The DRESUN Testbed for Research in FA/C Distributed Situation Assessment: Extensions to the Model of External Evidence*

Norman Carver
Department of Computer Science
Southern Illinois University at Carbondale
Carbondale, IL 62901
(carver@cs.siu.edu)

Victor Lesser
Department of Computer Science
University of Massachusetts
Amherst, Massachusetts 01003
(lesser@cs.umass.edu)

Abstract

This paper reports on extensions that have been made to the DRESUN testbed for research on distributed situation assessment (DSA). These extensions involve issues that have arisen in modeling the beliefs of other agents when dealing with inter-agent communication of incomplete and conflicting evidence, and evidence at multiple levels of abstraction. The extensions support highly directed exchanges of evidence among agents because they better represent the uncertainties that occur when DRESUN agents exchange incomplete and conflicting information. This is important in FA/C systems because agents must share results in order to satisfy their local goals as well as the overall system goals. Thus, sharing must be done efficiently for an FA/C approach to be effective. These issues will arise in any distributed problem solving application involving interacting subproblems, when agents must function without complete and up-to-date information.

Introduction

The *functionally accurate, cooperative* (FA/C) paradigm for distributed problem solving (Lesser & Corkill 1981; Lesser 1991) was proposed for applications in which tasks are naturally distributed but in which the distributed subproblems are not independently solvable. In the FA/C approach, agents produce tentative, partial results based on local information and then exchange these results with the other agents. The constraints that exist among the agents' subproblems are then exploited to resolve the local uncertainties and global inconsistencies that occur due to the lack of accurate, complete, and up-to-date local information.

In (Carver & Lesser 1991b) we described the capabilities of the initial implementation of DRESUN, a testbed for research on distributed situation assessment (DSA)[1] using an FA/C approach. DRESUN was developed to explore the implications of having agents with more sophisticated evidential representations and control capabilities than the agents that were used in earlier research with the Distributed Vehicle Monitoring Testbed (DVMT) (Lesser & Corkill 1983; Durfee & Lesser 1991). Because of agent limitations, that research did not adequately address several important issues that arise when sharing incomplete and inconsistent information among DSA agents. Furthermore, overall agent activities were not driven by an explicit need to produce local solutions that were globally consistent (let alone globally optimal).

This paper reports on extensions to the initial DRESUN testbed, related to modeling the beliefs/evidence of other agents. The basic DRESUN architecture provides a good basis for an FA/C approach because inter-agent subproblem interactions are explicitly represented and used to drive problem solving. However, our experiments showed that extensions to the model of *external evidence*[2] were necessary to make the most effective use of inter-agent communication of incomplete and conflicting evidence, and evidence at multiple levels of abstraction. The focus of these extensions has been on representing the uncertainties that occur when DRESUN agents exchange such information, determining precisely what information is needed to resolve global inconsistencies, and providing the ability to reformulate hypotheses to efficiently pursue alternative interpretations.

These issues are very important in FA/C systems because agents must share results in order to satisfy their local goals as well as the overall system goals.

*This work was supported by the Department of the Navy, Office of the Chief of Naval Research, under contract N00014-92-J-1450. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

[1] Situation assessment involves the fusion of sensor data, intelligence information, and so forth, and the interpretation of this information to produce a high-level description of the situation in the environment.

[2] In a DSA framework, results from another agent necessarily produces evidence as an integral part of the process of using the information—e.g., checking whether the information is consistent or inconsistent with the local interpretations. Evidence based on information shared by another agent is referred to as external evidence.

To resolve *data uncertainties*, an agent must be able to evaluate whether other agents' results are consistent or inconsistent with its own results, and integrate these other agents' results to revise and extend its local partial results (Lesser & Corkill 1981). A key assumption of the FA/C approach is that agents can do this without "excessive" communication among the agents.[3] However, DSA tasks can present several sources of difficulty for efficient results sharing: agents' local data may lead to solutions that are globally inconsistent; agent beliefs/results are uncertain and imprecise; results (e.g., interpretations) are complex structures; and beliefs are constantly being revised due to new data and further processing.

The scenario in Figure 1 is an example of a distributed vehicle monitoring situation in which local solutions are inconsistent and extended agent interactions are necessary to resolve the inconsistency. We will use this example to introduce some key issues in results sharing, and we will return to it in more detail in the Resolving Global Inconsistency Section, to explore the representation of external evidence in DRESUN.

