & Corkill 1981, Lesser 1991], agents need not have all the information necessary to completely and accurately solve each of their subproblems. The basic intuition behind this approach is that for many applications the subproblems that need to be solved by the different agents are not independent; there exist constraints among the subproblems. These constraints can be exploited to partially resolve the inconsistencies and uncertainties that occur in local problem solving due to the lack of accurate, complete, and up-to-date information. In the FA/C paradigm, each agent's local problem solving is organized so that partial and tentative results can be produced despite incomplete and uncertain information. When these partial results are exchanged among agents working on interdependent subproblems, the agents use them to constrain the possible solutions to their subproblems. This allows the local problem solving uncertainties which result from incomplete, incorrect, and inconsistent information to be partially resolved. Resolution can take the form of producing more complete partial results, resolving solution uncertainty due to competing, alternative partial solutions, detecting inconsistencies in previously generated results (either locally generated or received from other agents), and speeding up local problem solving because the space of possible solutions that needs to be examined is constrained.

The key question for FA/C problem solving is how cooperation among agents should be structured so that an acceptable answer can be converged upon within a reasonable amount of time, with limited communication between the agents. A subsidiary question, but one of equal importance, is what reasoning capabilities are required in agents in order to support such cooperation. To date, the exploration of the FA/C paradigm has been done with agents that did not have sophisticated evidential reasoning capabilities (e.g., the agents used in the DVMT [Lesser & Corkill 1983]). These agents had poor representations of the evidential relationships between competing, alternative hypotheses and they could not explicitly consider why existing evidence for hypotheses was uncertain nor what additional evidence they needed. In part, this was because the agents used very limited models of negative evidence and so could not consider events like ghosting that may provide alternative explanations for data. These weaknesses have limited the types of interactions among agents that could be supported; certain classes of solution errors have not been able to be resolved because this would have required exchanging large amounts of information. In addition, termination of network problem solving has been based on relatively simple criteria.

Previous FA/C implementations of distributed interpretation have either had only implicit representations of their goals for resolving solution uncertainty [Lesser & Corkill 1983] or else have had explicit representations of only very high level goals based on limited characterizations of uncertainty [Durfee & Lesser 1987] that did not provide detailed enough information. As a result, these systems cannot dynamically reason about the most important goals for generating a global solution and the best information to satisfy these goals. This has lead to the use of somewhat simplistic, static problem solving strategies. For example, the Partial Global Planning research [Durfee & Lesser 1987] uses heuristics like, "avoid redundant work" and "exploit predictive information." However, the appropriateness of such heuristics depends on the situation. If there is a great deal of uncertainty in overlapping solution areas then "redundant" work could be very useful. Likewise, whether predictive information should be exploited or not depends on the certainty of such information. In other words, in FA/C problem solving, strategies must be dynamically determined based on the current goals and state of problem solving. This requires that systems have good models of the state of problem solving in both the sending and receiving agents.

In this paper, we will show that agents with more sophisticated models of their evidence and their problem solving states can support the complex, dynamic interactions between agents that are necessary to fully implement the FA/C paradigm. We will do this in the context of a new distributed problem solving testbed, DRESUN, that simulates a distributed set of RESUN interpretation systems [Carver 1990] solving a DVMT-like aircraft monitoring problem. RESUN agents maintain explicit representations of the reasons why their hypotheses are uncertain and explicit representations of the state of their goals and the actions being taken to meet those goals. The RESUN architecture can support the sophisticated evidential reasoning that is crucial to the implementation of high level communication protocols that implement distributed differential diagnosis, true multi-sensor fusion, selective communication of information among nodes at different levels of detail, complex network-wide criteria for termination of problem solving, etc.

