

Distributed Sensor Interpretation: Modeling Agent Interpretations in DRESUN*

Norman Carver, Victor Lesser, and Qiegang Long
(carver, lesser, or qlong @cs.umass.edu)

CMPSCI Technical Report 93-75
September 1993

Abstract

The DRESUN testbed for research on distributed situation assessment (DSA) was developed to explore the implications of having agents with more sophisticated evidential representations and control capabilities than the agents that were used in our earlier research with the Distributed Vehicle Monitoring Testbed (DVMT). In DRESUN, communication among the agents is driven by the goal of determining the global consistency of local agent solutions, with unresolved global consistency questions viewed as sources of uncertainty about the correctness of local agent solutions. This paper reports on 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. Experimentation showed that extensions to the DRESUN model of external evidence were necessary to better represent the uncertainties that occur when DRESUN agents exchange such information. These are important issues since DSA agents typically must exchange much information in order to meet their goals. Furthermore, because DSA agents require evidential representations, this work is different from DAI work that has used justification-based representations of belief and has focused on methods for automatically maintaining (some level of) global consistency. Initial experimentation has been done with a variety of simulated distributed aircraft monitoring scenarios involving local solutions that are globally inconsistent and local solution uncertainty that can be resolved only through agent interactions. These experiments suggest that DRESUN agents have the flexibility to support the complex communication protocols and highly directed information exchanges that we believe are necessary to resolve global inconsistencies in DSA applications.

*This paper is a revised and extended version of the paper, "Resolving Global Inconsistency in Distributed Sensor Interpretation: Modeling Agent Interpretations in DRESUN" that appeared in the *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, May, 1993. 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 Introduction

In [Carver & Lesser 1991b] we described the capabilities of the initial implementation of DRESUN, a testbed for research on distributed situation assessment (DSA). 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 our earlier research with the Distributed Vehicle Monitoring Testbed (DVMT) (e.g., [Lesser & Corkill 1983, Durfee & Lesser 1991]). Because of the agent limitations, that research did not adequately address several important issues that may arise when sharing information among DSA agents, including: representing incomplete and inconsistent information from other agents, and determining what information is needed by other agents to resolve local uncertainties and global inconsistencies. Furthermore, overall agent activities were not driven by an explicit need to produce local solutions that were globally consistent.

This paper reports on issues related to modeling the beliefs of other agents that have arisen in the initial experimentation with DRESUN. These experiments showed that extensions to the model of external evidence were necessary to effectively utilize inter-agent communication of incomplete and conflicting evidence, and evidence at multiple levels of abstraction. The focus of these extensions has been: representing the uncertainties that occur when DRESUN agents exchange such information and providing the ability to reformulate hypotheses to efficiently pursue alternative interpretations. The extensions enhance the flexibility of the agents by improving their ability to evaluate the current state of agent beliefs, make better (local) use of incomplete information from another agent, and determine precisely what information is needed to resolve global inconsistencies.

Modeling issues are important because DSA agents typically must share information in order to satisfy their local goals as well as the overall system goals—since agent subproblems are interdependent [Lesser & Corkill 1981]. DSA tasks can present several sources of difficulty for information sharing: agents' local evidence may lead to solutions that are globally inconsistent; agent beliefs (interpretations of local data) are uncertain and imprecise; interpretations are complex structures; and beliefs are constantly being revised due to new data and further processing. When an agent shares information about its interpretations with another agent, that information necessarily produces evidence in a sensor interpretation framework 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. Interpretation evidence based on information shared by another agent is referred to as *external evidence*.

The scenario in Figure 1 is an example of a situation in which local solutions are inconsistent, and extended agent interactions are necessary to resolve the inconsistency. We will use this example to introduce the complexities of communicating information and representing external evidence