In the example, processing of only their own local data would cause agent A and agent B to form the track hypotheses T_a and T_b, respectively. Because the tracks extend through an area of overlapping interest, the agents can recognize that they must communicate to verify the global consistency of their local interpretations. These tracks are inconsistent since they imply that either a single vehicle is in different places at the same time or else two vehicles are in the same place at the same time. One important thing to note here is that while exchanging abstract results (the track hypotheses without their supporting evidential structures) allows the inconsistency to be detected, this level of information is not sufficient to allow the inconsistency to be resolved (i.e., there remains uncertainty about the correct global interpretation). This is because neither T_a nor T_b is significantly more likely than the other since each includes some good quality data and some poor quality data. (Even when partial results are "consistent," uncertainties may result when only abstractions are exchanged—see the Representing External Evidence Section.)

Resolving the inconsistency in favor of the most likely global interpretation requires an understanding of the level of belief provided by data supporting the different portions of each track. In other words, more detailed information about the interpretation hypotheses is required. One way to insure that agents have all the necessary information would be to always communicate the complete evidential structure associated with the hypotheses. However, because this structure can be very complex, communication and processing limitations will typically make it impractical to fully communicate this information.[4] Furthermore, complete communication is often not necessary. For instance, we will see that in this example each agent does not have to have complete, detailed information about each other's raw data in order to resolve the inconsistency.

Thus, what is needed is a system with the flexibility to request or respond with information at different levels of detail—based on the dynamic problem-solving requirements—as part of an incremental process of resolving inconsistency. This requires the ability to integrate incomplete results information, represent the resulting uncertainty, and use this representation to drive further actions. The DRESUN architecture provides these capabilities.

Because of its representation of inconsistency as a source of uncertainty and its emphasis on directed interactions to resolve inconsistency, DRESUN differs from most DAI work dealing with the global consistency of local agent beliefs. This work (e.g., (Courand 1990; Huhns & Bridgeland 1991)) has largely focused on methods for automatically maintaining (some particular level of) consistency and has used justification-based representations of belief (e.g., TMSs). DRESUN does not automatically enforce complete consistency because this can be very expensive both in terms of communication and computation, and it is usually not necessary. DRESUN uses an evidential (partial beliefs) representation rather than a justification-based representation of its beliefs because virtually all evidence and hypotheses in a DSA system are uncertain.

The next section reviews the DRESUN architecture, in which agent interactions are driven by the need to resolve uncertainty about the global consistency of local solutions. This is followed by a section that examines some of the issues that have arisen in representating external evidence in DRESUN. The example introduced in this section is then explored in more detail, and the paper concludes with a brief summary and a discussion

[3]Another current research direction of ours is formal analysis of the FA/C model in terms of the domain characteristics necessary for the approach to be effective and the quality of solutions that can be produced—see (Carver & Lesser 1994; Carver 1995). For instance, effective FA/C problem solving requires one or more of the following: 1.) only a subset of each agent's subproblems interact with those of other agents; 2.) it can be determined what local data/abstractions are relevant to which other agents; 3.) data abstractions (i.e., the tentative, partial results) can substitute for the raw data in determining the global consistency and likelihood of local solutions. If these conditions are not satisfied then a centralized approach may have better performance (though a distributed approach may still be preferred due to factors like tighter coupling of the processor and sensor, or increased reliability and graceful performance degradation).

[4]The example shown here is simplified to allow us to focus on our main points. It shows only a small fraction of the data that would be faced by most real-world DSA systems, and it does not show the numerous alternative interpretation hypotheses that would be interrelated with any solution hypotheses.