The key to achieving the necessary complex and dynamic interactions between agents is to make the solution convergence process explicit. In our approach, this has been done by giving each agent an explicit representation of the goals that must be satisfied in order to meet the criteria for termination of (global) problem solving. Termination criteria that are not satisfied or have not been verified as satisfied, are viewed as sources of uncertainty about the global correctness of local solutions. Goals representing the need to resolve these uncertainties are posted and drive the overall problem solving process. Communication between agents results from the agents taking actions to meet these goals. Because the goals are explicit and detailed, communication between agents can be very directed. That is, instead of simply exchanging information about partial solutions, agents communicate specific *evidence* that can be used to satisfy goals of resolving particular uncertainties. Another way of viewing our approach is that we have made explicit the need to enforce constraints between possibly interdependent subproblems of the agents. We recognize (possibly) interdependent subproblems and post goals to resolve uncertainty about whether the relevant partial solutions are consistent.

In the next section we present an example scenario to show the kinds of agent interactions that must occur to converge on solutions. The following two sections give a brief description of the RESUN framework and the extensions that have been necessary for FA/C distributed interpretation. In the next section we contrast our approach with related approaches in distributed problem solving. The next to the last section contains a detailed trace of the way DRESUN handles the example discussed in the earlier section. Finally, the paper concludes with a summary of the key points.

2 Agent Interactions in the FA/C Paradigm

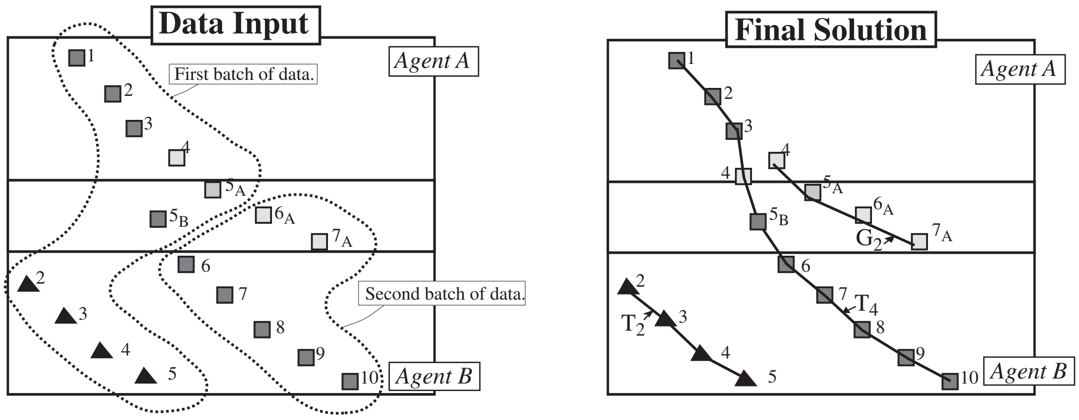

Figure 1: Example scenario data and correct interpretations.

To get an idea of the kinds of interactions that must occur between FA/C agents in order to converge on correct solutions, we will consider the aircraft monitoring scenario in Figure 1. There are two agents whose regions of interest overlap. Each agent receives data only about its region, from its own acoustic sensor. The data point symbols in Figure 1 represent the positions of groups of acoustic signals detected by the sensors. The numbers associated with the data points give the times that these signals were generated. Data points include the position of the signal source and the frequency class of the signal. Each type of aircraft produces a characteristic spectrum of acoustic frequencies. The goal of the system is to identify any aircraft that are moving through the regions of interest, determine their types, and track them through the regions.

Solution uncertainty arises from several sources, including improperly sensed signals, ghosting, and environmental noise. As a result of acoustic signal propagation and limitations in the acoustic sensors, not all acoustic signals emanating from an aircraft are properly sensed; some or even all of the frequencies in the spectrum may be missing and others may be shifted into the wrong frequency class. Ghost signals may appear as a result of environmental reflections of signals. Non-aircraft sources of acoustic signals may also be detected—these are referred to as noise. As a result of these factors, it is not possible to immediately determine whether sensor data results from an actual aircraft or whether it is the result of ghosting or environmental noise.

Resolving uncertainty about the correct explanations for data requires that the system gather evidence for and against the alternatives. This is done with models of how the different events should constrain the observed data. For example, aircraft cannot simply come to a stop, so tracks that suddenly disappear are unlikely to be from aircraft (though data may be missed by sensors). Ghost tracks, on the other hand, must have limited length. Ghost tracks must also originate from some source aircraft track, will typically be detected as incomplete frequency spectra, and may not be detected by different sensors (at different positions). Environmental noise data will not typically correspond to valid aircraft frequency spectra and cannot be correlated over time (into a track). Of course, the normal variations in the sensing of events means that definitive interpretations cannot be produced from small numbers of data points even when all the possible alternative explanations can be considered.