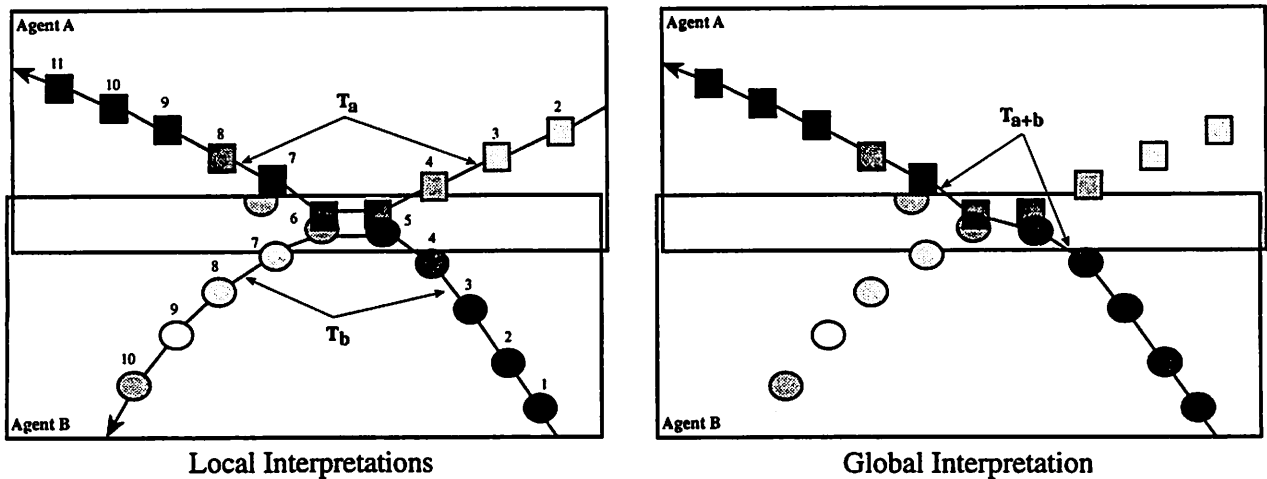


Figure 1: An example of inconsistent local interpretations.

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. "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 B 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. 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. This inconsistency cannot be immediately resolved because neither T_a nor T_b is significantly more likely than the other (each includes some good quality data and some poor quality data). 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).

in DRESUN, and we will return to it in more detail later. In the example, their own local data will cause agent A and agent B to form track hypotheses (T_a and T_b , respectively) that are inconsistent with each other. Because the tracks extend through an area of overlapping interest, the agents recognize that they can communicate to try to verify the global consistency of their local interpretations. Simply exchanging the partial solutions—i.e., the track hypotheses without their supporting evidential structures—is sufficient to allow the inconsistency to be detected, but this level of information is not sufficient to allow the inconsistency to be resolved. Resolving the inconsistency in favor of the most likely global interpretation requires an understanding of the quality of the supporting data for the different portions of each track in order to be able to identify the most likely overall interpretation of the data.

Obviously, one way to insure that agents have the information necessary to resolve global inconsistencies would be to always communicate the complete evidential information associated with the solution hypotheses. However, because interpretation hypotheses are complex structures that may be interrelated with numerous other hypotheses, communication and processing limitations

typically make it impractical to fully communicate.¹ Furthermore, complete communication is usually not necessary—e.g., in this example each agent does not have to know the other's actual data. 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 extended process of resolving inconsistency. This requires the ability to integrate incomplete information, represent the resulting uncertainty, and use this uncertainty to drive further actions.

We believe that the DRESUN architecture provides the basis for such a flexible and reactive approach to the communication and use of external evidence. Initial experimentation has been done with a variety of simulated distributed aircraft monitoring scenarios involving local solutions that are globally inconsistent and local solution uncertainty that can be resolved only through agent interactions. These experiments suggest that DRESUN agents can indeed support the complex communication protocols necessary to resolve global inconsistency using highly directed approaches to information exchange like distributed differential diagnosis. The experiments have also shown that such protocols need not be scripted, but can be driven by goals for resolving the uncertainties that arise as information is exchanged among agents.

Because of its representation of inconsistency as a source of uncertainty and its emphasis on directed interactions to resolve inconsistency, DRESUN is different from most DAI work dealing with global inconsistency of local agent beliefs. Much of this work (e.g., [Bridgeland & Huhns 1990, Courand 1990, Huhns & Bridgeland 1991]) has 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 consistency because this can be very expensive both in terms of communication and computation, and it is not always necessary. A DSA system must use an evidential representation (with partial beliefs) rather than a justification-based representation of its beliefs, since the interpretations in a sensor interpretation system are virtually always uncertain.

An evidential belief representation provides both complications and opportunities as compared with justification-based approaches. For instance, in an evidential approach it is not usually possible to attain complete consistency because of the uncertain evidence that underlies each hypothesis. While this means that agent beliefs are never completely consistent, it also means that there is rarely a problem with “hard inconsistency” among agent beliefs. Inconsistencies produce negative evidence and reduced belief (or at least greater uncertainty) in the interrelated hypotheses (i.e., those supporting the inconsistent hypotheses). Thus, instead of having to simply disbelieve anything from the “inconsistent agent” that cannot be proven to be independent of the inconsistency as

¹The 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 need to be processed by most real-world DSA systems.

with [Huhns & Bridgeland 1991], the belief ratings of affected information from the agent will appropriately reflect the strength of the inconsistency and the likelihood that the agent's information is correct.

The next two sections review the RESUN sensor interpretation architecture and the original DRESUN extensions that drive agent interactions based on the need to resolve uncertainty about the global consistency of local solutions. The conceptual basis for the extensions that drive interactions is analyzed in Section 4. Section 5 then examines some of the issues that arise in representing incomplete external evidence, and the example introduced in this section is explored in more detail in Section 6. The paper concludes with a brief summary and a discussion of current research issues.