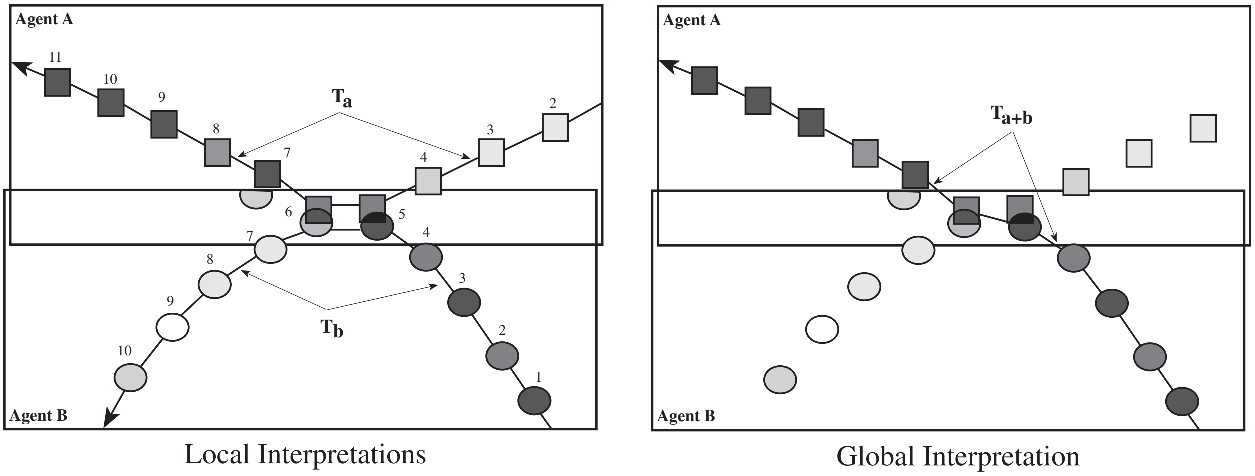

Figure 1: An example of inconsistent local interpretations.
The application is vehicle monitoring. Agent A and Agent B receive data only from their own individual sensors, whose coverage regions overlap. Agent A's data is represented by squares and agent B's by circles, with positions as indicated at the times denoted by the associated numbers. The grey density of the data points corresponds to the relative "quality" of the data—i.e., the a priori likelihood that the data would have resulted from a real vehicle. and the spectral content of the acoustic signals. "Empty" points denote data whose existence has been assumed by the agents. Based on its own data, each agent would form the local interpretations shown: agent A would hypothesize vehicle track T_a and agent would hypothesize vehicle track T_b. T_a covers agent A's data from times 2 through 11, and T_b covers agent B's data from times 1 through 10. The preferred global interpretation—given a complete view of the data from both agent A and agent B—is T_{a+b} because it covers more high quality data than either of the local tracks (the remaining uninterpreted data is due to ghosting phenomena and may or may not be explicitly interpreted depending on the termination criteria of the system). T_{a+b} covers agent B's data from times 1 through 6 and agent A's data from times 5 through 11 (it covers both agents' consistent data at times 5 and 6).

of current research issues.

DRESUN

DRESUN agents are RESUN interpretation systems (Carver & Lesser 1991a; Carver & Lesser 1993). One of the key ideas in RESUN is the use of symbolic *source of uncertainty* statements (SOUs) in the evidence for the interpretation hypotheses. The SOUs allow agents to understand the reasons why their hypotheses are uncertain and why their termination criteria remain unmet. RESUN also supports the use of satisficing control and heuristically controlled, approximate evaluation of hypothesis belief to deal with the computational complexity of DSA problems (Carver & Lesser 1994).

In an FA/C system, there must be some mechanism to drive the exchange of results among the agents so that incorrect and inconsistent local solutions can be detected and dealt with. Ideally, this would be accomplished with a mechanism that allows agents to understand where there are constraints among their subproblems, so that information interchange could be highly directed. DRESUN provides just this capability for DSA applications. DRESUN agents create *global consistency* SOUs whenever it is determined that a local hypothesis can obtain evidence from another agent—i.e., whenever a subproblem interaction (constraint) is detected. These SOUs are viewed as sources of uncertainty about the correctness of an agent's local solution because they represent unresolved questions about the global consistency of the solution. Examples of situations involving each of the global consistency SOUs are shown in Figure 2.

DRESUN's global consistency SOUs make explicit the possible interrelationships between agents' local subproblems, and provide an integrated view (in conjunction with the standard RESUN SOUs) of both the local and global problem-solving goals, which drive agent control decisions. Thus in DRESUN, agents exchange results based on the goal of insuring the global consistency of their local solutions. So far, though, we have not discussed what "resolving" a global SOU means. Resolution of a global SOU involves exchanging information among the associated agents so as to effectively propagate evidence between their hypothesis-evidence networks.[5] An example of the resolution of a global SOU is shown in Figure 3. Resolution of global SOUs is analogous to (intra-agent) evidence propagation, and as with evidence propagation there are a range of strategies that may be used to determine which global SOUs to pursue and how

[5]While there are significant differences between these networks and Bayesian or belief nets, for our purpose here the reader can consider them to be similar.