Because each agent has a limited view from its own sensor, individual agents cannot verify these kinds of constraints without communicating with each other. For example, both aircraft tracks and ghost tracks may continue from one agent's region into another; determining whether tracks are continuous requires communication. Likewise, the source of an agent's ghost track may be outside the agent's region. In other words, each local agent's subproblems may be interdependent with other agents' subproblems (the subproblems here are determining the correctness of interpretation hypotheses).

In the example in Figure 1, the two agents must communicate in order to converge on the correct solution and in order to produce reasonable levels of certainty in their solutions. Without any communication, agent A would incorrectly interpret its input data (for times 1 through 7) as a ghost track. This would happen because agent A's sensor has failed to detect any signals from track T_4 at times 4 and 5 (i.e., at T_4 points 4 and 5_B in the final solution of Figure 1). Were this data available to agent A, it would suggest the alternative (correct) explanation of agent A's time 1 through 3 data as being due to an actual aircraft (that produces T_4). Without any communication, agent A would also continue to be very uncertain about its ghost track explanation for the data; it would not be able to find a source for the ghost track and could not be sure that the ghost track did not continue beyond its border with agent B (since this might suggest that the data was really due to an actual aircraft). Likewise, agent B's confidence in its interpretations of its data (track T_2 and the time 5 through 10 portion of track T_4) would also be somewhat limited. For instance, while the time 5 through 10 data of T_4 may be quite good in terms of its match to likely aircraft frequency spectra, B's confidence would still be limited because of the limited time (number of points) over which it is able to track the vehicle.

This example also shows that a complete answer map could not easily be created from the agents' independent solutions; there would have to be major adjustments of some of the individual interpretations. This adjustment process requires back and forth communication between the agents rather than simply having one agent's "better" solutions override the others. Here, the portion of track T_4 constructed by agent B is not so strongly supported that it can be forced into the global solution without some corroboration from agent A. This requires that agent A use agent B's portion of track T_4 as predictive information, allowing agent A to make assumptions about its sensor having missed signals at times 4 and 5 that could complete track T_4. Agent A must also be able to produce an acceptable interpretation for the remainder of its original ghost track (the time 4 through 7 data). Once again, communication with agent B helps to confirm most of this data (times 5 through 7 in the overlapping region) as ghost data and can provide a source for that ghost track (T_2).

3 RESUN Agents

In the DRESUN testbed, individual agents are RESUN interpretation systems [Carver 1990, Carver & Lesser 1991]. Interpretation hypotheses are maintained on a blackboard database, but RESUN extends the conventional blackboard representation of hypotheses. The most important extension involves the use of symbolic statements of the *sources of uncertainty* (*SOUs*) in the evidence for the hypotheses. Symbolic SOUs are attached to hypotheses as they are created or refined. Having the symbolic SOUs makes it possible for the system to understand the reasons why hypotheses are uncertain. For example, a track hypothesis in an aircraft monitoring system may be uncertain because its supporting sensor data is incomplete or because this data might have alternative explanations (e.g., it is ghost data or it is from a different aircraft).

Control decisions are made by a script-based, incremental control planner with context-specific focusing. The hierarchical goal/plan/subgoal structure created by the control planner provides the system with an explicit representation of the system's current goals, the relationships between alternative goals, the relationships between goals and actions, and the status of the methods being used to pursue goals. Because of this, control decisions can be highly context-specific and can explicitly consider the current state of problem solving. A major innovation of the control planner is its *refocusing mechanism*. Refocusing can be used to handle decision nondeterminism and can provide the goal-directed planning mechanism with opportunistic control capabilities.