2 RESUN

Individual DRESUN agents are RESUN interpretation systems [Carver & Lesser 1991a, Carver & Lesser 1993]. This architecture was selected because our earlier research suggested that a sophisticated agent architecture is necessary to support the complex interaction protocols that are required in real-world DSA systems [Lesser 1991]. DRESUN agents are "sophisticated" because they are self-aware and can implement highly context-specific problem-solving strategies. One of the key ideas in RESUN is the use of symbolic *source of uncertainty* statements (SOU) in the evidence for the interpretation hypotheses. The SOUs allow the agents to understand the reasons why their hypotheses are uncertain so they can select methods like *differential diagnosis* that directly resolve uncertainty instead of being limited to (indirect) *incremental hypothesize and test* methods (as in most blackboard-based interpretation systems).

Control decisions are made by a script-based, incremental control planner. The hierarchical goal/plan/subgoal structure created by the control planner provides each agent with an explicit representation of its current goals, the relationships between alternative goals, the relationships between goals and actions, and the status of the methods being used to pursue goals. Planning decisions are made by applying focusing knowledge associated with the subgoals and variables of the particular plan schemas. This design supports highly context-specific control decisions that can explicitly consider the current state of problem solving. The control planner also supports opportunistic and reactive control through its *refocusing mechanism*, which allows focusing decisions to be dynamically reconsidered.

RESUN views interpretation as an incremental process of gathering evidence to resolve particular sources of uncertainty in the interpretation hypotheses: consider what SOUs keep the current answers from being certain, select one or more SOUs to pursue next, take actions appropriate to resolve these SOUs, and repeat this cycle until the termination criteria are met. The overall inter-

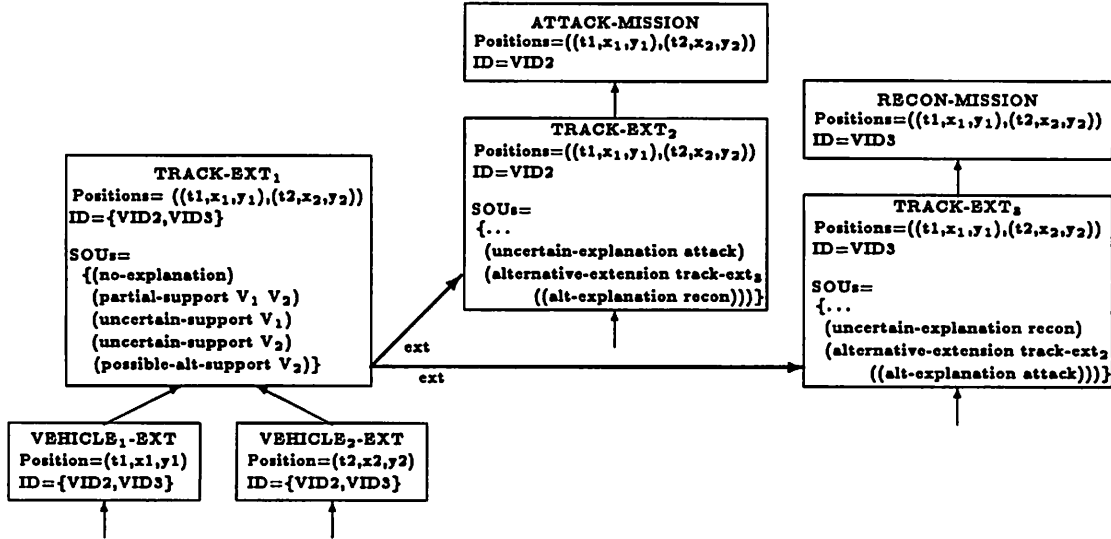


Figure 2: Example representation of RESUN interpretation hypotheses.

pretation process is driven by a high-level model of the state of problem solving, called PS-Model. PS-Model consists of the system's current interpretation solution (the answer-level hypotheses that are currently believed) and a set of SOUs representing the reasons why the current solution is insufficiently believed for termination of problem solving. For example, PS-Model SOUs may denote that some potential answer hypotheses are insufficiently supported or that no evidence has been gathered for a portion of the region of interest (termination in interpretation problems requires not only that existing hypotheses are sufficiently proved or discounted, but that enough of the data has been examined to be sufficiently sure that there are no additional answers).

One reason that RESUN was considered a good base for building a DSA testbed is that RESUN interpretation hypotheses involve multiple levels of representation. This potentially provides a distributed system with great flexibility in what it can communicate. An example of a portion of a RESUN hypothesis is shown in Figure 2. Each hypothesis is maintained as a set of multiple possible versions called *extensions*. Each extension is kept linked to its lower-level, *support* evidence and its higher level, *explanation* evidence. Extensions have different sets of support and explanation evidence and different attribute values; they allow the system to efficiently maintain and pursue alternative versions of a hypothesis. Extension attributes may be imprecise, reflecting imprecise and/or incomplete supporting evidence.