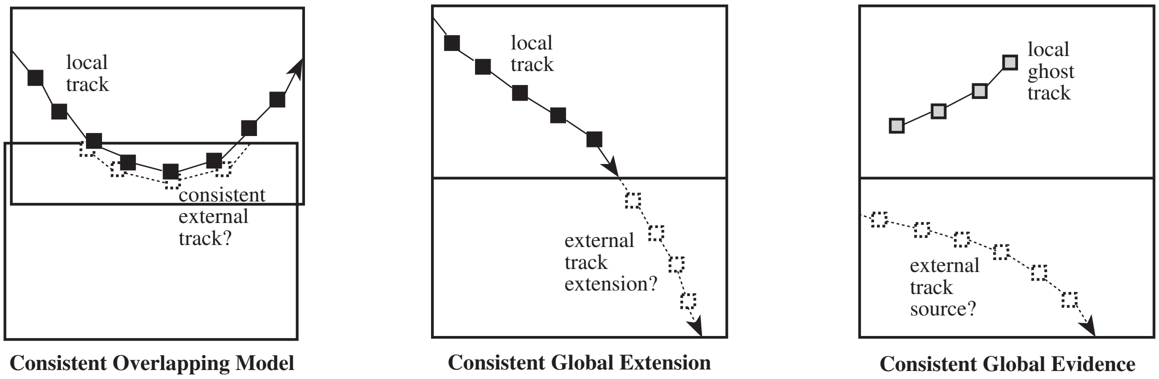

Figure 2: Examples of the global consistency SOUs for vehicle monitoring.
There are three types of global consistency interactions in sensor interpretation problems: interpretations in regions of overlapping interest among agents must be consistent, "continuous" hypotheses (e.g., vehicle tracks) that would extend into other agents' areas must have consistent external extensions, and hypotheses that require evidence that could be in another agent's area (e.g., the source for a ghost track hypothesis) must have appropriate external evidence. Instances of these situations can be detected given the domain model and knowledge of the organization of agent interest areas. DRESUN uses three global consistency SOUs to denote instances of these global interactions: *consistent-overlapping-model*, *consistent-global-extension*, and *consistent-global-evidence*.

completely to propagate their effects. We will not explore the issues involved in these choices here, but see (Carver & Lesser 1994).

Representing External Evidence

The DRESUN architecture provides the basis for effective FA/C problem solving because it explicitly represents the inter-agent subproblems interactions and uses this information to drive the exchange of partial results among the agents. However, our initial experimentation with DRESUN found that there were some restrictions on the coordination strategies that could be supported because of an inability to represent the uncertainties that arise when using incomplete results information and results information at multiple levels of abstraction. This section will try to make clear the differences between representing local evidence and external evidence, by focusing on the problems that arise when evaluating the effect of (both consistent and inconsistent) incomplete external evidence.

First, we must make the concept of global consistency/inconsistency more precise. In interpretation problems, data and hypotheses[6] are consistent if they can be *merged* into a single valid interpretation hypothesis or if they relate only to completely independent top-level hypotheses. For example, two vehicle track hypotheses that overlap in time are consistent if their (imprecise) vehicle type parameters are consistent, if their (imprecise) positions intersect at the overlapping times, and if their positions for the non-overlapping times are consistent with vehicle movement constraints. Consistency checking is straightforward in DRESUN.[7]

When consistent local hypotheses are merged, agents have all the evidence necessary to construct a *complete* new hypothesis. For example, in the consistent local evidence example in Figure 4, the supporting data of T_1 and T_2 can be used to create a new hypothesis T_3. Now, consider the case in which T_2 is an external hypothesis, and the local agent does not have (immediate access to) any of T_2's supporting evidence. In this case, the local agent can also create a new hypothesis T_3, which has the same attributes (e.g., positions and vehicle ID) as the T_3 created in the local evidence case. However, without access to the evidence for the external T_2, the belief in this T_3 cannot be properly evaluated. Evaluating the belief in T_3 requires knowledge of the quality of the data for each of supporting vehicle positions, but all the local agent has access to is the overall belief in T_2—which depends on the quality of the data from the overlapping positions as well as the positions that extended T_1. While the belief in T_3 might be estimated from this evidence (assuming, for instance, that T_2's overlap data is of about the same quality as T_1's), the resulting belief rating will be uncertain.

One of the characteristics that makes sensor interpretation difficult is that inconsistency (i.e., alternative interpretations of the data) leads to complex evidential/belief interrelationships among hypotheses. In the

[6]When we talk about "hypotheses" in this paper, we are really referring to what are called hypothesis *extensions* in RESUN/DRESUN (Carver & Lesser 1991a).

[7]DRESUN requires that hypotheses have sufficient attributes to be able to judge the consistency of new evidence without having to have access to all the existing evidence for the hypotheses. Thus, the consistency of two track hypotheses can be judged without requiring access to their supporting evidence.