In RESUN, interpretation is viewed as an incremental process of gathering evidence to resolve particular sources of uncertainty in the interpretation hypotheses. In other words, the problem solving process iteratively considers what the sources of uncertainty are that keep the current answer from being sufficiently certain for termination and then takes actions appropriate to resolve this uncertainty. This process is repeated until the termination criteria are met. Having the symbolic SOUs allows the system to identify and use methods that can directly resolve the uncertainties. By contrast, most blackboard-based interpretation systems are limited to (indirect) *incremental hypothesize and test* methods. In particular, the SOU representation permits the use of *differential diagnosis* techniques because the possibility of alternative explanations for hypotheses and data are explicitly represented.

The overall interpretation process is driven by a high-level model of the state of problem solving, called PS-Model. PS-Model includes a statement of current interpretation “answer” in terms of believed hypotheses and symbolic statements of the sources of uncertainty (SOUs) that keep the current answer from being sufficiently believed for termination of problem solving. For example, PS-Model SOUs may denote that no evidence has been gathered for a portion of the region of interest, that there is data which has not been examined to see if it can support an answer, and that some existing potential answer hypothesis is insufficiently supported. Termination in interpretation problems requires that the system not only consider whether existing hypotheses are sufficiently proved or discounted, but must also consider whether enough of the data has been examined to be sufficiently sure that no additional answers may be found—without having to examine all of the data.

The RESUN evidential representation system also includes a scheme for numerically summarizing the symbolic SOUs. This process produces a composite characterization of the uncertainty in a hypothesis in terms of an overall belief rating and the relative uncertainty contributions of the different classes of SOUs. This summarization is used in evaluating the satisfaction of termination criteria and when reasoning about control decisions; the composite rating allows for more detailed reasoning than would be possible with a single number rating. The RESUN model of interpretation uncertainty includes the following SOU classes (that are used in the composite summary): partial evidence, possible alternative explanations, possible alternative support, alternative extensions (hypothesis versions), negative evidence, and uncertain constraints.

4 Extending RESUN for DRESUN

In order to use RESUN agents for distributed problem solving, the (single-agent) RESUN model has had to be extended. For example, DRESUN agents have to represent: global consistency termination criteria, inter-agent communication dialogues, and evidence from other agents (“external evidence”). The set of control plans of the individual agents also have to be extended to be able to respond to these additional features.

In keeping with the basic RESUN model of control being driven by the need to resolve uncertainty, verification of global consistency is driven by adding appropriate SOUs to the PS-Model. These SOUs effectively represent the uncertainty over the global consistency of an agent's local solutions. They are created when an agent recognizes that his solutions (subproblems) potentially interact with those of other agents (based on the organization of agent areas). There are three types of global consistency: solutions involving overlapping regions of interest among agents must be consistent, ”track” hypotheses that can extend into other agents' areas must be consistent, and agents must be able to find appropriate external evidence when the hypotheses require evidence which could be in other agents' areas–e.g., ghost track source (explanation) or attack scenario involving multiple aircraft over a large area

Consistency in overlapping areas is handled by adding *consistent-overlapping-model* SOUs to PS-Model. These SOUs keep track of the fact that a particular portion of the overlapping region of the PS-Model has not been checked to verify that it is consistent with the model of an overlapping external agent. Once information is obtained from the other agent, this external evidence will be integrated into the agent's hypotheses and any uncertainty due to actual inconsistency will be represented at that level.

Consistency of hypotheses that involve continuous "tracks" of supporting evidence is handled as an extension of the method that is used for judging the completeness of these tracks for single agents. When tracks cannot be extended further using an agent's own data and the extension region for the track involves another agent's area, then a *consistent-global-extension* SOU will be added to the track's model in PS-Model. Once again, when evidence is exchanged to resolve these SOUs, any resulting uncertainty due to inconsistency will be represented at the level of the corresponding track hypotheses.

Consistency of hypotheses that may require evidence from other agents' areas is handled in a manner similar to "track" extension consistency. When evidence for a hypothesis cannot be found in an agent's own region and it is possible that the evidence could be in another agent's region, negative evidence will be added to the hypothesis, but with SOUs denoting the possibility that this evidence could be gathered from another agent. These external evidence SOUs then trigger the creation of *consistent-global-evidence* SOU in PS-Model (associated with the model of the relevant hypothesis).