Symbolic SOUs are associated with each hypothesis, extension, and inference in RESUN. The classes of SOUs are based on a model of the uncertainty in abductive inference (the basis for sensor interpretation [Carver & Lesser 1991a, Carver, Cvetanovic, & Lesser 1991]) and the effect that gathering complete evidence has on resolving uncertainty (the basis of hypothesize and test

approaches). The RESUN model of interpretation uncertainty includes the following SOU classes: partial evidence, possible alternative explanations, possible alternative support, alternative extensions (hypothesis versions), negative evidence, and uncertain constraints.

RESUN also includes a scheme for numerically summarizing the evidence for hypotheses and extensions using the symbolic SOUs. This process produces a composite characterization of the uncertainty in a hypothesis or extension in terms of an overall belief (probability) and the relative uncertainty contributions of the different SOU classes. The summary is used in evaluating the satisfaction of termination criteria and in making control decisions. Having a composite rating allows for more detailed reasoning than would be possible with a single number rating. For example, it can be used to distinguish between a hypothesis with low belief due to limited evidence having been gathered so far and one with low belief due to negative evidence—such distinctions can be critical when deciding whether to pursue a hypothesis further.

One of the major sources of uncertainty in (abductive) interpretation inferences is the possibility of alternative explanations for data or hypotheses. For example, while sensor data may be able to be explained as originating from a vehicle (track), it might also be due to ghosting or it might just be noise (even if it is probably from a vehicle, it might not be clear which vehicle it is from). When alternative explanations are pursued, they lead to the creation of alternative hypothesis extensions in RESUN. Figure 2 shows *track-ext₂* and *track-ext₃* which are alternative extensions of *track-ext₁* as a result of two alternative explanations for *track-ext₁*: *attack-mission* and *recon-mission*. These two alternative track extensions are linked via their *alternative-extension* SOUs, which are used during the summarization process to identify alternative hypothesis extensions and propagate belief. For example, as evidence is gathered for *track-ext₂* this causes *track-ext₃* to be less believed—and vice versa.

3 DRESUN

To use RESUN agents for distributed problem solving, the (single-agent) RESUN model had to be extended to drive agent interactions. In DRESUN, interactions result from the goal of insuring the global consistency of local agent solutions. Keeping with the RESUN model, verification of global consistency is driven by additional SOUs in the PS-Model. Unresolved global consistency questions are viewed as sources of uncertainty about the correctness of an agent's solutions. For instance, in the example of Figure 1, when agent A recognizes that its solution intersects an area of overlapping interest with agent B, an SOU is created to represent the uncertainty about the consistency of this agent's and the other agent's interpretations of what is occurring in the overlap area. DRESUN's global consistency SOUs make explicit the possible interrelationships (constraints) between local




Figure 3: Examples of global consistency uncertainties.

subproblems, and provide an integrated view (in conjunction with the standard RESUN PS-Model SOUs) of both the local and global problem solving goals to drive agent control decisions.

Global consistency SOUs are created when an agent recognizes that one or more of its interpretations potentially interact with those of other agents (based on knowledge of the organization of agent areas of interest). 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/explanation for a ghost track hypothesis) must have appropriate external evidence. Examples of situations involving each of these global consistency uncertainties are shown in Figure 3.

When an interpretation includes support from data that intersects a region of overlapping interest, a *consistent-overlapping-model* SOU is added to the agent’s PS-Model. When hypotheses that involve continuous “tracks” of supporting evidence cannot be extended further using an agent’s own data and the extension region for the track intersects another agent’s area, a *consistent-global-extension* SOU will be added to the agents PS-Model. 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 an SOU denoting the possibility that this evidence could be gathered from another agent. This triggers the creation of a *consistent-global-evidence* SOU in the PS-Model.

The three SOUs just described are used to represent unexamined global consistency. For example, a *consistent-overlapping-model* SOU represents the fact that the agent has not yet checked to verify that the associated interpretation is consistent with the external agent’s interpretations (in

the overlap region). Once information is actually obtained from the another agent, it is integrated into the agent's hypothesis structure (see Sections 5 and 6). At that point, any uncertainty due to incomplete knowledge or outright inconsistency would be represented at the hypothesis level. In addition, if inconsistency between local and global evidence is found, a *global-inconsistency* SOU is created in the agent's PS-Model.

4 A Conceptual Description of the DRESUN Approach

One of the goals in developing a distributed version of the RESUN system was to provide a framework that could produce the same global interpretation that would be produced by a centralized system. In other words, the new DRESUN SOUs have to be able to drive agent problem solving activities and the interactions among the agents so that the globally "best" composite interpretation is developed.² Since the new SOUs enforce the global consistency of local solutions, it is instructive to discuss how and to what extent this approach guarantees that the globally best interpretation is created.