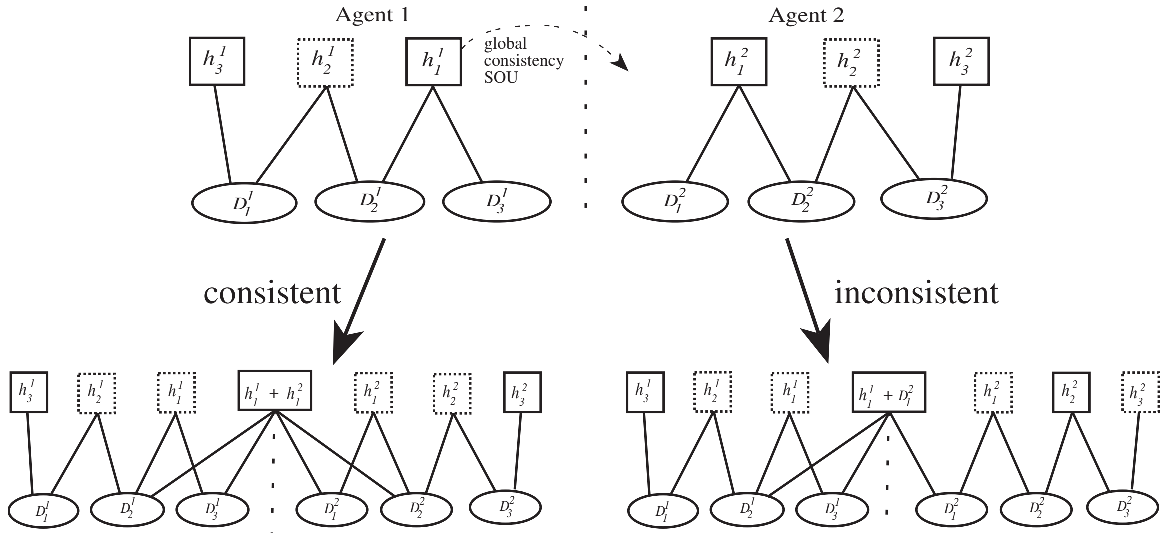

Figure 3: An example of the resolution of a global consistency SOU.
When there is a consistent explanation in the external agent, resolution of the global SOU associated with h_1^1 results in the creation of a merged hypothesis as a new alternative explanation in each agent. When the local hypothesis is inconsistent with hypotheses in the external agent, new alternatives may be created (as shown here). When the local hypothesis is inconsistent with the data in the external agent, new evidential links are created to represent the contradictory evidence.

case of only local evidence, these relationships can at least be properly evaluated. For example, in the inconsistent local evidence example in Figure 4, suppose that T_1 and $T_{2'}$ overlap at V_3 and V_4, but are inconsistent. The inconsistency is recognized because T_1 and $T_{2'}$ are alternative explanations for the shared V_3 and V_4 support, which allows the negative evidential relationship between T_1 and $T_{2'}$ to be evaluated properly.

Now, consider the case in which T_1 and $T_{2'}$ are inconsistent, but $T_{2'}$ is an external hypothesis. It is still straightforward to detect the inconsistency. However, because the local agent does not have any of $T_{2'}$'s supporting evidence, this inconsistency can be represented only as negative evidence for T_1—which again makes it impossible to precisely evaluate the belief in T_1. First, the effect that alternative interpretations have on each other's belief depends on the relative belief of the shared and non-shared portions. For instance, if the belief in $T_{2'}$ is largely due to the quality of the overlap data, then $T_{2'}$ does not represent strong belief against T_1. Second, evidential propagation now does not automatically reflect correct beliefs if there are other interrelated hypotheses. Suppose, for example, that there are additional local hypotheses that are inconsistent with T_1 (because they are alternative explanations for some of T_1's support). These hypotheses may also be inconsistent with the external $T_{2'}$, or they may be consistent with it. Unless these relationships are explicitly examined, this will lead to additional uncertainty in the belief in T_1.

This brief example shows that communication of abstract, incomplete hypothesis information can lead to uncertainty about the effect of external evidence on local hypotheses. Communicating incomplete information can still be useful, however. In some situations the uncertainty may not be critical to resolve, and if it is, it may be able to be used to guide further communcation. We have addressed this and a number of other representation issues in our extensions to DRESUN. Several of these extensions will be described in the example in the next section. As part of our extension of the model of external evidence, we have given DRESUN agents the ability to:

- link multiple views of a hypothesis based on external evidence at different levels of abstraction;
- link multiple hypothesis extensions that are being used in alternative local interpretations (e.g., different portions of a single external vehicle track hypothesis);
- locally create alternative versions of external hypotheses based on incomplete information and represent the resulting uncertainties;
- reformulate hypotheses for more efficient exploration and representation of alternatives;
- communicate back results of integrating information that was sent by another agent;
- avoid circular reasoning when exchanging evidence among agents;
- identify when shared information should be updated.

Resolving Global Inconsistency: An Example

In this section, we will return to the example of Figure 1. Our purpose will be to show how DRESUN's