As we have stated above, communication between DRESUN agents does not simply involve exchanging solutions, but is directed toward the exchange of *evidence* to resolve particular uncertainties. In order to best understand how to integrate evidence from another agent, it is useful to have a context for the received information. This is provided through the concept of a dialogue. When a DRESUN agent initiates a request for evidence, it is effectively initiating a new dialogue. The control plan instance that started the communication implicitly understands the purpose of the dialogue and all further communications related to that dialogue (communications identify the dialogue they result from) are handled by that same control plan instance—rather than by some general communication handling plan.

In single-agent RESUN systems, when a hypothesis is used as evidence, all of its supporting substructure (i.e., *its* evidence) is available. When using evidence from another agent this is typically not the case because communicating all of this information would be too expensive. As a result, hypotheses supported by evidence from other agents cannot be constructed as normal RESUN hypotheses with only *support* and *explanation* evidential inferences. Instead we must add a new evidence category, *external evidence*, that allows us to directly support hypotheses with information from another agent (and we add another SOU class to the composite summary of SOUs: external-evidence-sous).

Since most evidence is uncertain when it is sent to another agent, another important aspect of dialogues is the need to update external evidence as hypotheses evolve. For example, while a track hypothesis from another agent can explain a given agent's ghost track, the other agent may itself be uncertain about the correctness of the track. As additional evidence is gathered by the other agent, it may decide that the track it sent is actually incorrect. Conversely, the initiating agent may find that its ghost track is not a ghost track after all. In either case, the agents will need to initiate a new dialogue to resolve the uncertainty over the external evidence. Recognizing the need for updating is done through the use of *external-evidence-uncertainty* SOUs that are associated with an agent's external evidence.

The integration of external evidence shows why communication of information between agents is not just a matter of exchanging information. Sometimes external evidence may be consistent with an agent's own evidence either immediately or through refinement of uncertain parameter values. In these cases, integration is relatively straightforward. However, there may also be cases that require a complex dialogue to handle—e.g., overlapping, "partially consistent" vehicle track hypotheses. In these cases there are many possible explanations for the data: the tracks might actually be due to different vehicles (they only appear to overlap due to limited sensor resolution), one track is right and the other is wrong (the non-overlapping data of the incorrect track has other explanations), each track is wrong (alternative correct tracks can be identified when all the data is analyzed), etc. In single-agent RESUN systems this uncertainty is represented by SOUs associated with the supporting substructure; all the data is available to the agent so it is possible to see that there are alternative track explanations for the data. Without direct access to the substructure, inconsistency in external evidence must be resolved via an inter-agent differential diagnosis process.

5 Relationship to Other Research

Resolving global consistency can be viewed as a form of "consensus formation" [Courand 1990]. However, in the consensus formation framework, agents start dialogues in order to eliminate conflicts they have about their joint plans. By contrast, in the DRESUN approach, agents communicate not only when conflicts emerge, but when there are any sources of global uncertainty; conflicts are just viewed as one particular reason for uncertainty. There is another distinction between DRESUN and most other approaches to cooperation that emphasize settling on appropriate plans and goals. In DRESUN, it is the current sources of uncertainty that drive control by determining what goals and plans are currently applicable.

Because of its concern with solution uncertainty the DRESUN approach is closer in some ways to systems based on belief revision. Such systems include: DTMS [Bridgeland & Huhns

1990], DATMS [Mason & Johnson 1989], and RDRMS [Doyle & Wellman 1990]. The non-monotonic DTMS employs an algorithm that guarantees local consistency for each agent and global consistency of shared information. DATMS permits inconsistency to exist among different knowledge bases. RDRMS relates belief revision to revisions of large plans, and uses a decision-theoretic model to make rational decisions about typical belief maintenance choices. RDRMS is more flexible in finding the supporting arguments or pursuing consequences and therefore it is closer to DRESUN than the other belief revision systems. Unlike these belief revision systems, the agents in DRESUN are driven to resolve the global inconsistencies as explicit cases of solution uncertainty. As a result, DRESUN agents make use of a variety of methods—e.g, differential diagnosis techniques that reason about alternative support, explanation, and external evidence. Furthermore, use of an evidential reasoning system (like that based on the SOUs) allows for hypotheses to have degrees of belief instead of just IN and OUT belief values as in typical TMSs.