For this discussion, it will be useful to think of interpretation as a constraint satisfaction problem (CSP) [Mackworth 1992] and distributed interpretation as a distributed constraint satisfaction problem (DCSP) [Yokoo, Durfee, Ishida, & Kuwabara 1992]. When viewing sensor interpretation as a CSP, each piece of sensor data (or other interpretation evidence) represents a constraint—in conjunction with the causal model of the domain that determines the legality and likelihood of interpretations of the data. Effectively, each piece of data is a constraint of the form: the composite interpretation must have an explanation of what caused this datum. Application of each such constraint may involve considerable processing to determine its implications (i.e., using the causal domain model to identify the consistent and likely interpretations).

There are several characteristics of interpretation when viewed as a CSP that are important for understanding the complexities of the problem. First, interpretation constraints are not simply predicates, they provide numeric ratings of "how well" a solution meets the constraints. These ratings represent the conditional probability that the interpretation is correct, given the constraints (data) examined. Second, interpretation problems are nearly always *underconstrained*. By this we mean that there is residual uncertainty about the correctness of any interpretation even if all of the available constraints are applied. Third, as in most CSPs, it is computationally infeasible to generate all possible solutions of a set of interpretation constraints. Fourth, it is often computationally

²Obviously, this goal has to be met in such a way that it is possible to take advantage of the distributed system's ability to process data concurrently and limit the amount of data that has to be communicated. This requires the ability to support coordination strategies that can sequence actions in different agents so that results are available when needed, limit the amount of redundant work, and minimize the communication of information among the agents.




Figure 4: Examples of why data (constraint) independence is difficult to judge.

Missing Data Connection: Because sensors may sometimes fail to detect a vehicle over a portion of its track, there can be uncertainty about whether/which data is associated with a track—without evaluating the relative likelihood of alternative explanations for the data. While track extension through three missing data positions may be improbable, it could lead to a more likely composite interpretation than the alternative track extension.

Possible Ghost Source: Vehicle tracks may produce ghost tracks (correlated noise), but models of such phenomena may not be sufficient to predict the occurrence of these events nor their exact characteristics (e.g., position). Thus, certainty about the occurrence of a vehicle track may only imply that it is possible for there to be related data within an area around the track, and this area might be large and encompass much data that is not related to the track.

Multi-Vehicle Pattern: Multiple vehicles may be engaged in activities that are part of a common scenario (e.g., target attack), but since the vehicles might be of very different types, widely distributed in both space and time, relations are not obvious even at the track level.

infeasible even to apply all of the constraints to a potential solution (since sensors may generate very large amounts of data). Fifth, constraints can interact in complex ways that are dependent on their solutions—e.g., the data subset $\{d_i\}$ may or may not constrain the interpretation of the data subset $\{d_j\}$ depending on what the correct interpretation of $\{d_i\}$ is determined to be.

Because of the probabilistic nature of constraint satisfaction in sensor interpretation and because interpretation problems are virtually always underconstrained, the goal of an interpretation system is to find the *most likely* interpretation for a set of data. Thus, interpretation involves constraint *optimization* rather than (Boolean) constraint satisfaction [Mackworth 1992]. This makes distributed interpretation different from the DCSPs considered in [Yokoo, Durfee, Ishida, & Kuwabara 1992], and it means that the algorithms presented there are insufficient for distributed interpretation. The interpretation problem is also more complicated than many other CSPs because it is generally impossible to judge interactions between constraints without evaluating all possible solutions to the complete set of constraints.³ Saying this from an evidential point of view, it is impossible to judge

³Consider, for example, a standard CSP with three variables, x_1 , x_2 , and x_3 , and two constraints, $C_1(x_1, x_2)$ and $C_2(x_2, x_3)$. It is straightforward to determine that these constraints interact—in the sense that they result in the possibility of solutions to x_1 constraining the solutions to x_3 (and vice versa). This is the basis of consistency algorithms (e.g., constraint propagation) for solving CSPs. If one views the solutions to individual variables as *components* of the overall CSP solution, another way of describing the consequence of this constraint interaction is to say that the x_1 and x_3 solution components are not independent. Determining the independence of constraints and/or solution components can be critical for interpretation because it is impractical to consider all possible constraints (data) and all possible solutions. However, nearly every interpretation constraint has the potential of interacting with any other constraint—i.e., new evidence might affect the belief in any component of the solution. This can

the independence or conditional independence of evidence from the data level in this domain. To understand why this is so, consider the examples in Figure 4. This complicates problem solving because it means that a considerable amount of work may be required to determine if uninterpreted data can affect the belief in individual hypotheses.