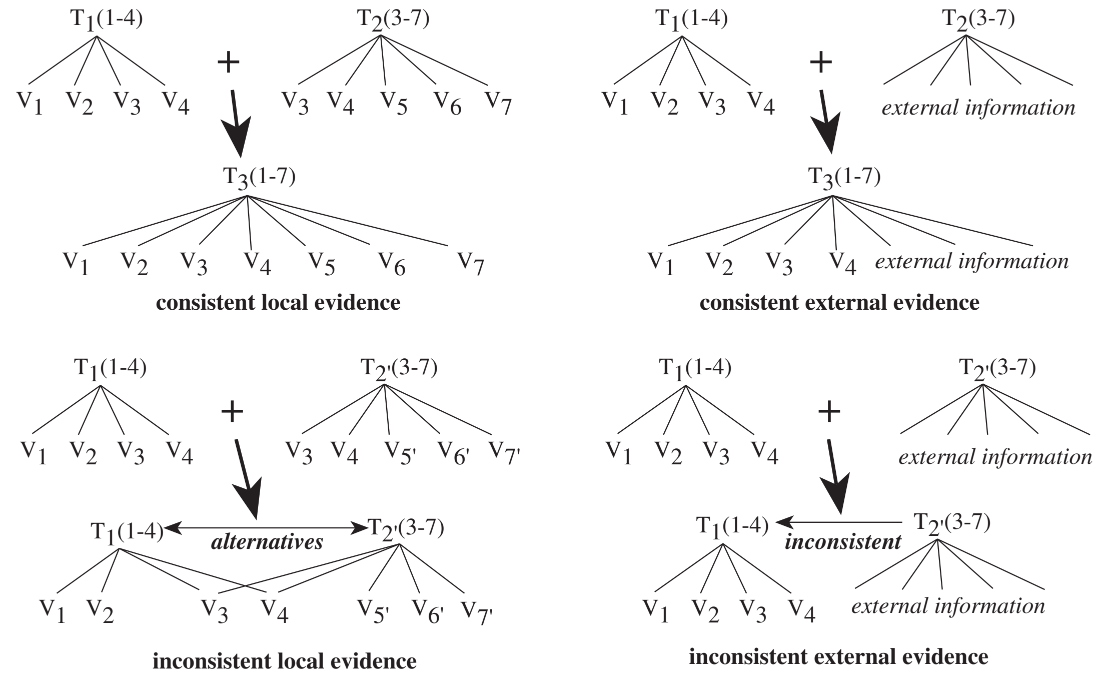

Figure 4: Examples of the differences in representing local and external evidence.

extended representation of external evidence provides the flexibility to be very efficient about the information that must be communicated among agents to resolve global inconsistencies. For the purpose of our presentation, we will assume that agent A and agent B have already formed their local track hypotheses (T_a and T_b) that extend through the overlap region. This results in *consistent-overlapping-model* SOUs being posted in each agent's PS-Model. We also will assume that these SOUs are not pursued until some level of confidence in the track hypotheses is reached (based on the local evidence), and that agent A is the first to communicate about its SOU.

When agent A starts initiates a dialog (with agent B) to resolve its "overlap" SOU, there are two options depending on whether agent A thinks the bulk of the processing to check consistency should be done by itself or by agent B: it could request agent B to send its best interpretations that cover the overlap region and then check consistency itself or it could send track T_a to agent B and let that agent check consistency. Likewise, if agent A chooses to send T_a it has several options in terms of the amount of detail it sends about T_a, or if agent B is requested to send back its interpretations it has several representation options. Here, we make the assumption that agent A will handle consistency checking and that potential solutions will initially be communicated at their most abstract level: sending only the attributes and degree of belief in the most likely top-level, track hypotheses.

Given these decisions, agent A requests that agent B send it any relevant potential solutions and agent B responds with track T_b. Agent A finds that track T_b is inconsistent with its own track T_a since the tracks overlap but cannot be merged. Because T_b is inconsistent with T_a, *negative external evidence* is created for T_a. This is the second stage of agent A's representation shown in Figure 5. The creation of this negative external evidence will cause a *global-inconsistency* SOU to be added to agent A's PS-Model. Whether or not this "inconsistency" SOU results in further communication or other processing depends on several factors, including: the original belief in T_a, the uncertainty about the magnitude of the (negative) effect that T_b has on T_a due to incomplete information about the external hypothesis (as described in the Representing External Evidence Section), the ability of agent A to pursue other sources of uncertainty in T_a (to locally resolve the uncertainty in T_a), the general classes of uncertainty affecting other hypotheses, the global consistency termination criteria, and so on.

Assuming that the agent chooses to pursue the "inconsistency" SOU, it first identifies plans that are relevant to resolving the SOU. One plan that we have developed for resolving a *global-inconsistency* SOU is applicable only when the inconsistency involves track hypotheses that are *partially consistent*. This plan attempts to construct a new extension of the local track that is consistent with a portion of the external evidence. Exactly which of the possible alternatives it initially chooses to create depends on the information it has about the relative credibility of the various por-