6 An Example of the DRESUN Approach

In this section, we will give a brief trace of the kind of agent actions that are necessary to deal with the scenario that was discussed in an earlier section. We will indicate how global uncertainty drives overall problem solving and how methods that involve inter-agent communication are also used to resolve an agent's local uncertainty. Figure 2 shows a chronological sequence of the important decision points:

Scene 1: The agents receive a batch of data for times 1-5. Driven by their local goals of resolving uncertainty about possible interpretations in their areas, they begin identifying possible vehicle tracks and extending these tracks.

Scene 2: At this point, agent A has created a single track hypothesis, T_1. T_1 is quite uncertain due to the poor quality of its time 4 and 5 data. Agent B has created track hypotheses T_2 and T_3. T_2 is fairly certain, because there are few inconsistencies in its supporting data. T_3 is a very preliminary track hypothesis which is based on a single position. Because T_3 is so uncertain, agent B does not communicate with agent A at this point to verify the global consistency of the 5_B data (agents don't communicate about each little bit of data since it could just be noise).

At this point, agent A has two major goals: resolving its uncertainty over the correctness of T_1 (based on its local data) and resolving its uncertainty over the global consistency of T_1 because its time 5 supporting data is in the region that overlaps with agent B. These goals are represented by *uncertain-answer* and *consistent-overlapping-model* SOUs in PS-Model. Because it is waiting for additional data to pursue T_1 and because the overlapping hypothesis T_1 is uncertain, agent A decides to pursue the global consistency SOU. It does this by requesting agent B to verify the 5_A data (in T_1). In reply, B informs A that it could find no evidence to support 5_A. A records negative external evidence in T_1, reducing A's