One consequence of these issues is the way that the overall problem-solving goals are mapped into SOUs in the (centralized) RESUN system: uninterpreted data leads to PS-Model SOUs that represent uncertainty about the correctness of the *composite interpretation* rather than uncertainty about individual hypotheses. In other words, the belief ratings associated with hypotheses represent the conditional probability of the hypotheses given the current interpretation evidence, but not given the (known) existence of uninterpreted data.⁴ This can lead to situations in which a hypothesis may be accepted on the basis of partial information, but then rejected after additional data is interpreted (or vice versa).⁵ As a result, potential components of the composite interpretation cannot truly be accepted as meeting the overall problem-solving (termination) goals until all uninterpreted data has been “satisfactorily” processed.⁶

As the example in the introduction section showed, local agent problem solving may not lead to the globally best composite solution. From a DCSP view, it is clear that this is because each agent has only a portion of the globally available constraints and these constraints may interact. Thus, identifying the globally best composite interpretation requires the ability to identify and apply all the globally available constraints that are relevant to each solution component. However, the previous discussion has pointed out why this is difficult in distributed interpretation problems. The key issues are how to identify the constraints among the set of agents that are relevant to each local interpretation component, how to evaluate which of these constraints are critical to producing the correct solution (constraints differ in their power and it is often infeasible to apply all interpretation constraints), and how to do this without excessive communication.

In the DRESUN approach, we do not deal directly with the data-level constraints in attempting to develop the globally best interpretation. Instead, interpretation is driven by ensuring the global consistency of local agent solutions. The SOUs that have been added to DRESUN to extend the RESUN model are sufficient to drive the creation of global composite interpretations that

be understood by the fact that in a conventional CSP formulation of interpretation, the composite interpretation is represented as a *single variable*—whose solution must satisfy all of the constraints—since the solution components are not known ahead of time.

⁴Hypothesis uncertainty due to the possibility of alternative interpretations of the data supporting the hypothesis are accounted for in the belief ratings, but this is done using only the a priori probabilities.

⁵To handle the potential intractability of finding the best interpretation, we accept solution components using a threshold rule: if the belief rating of an answer-level hypothesis meets a certain probability threshold, it is accepted as part of the composite solution.

⁶An important consequence of this is that the termination threshold for “nonanswer” phenomena affects the (true) probability of the correctness of solution components since it determines the extent to which alternative interpretations will be considered.