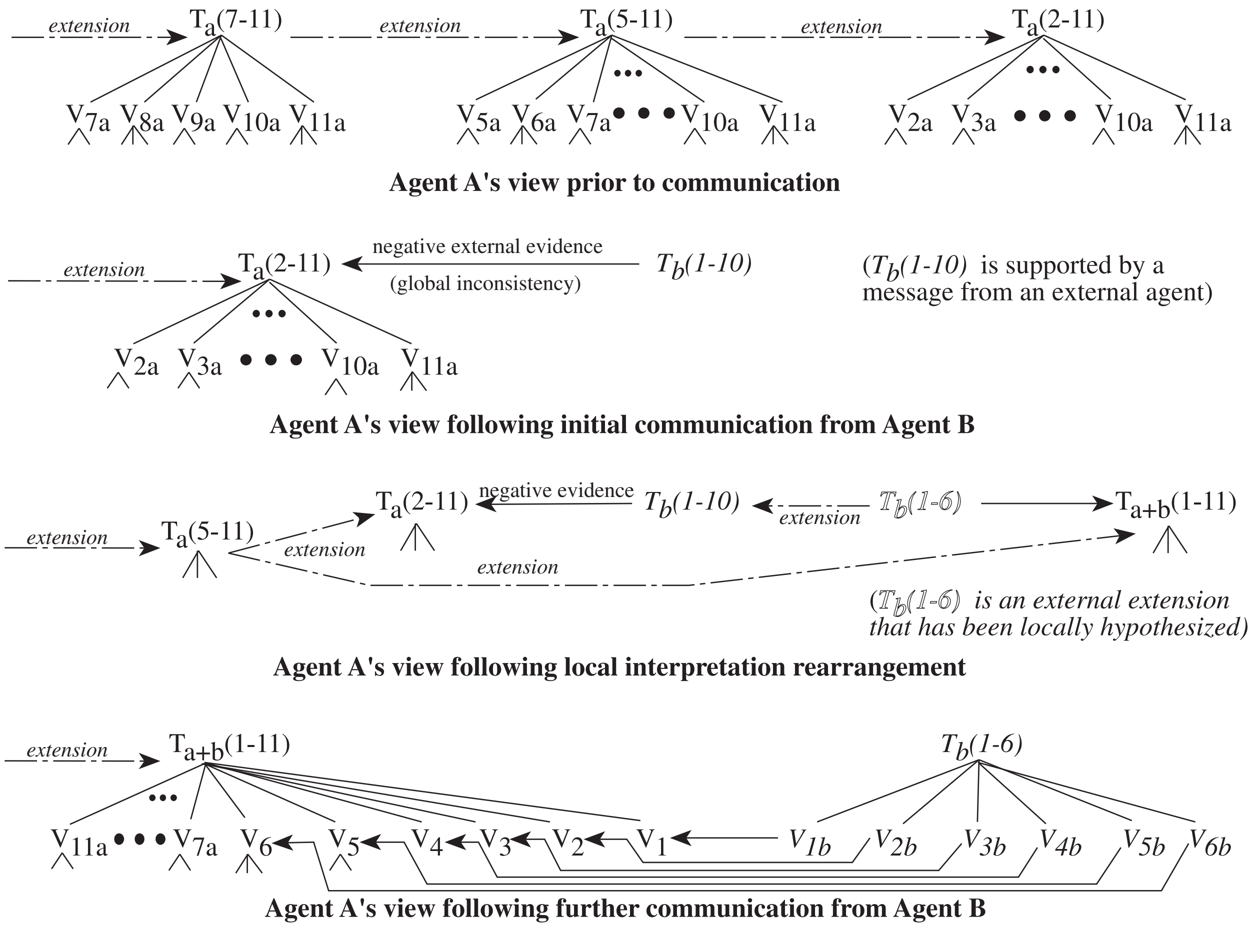

Figure 5: Agent A's evolving representation of evidence for the example.

tions of the inconsistent tracks. Here, agent A knows about the credibility of portions of only its own track T_a: the support from times 5 through 11 is quite strong and that from times 2 through 4 is weak (see Figure 1). Starting with the consistent portion of T_a at times 5 and 6, and the better supported portion from times 7 through 11, agent A decides to create a new track extension using local evidence from times 5 through 11.

In the third stage of Figure 5, agent A has created track T_{a+b} based on $T_a(5-11)$ (one of the intermediate extensions of T_a) and $T_b(1-6)$. $T_b(1-6)$ is an intermediate extension of the external track T_b that has been *locally hypothesized* by agent A—i.e., agent A has created this extension without knowing whether agent B has a representation of this version of T_b and without knowing the degree of belief in this portion of T_b.[8] Because agent A lacks both the supporting evidence and the overall belief rating for $T_b(1-6)$, there is considerable uncertainty about the belief rating for T_{a+b} (remember, it has only the overall rating for $T_b(1-10)$). The reasons for its uncertainty are represented by an SOU that is posted with T_{a+b}.