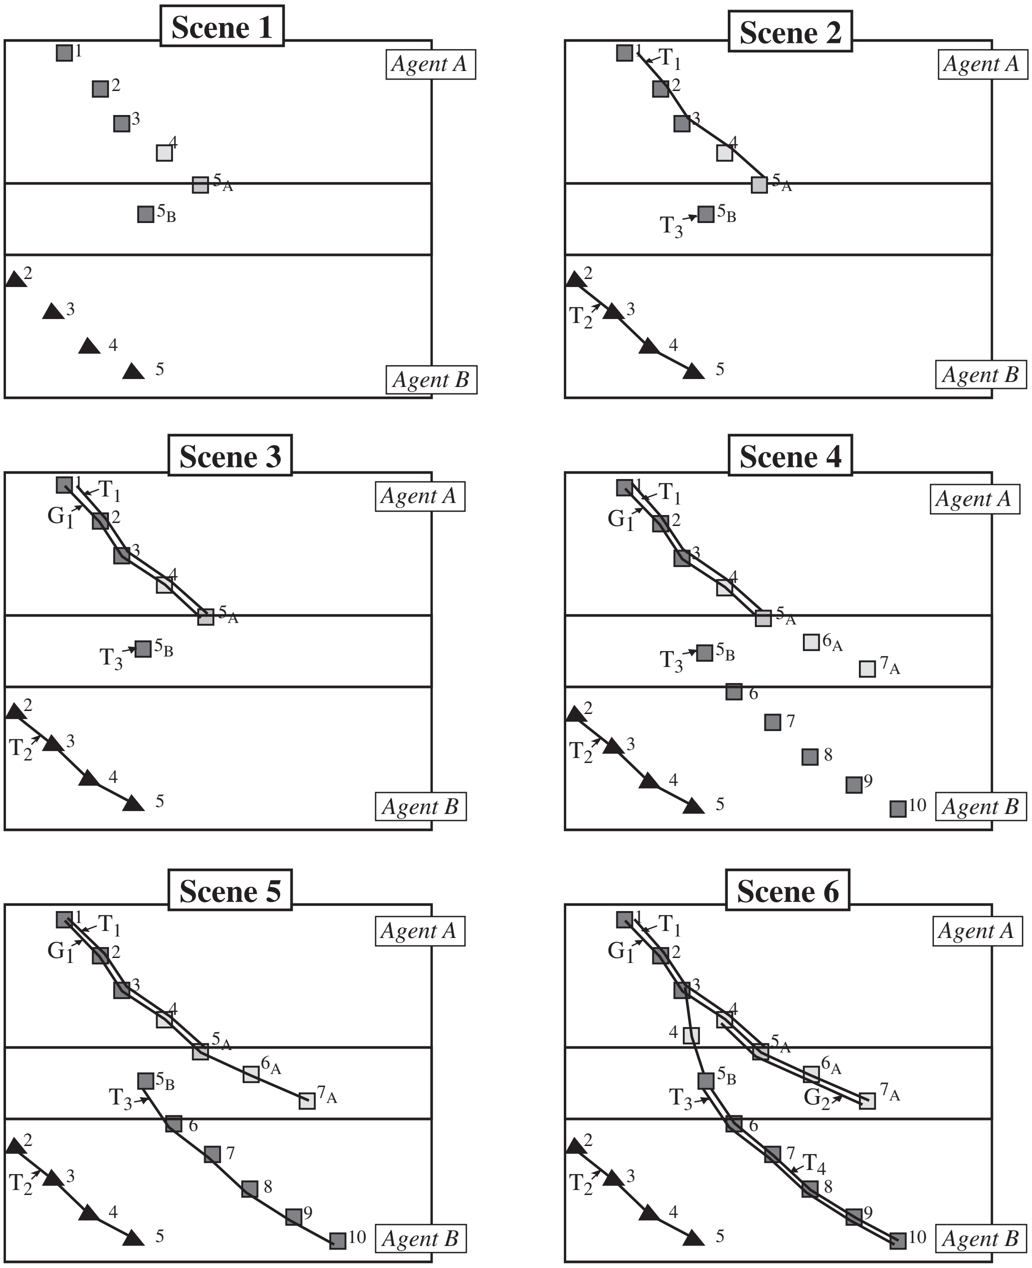

Figure 2: A chronological sequence of scenes depicting important problem-solving interactions for the example scenario.

belief in T_1. A now returns to its (local) goal of resolving uncertainty over the possibility that T_1 is an answer hypothesis. The uncertainty in T_1 as a result of its poor data and negative external evidence causes agent A to consider pursuing alternative explanations for T_1's data. Examining the symbolic SOUs for T_1's supporting evidence, A finds that the data could be due to a ghost track. Since the negative external evidence further supports this possibility, A decides to pursue it.

Scene 3: A has created ghost track G_1 to pursue as an alternative to track T_1. One source of uncertainty for G_1 is its lack of an explanation: a track hypothesis that is the source of the ghost. In order to resolve its uncertainty in G_1, agent A examines its hypotheses for a possible source, but finds none. While this generates negative explanation evidence for G_1, this evidence is weak and uncertain because it is possible for the source track to be outside of A's region of interest—in B's region. This results in the creation of a *consistent-global-evidence* SOU in PS-Model. While this global consistency SOU can cause communication, communication here actually occurs as a result of A continuing to resolve its uncertainty over G_1: agent A requests a source for G_1 from agent B in order to resolve uncertainty over the negative explanation evidence. This shows that similar sorts of communications between agents can be driven by both local and global goals. B's track T_2 is consistent with the criteria for being a source track of the ghost G_1. A uses T_2 as further evidence for G_1 by recording T_2 as a possible explanation for G_1. Note though, that agent A maintains information about the uncertainty associated with this evidence.

Scene 4: At this point, a new batch of data comes in for times 6-10. Driven by their local goals of resolving uncertainty, the agents pursue their existing hypotheses: agent A extends ghost G_1, while B extends track T_3.

Scene 5: The characteristics of the 6_A and 7_A data conform well to the model of ghosting and contribute to increased confidence in G_1 as does the fact that G_1 cannot be continued (the ghosting model is explained in an earlier section). The extension of G_1 with more data in the overlapping region results in a new *consistent-overlapping-model* SOU associated with G_1 in PS-Model. This once again causes agent A to request B to confirm G_1's support in the overlap region; confirmation failure increases confidence in G_1.