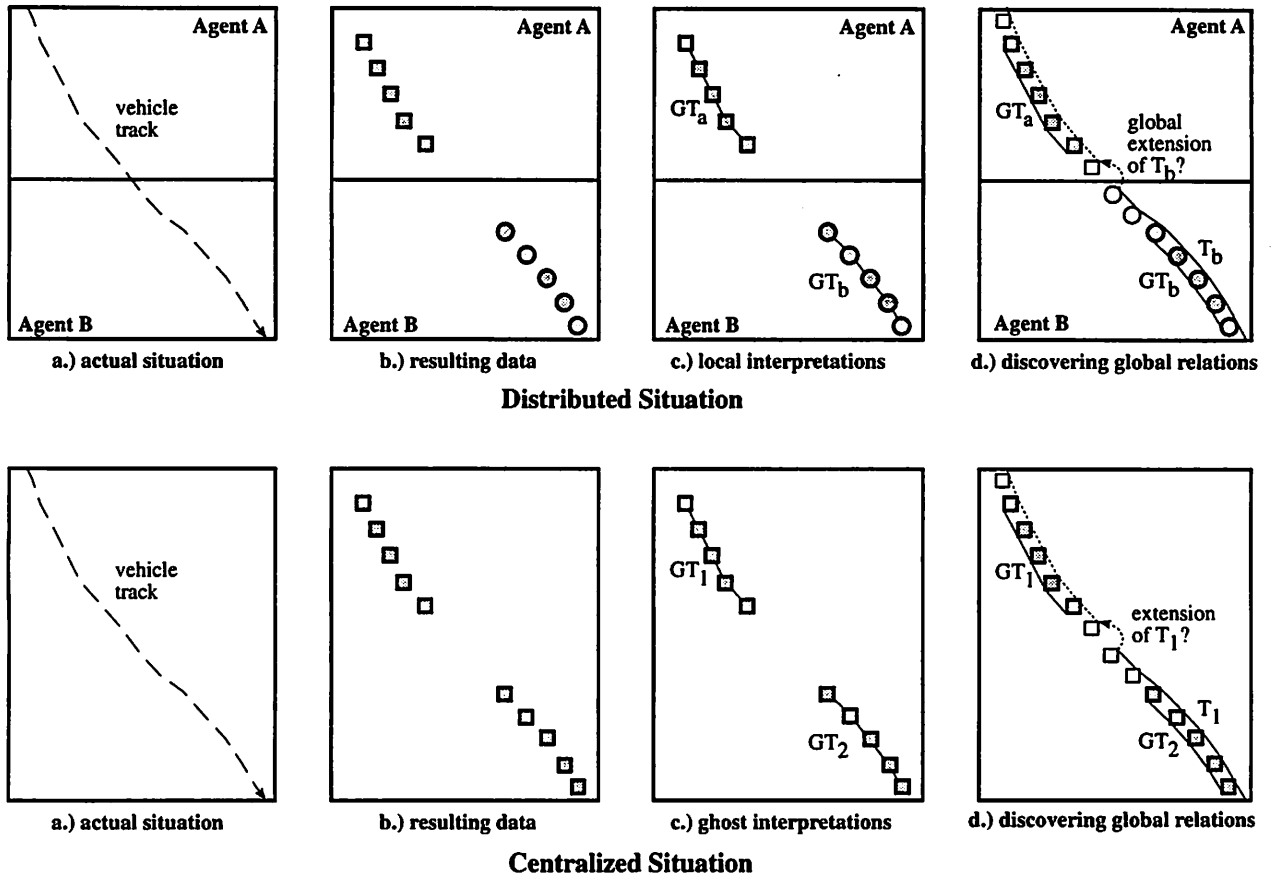


Figure 5: An example where an incorrect global interpretation might be produced.

First consider a multi-agent, distributed situation:

- a.) Assume that this is what actually happens in the environment: a vehicle moves through the regions monitored by each of the agents (for simplicity, we will assume that the agent regions do not overlap).
- b.) Suppose that the data produced by the above event is of very "poor quality" (because of environmental disturbances, poor sensor performance at region boundaries, etc.).
- c.) Because of the quality of the data, each agent locally interprets its portion of the data as most likely being ghost data (for this example we will ignore the issue of finding a source explanation for the ghost tracks). If these interpretations are sufficiently highly believed that the termination criteria are met, then no potential global interactions will be investigated and the "incorrect" global interpretation will be the system solution.
- d.) However, if either agent's belief in its ghost track hypothesis is insufficient to meet the termination criteria, alternative interpretations will be pursued and the potential global interaction discovered. Here Agent B decides to pursue an alternative track interpretations for its "ghost data" and in this process of developing the track interpretation (making assumptions about missing data), it will discover that the potential track could be extended into agent A's region. This would result in communication and development of the globally complete track.

Now consider the centralized situation: a. and b.) The same event occurs and that the same quality data is produced, but this data is now directly available to a single agent.

- c.) The fact that all of the data is visible to the agent may not lead to a different interpretation. The agent begins to interpret its data by identifying some uninterpreted data and developing an interpretation for it. For example, if it begins with some of the GT_1 data, it may develop the ghost track explanation for that portion of its data and be satisfied with that interpretation (the GT_2 data would be similarly processed). Thus, if the ghost interpretations meet the termination criteria, the fact that all of the data is visible will not affect the interpretations that are produced.
- d.) Just as in the distributed situation, development of the "correct" track interpretation requires that the ghost interpretations be too uncertain for termination, so the agent is forced to investigate alternative interpretations.

are the same as what would be created by a centralized RESUN system. To see that this is so, consider first that the global PS-Model SOUs described in the last section identify all the potential global constraint interactions, based on the hypotheses that are explicitly created by the individual agents.⁷ This means that whenever constraints may interact between agents—based on a local interpretation that is deemed possible by at least one of the agents—the points of interaction are identified by PS-Model SOUs that must be resolved to meet the termination criteria. To make this more clear, consider the example in Figure 5. This example appears to be a case in which the “correct” global interpretation cannot be found by our approach. However, the exact same problem occurs in a centralized system. Our approach does not guarantee that the correct interpretation will be developed, but this is not guaranteed by a centralized approach either. Whether or not the correct interpretation is developed depends on the belief ratings (conditional probabilities) that are locally computed for each possible interpretation. If the data is so unusual that incorrect interpretations are judged sufficiently conclusive, then neither the distributed nor the centralized systems will reach the correct conclusions.

Once potential global constraint interactions are identified, DRESUN may initially exchange partial solutions rather than the constraints themselves. Of course, each local interpretation is only an abstract, approximate representation of a particular subset of the local constraints. One consequence is that there may not be enough information about the constraints to be certain about the consistency of the local solutions or to be able to identify the (best) global solution when the local solutions are inconsistent. Another consequence is that when local solutions are revised as a result of further local interpretation (since constraint independence is not readily judged) revisions may be required in the global interpretation. As a result, multi-step dialogs may be required to determine the best interpretation—i.e., the solution that best meets a set of distributed constraints. This process can be viewed as a form of multistage negotiation [Conry, Kuwabara, Lesser, & Meyer 1991].⁸ In [Conry, Kuwabara, Lesser, & Meyer 1991], though, the multiple stages of negotiation were required to develop alternative solutions that were consistent with interacting constraints among the agents (e.g., if agent A has a constraint that links agents A and B, and agent C has a constraint that links agents B and C, the decisions of agents A and C can interact, but this interaction will not be locally obvious to each agent). Here, dialogs also can result from the need to gather more detailed information about constraints and from the need to deal with new constraint relations