Assuming that agent A decides to pursue T_{a+b} further, since it is a credible globally consistent solution, this SOU drives the selection of a plan that requests agent B to communicate the support belief summaries for $T_b(1-6)$.[9] When the requested information is received, it is integrated into agent A's incomplete representations of $T_b(1-6)$ and T_{a+b}. The result of this process is shown in the final stage of Figure 5. With this level of information, agent A can evaluate the likelihood of T_{a+b} (given the evidence gathered so far by itself and agent B). Depending on the results of this evaluation and the termination criteria, agent A may then consider T_{a+b} to be a (likely) solution or to not be a solution, or it may need to try to gather additional

[8] Agent B may or may not have created this particular extension depending on how it gathered the evidence to construct $T_b(1-10)$. Agent B may have to reformulate its view of T_b in order to create $T_b(1-6)$.

[9] This information is provided in our application by the belief ratings at the vehicle (position) level and in the *uncertain-support* SOUs at the track level.

evidence to resolve the remaining uncertainty.

This example shows how DRESUN agents can carry on dialogs in order to resolve global inconsistencies. It also shows that these dialogs can be directed, using information at appropriate levels of detail, in order to limit the amount of information that must be communicated (and integrated). In this example, agent A does not have to have complete knowledge of T_b's supporting evidential structure or its raw data. All that is needed is information about the "quality" of the data sets supporting T_b. In fact, because agent A was able to to construct alternative hypotheses based on its local data and an incomplete view of T_b, it was able to limit the information it required to just a portion of T_b's support. The flexibility to do this sort of local processing is possible because DRESUN agents represent the uncertainty that results from the use of incomplete external evidence.

Conclusions

In this paper, we have discussed some of the issues that can arise in a DSA system when sharing incomplete or inconsistent information, or information at different levels of detail. These are important issues since the sharing of partial results is a critical element of an FA/C approach to distributed problem solving. We have also shown how the agent architecture of the DRESUN testbed has been extended to give the agents the flexibility to communicate information in a very directed manner. These extensions have focused on representing the uncertainties that occur when DRESUN agents exchange incomplete information, determining precisely what information is needed to resolve global inconsistencies, and reformulating hypotheses to more efficiently pursue alternative interpretations.

Because DRESUN supports a range of methods for resolving interpretation uncertainty and global inconsistency, coordination strategies must consider a variety of questions about whether/when/how to pursue interpretations and SOUs (the example in the previous section mentioned a number of options faced by the agents). We are pursuing both analytical and experimental approaches to determine appropriate coordination strategies (Decker & Lesser 1995), are analyzing the quality of solutions that can be produced by FA/C-based DSA systems (Carver & Lesser 1994), and are developing methods for analyzing the inherent complexity of interpretation scenarios (Whitehair & Lesser 1993). Since it is difficult to evaluate a framework independently of the strategies that are encoded within it, the development of suitable coordination strategies is a major focus of our current research.

References

Norman Carver and Victor Lesser, "A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty," *Proceedings of AAAI-91*, 724–731, 1991.

Norman Carver, Zarko Cvetanovic, and Victor Lesser, "Sophisticated Cooperation in FA/C Distributed Problem Solving Systems," *Proceedings of AAAI-91*, 191–198, 1991.

Norman Carver and Victor Lesser, "A Planner for the Control of Problem Solving Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Planning, Scheduling and Control, vol. 23, no. 6, 1519–1536, 1993.

Norman Carver and Victor Lesser, "A First Step Toward the Formal Analysis of Solution Quality in FA/C Distributed Interpretation Systems," *Proceedings of the 13th International Workshop on Distributed Artificial Intelligence*, July, 1994.

Norman Carver, "Examining Some Assumptions of the FA/C Distributed Problem-Solving Paradigm," *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Society Conference*, April, 1995.

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

Keith Decker, and Victor Lesser, "Designing a Family of Coordination Algorithms," *Proceedings of the International Conference on Multiagent Systems*, June, 1995.

Edmund Durfee and Victor Lesser, "Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation," *IEEE Transactions on Systems, Man, and Cybernetics* vol. 21, no. 5, 1167–1183, 1991.

Michael Huhns and David Bridgeland, "Multiagent Truth Maintenance," *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Distributed Artificial Intelligence, vol. 21, no. 6, 1437–1445, 1991.

Victor Lesser and Daniel Corkill, "Functionally Accurate, Cooperative Distributed Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 1, 81–96, 1981.

Victor Lesser and Daniel Corkill, "The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks," *AI Magazine*, vol. 4, no. 3, 15–33, 1983 (also in *Blackboard Systems*, Robert Engelmore and Tony Morgan, editors, Addison-Wesley, 1988).

Victor Lesser, "A Retrospective View of FA/C Distributed Problem Solving," *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Distributed Artificial Intelligence, vol. 21, no. 6, 1347–1362, 1991.

Robert Whitehair and Victor Lesser, *A Framework for the Analysis of Sophisticated Control*, Technical Report 93-53, Computer Science Department, University of Massachusetts, 1993.