Agent B's track hypothesis T_3 has become quite well supported and is found to be complete for B's region of interest. This leads to the posting of a *consistent-global-extension* SOU in PS-Model (as well as an increased importance rating of the *consistent-overlapping-model* SOU due to the 5_B data). The *consistent-global-extension* SOU causes agent B to request agent A to look for extensions of T_3. This request initiates a dialogue between the agents that eventually results in agent A developing new interpretations for its data (see Scene 6).

Scene 6: When A first looks for extensions for T_3 it finds none, but given the level of belief in T_3, agent A is willing to assume that its sensor has missed the time 5 data in the overlapping area. However, this still fails to produce an extension and agent A reports this to B, telling it that must further resolve its uncertainty to convince A to make further assumptions. Agent B does this by looking for alternative extensions for T_3 (since the particular version of the track hypothesis, T_3, is less certain than it is that there is some correct hypothesis). B fails to find any alternative extensions of T_3 and is now able to convince A to make assumptions about missing data at both times 4 and 5. This leads to the creation of the complete track T_4 using agent A's time 1 through 3 data which was less well explained by ghost G_1 due to the characteristics of its frequency spectra.

This results in ghost G_1 becoming disbelieved, which forces agent A to pursue it further. Agent A finds that G_1 has become disbelieved because there is a more highly believed explanation for some of its supporting data (T_4). This causes agent A to look for a new explanation for the remainder of the data that was supporting G_1. It finds that this data can still be explained as ghosting and it creates the new ghost hypothesis G_2. This ghost hypothesis is strongly supported by the remaining data of G_1 due to the nature of its frequency spectra and the absence of corroborating data from agent B. Pursuing G_2 further, agent A finds that T_4 is a likely source/explanation for this ghost track.

Here we see the final state of the possible solutions that result from the combination of evidence from both agents. The solutions with acceptable certainty for termination are tracks T_2 and T_4, and ghost track G_2.

7 Conclusions and Status

The example shows that FA/C distributed problem solving can require complex interactions between agents in order to converge on correct solutions. The DRESUN framework makes it possible for agents to have directed dialogues for distributed differential diagnosis, make use of a variety of problem solving methods in response to changing situations, transmit information at different levels of detail as appropriate, and drive local and global problem solving using the notion of the global consistency of local solutions. These capabilities have not been part of previous implementations of the FA/C paradigm. The implementation of the DRESUN framework is currently undergoing testing and we expect to have detailed performance results in the near future.

8 References

Bridgeland, D. and M. Huhns. 1990. Distributed Truth Maintenance. In *Proceedings of AAAI-90*, 72–77.

Carver, N. 1990. Sophisticated Control for Interpretation: Planning to Resolve Uncertainty. Ph.D. diss. Computer and Information Science Department, University of Massachusetts.

Carver, N. and V. Lesser. 1991. A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty. *Proceedings of AAAI-91.* Forthcoming.

Courand, G. 1990. Cooperation via Consensus Formation. In *Proceedings of the 10th International Workshop on Distributed Artificial Intelligence*, Ch.10.

Doyle, J. and M. Wellman. 1990. Rational Distributed Reason Maintanance for Planning and Replanning of Large-Scale Activities (Preliminary Report). In *Proceedings on Inovative Approaches to Planning, Scheduling and Control Workshop sponsored by DARPA*, 28–36. Morgan Kaufmann.

Durfee, E. and V. Lesser. 1987. Using Partial Global Plans to Coordinate Distributed Problem Solvers. In *Proceedings of IJCAI-87*, 875–883.

Lesser, V. and D. Corkill. 1981. Functionally Accurate, Cooperative Distributed Systems. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1):81–96.

Lesser, V. and D. Corkill. 1983. The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks. *AI Magazine*, 4(3):15–33.

Lesser, V. 1991. A Retrospective View of FA/C Distributed Problem Solving. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed AI.* Forthcoming.

Mason, C. and R. Johnson. 1989. DATMS: A Framework for Distributed Assumption Based Reasoning. In Les Gasser and Michael Huhns, eds. *Distributed Artificial Intelligence*, Vol. II, 293–317. Morgan Kaufmann.