⁷The global SOUs are derived from local counterparts. For example, *consistent-global-extension* is related to the local *partial-support* SOU and *consistent-overlapping-model* is related to the local *partial-support-source* SOU. Separate SOUs are used to make the global nature of these relations explicit, and to represent the fact that the local agent initially has a complete lack of knowledge about the relations—e.g., even if it were to process all of its data, it could not resolve the *consistent-global-extension* SOU as it could for a *partial-support* SOU.

⁸The connection between resolution of inconsistency in cooperative distributed problem solving and work on negotiation was noted in [Lesser 1991].

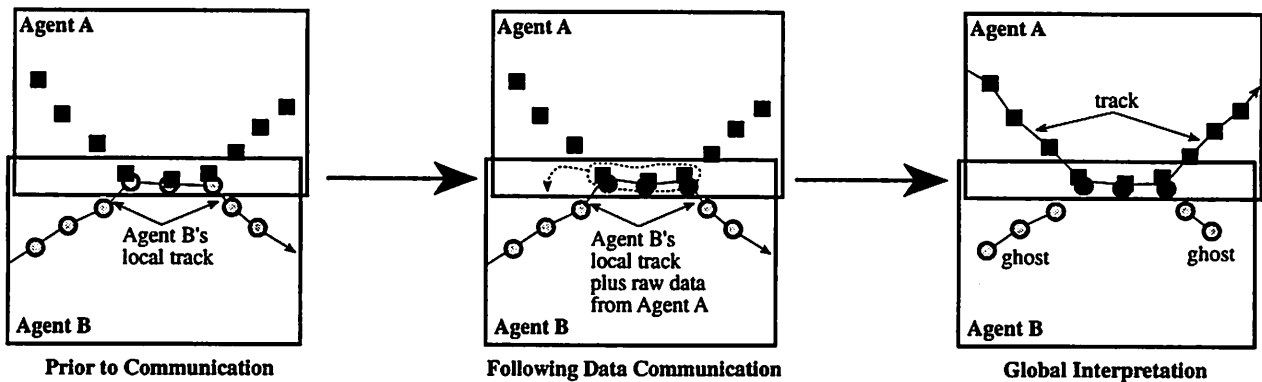


Figure 6: An example of global interpretation uncertainty when transferring raw data.

Here, local agent interpretations interact because there is data in the overlapping region. We assume that agent B is the first to form an interpretation of this data as part of a track based on its local data. When the agent reasons about how to best confirm the global consistency of this interpretation, it may appear that the best way is simply for agent A to transmit its overlap region data to agent B (particularly if agent A has other important work to do and has not yet processed this data). However, while this data corroborates agent B's track interpretation, transmission of only this data means that agent B will not be aware that agent A has additional data that could affect the likelihood of this interpretation. As a result, agent B's track interpretation appears to be globally consistent at this point. That this is incorrect will only become apparent when agent A interprets its other data and finds that there is an alternative interpretation of the overlap data which is both globally consistent and more likely given all the available evidence (note that discovering this alternative requires that agent A's termination criteria force it to gather sufficient evidence to be willing to overturn agent B's highly believed track interpretation).

that are discovered during further problem solving (so multi-stage processes may be required even in two-agent systems).

Ensuring the global consistency of local solutions instead of dealing directly with global constraints has several advantages in this domain. First, it is easier to identify instances where local interpretations must be globally consistent than it is to identify individual constraints that interact across agents. Second, judging consistency of interpretations can be more efficient than applying constraints across multiple agents since much of the work of applying the constraints need not be duplicated and data constraints are relatively costly to apply. Third, constraints may not be visible to both agents at the same time because of how they examine their data. In fact, certain "global constraints" may never become apparent to some agents. For example, a ghost may appear only to one agent's sensor and vehicle tracks may be incompletely sensed or entirely missed by some sensors. While the "absence of data" is clearly an interpretation constraint that may have global implications, lack of data at each such point in space-time cannot practically be represented as an explicit constraint that might trigger global interactions. Thus, global constraints may only be able to be identified by a single agent. This means that the relations represented by the global PS-Model SOUs always represent evidential relations (i.e., they can always lead to evidence for or against the associated hypotheses) even when they do not represent relations between explicit subproblems of the agents.

While uncertainty about constraint interactions (i.e., uncertainty about the independence of hypotheses and data) affects both centralized and distributed vehicle monitoring, its importance has become more apparent to us as a result of work with DRESUN because of some of the differences between the centralized and distributed systems. In particular, agent actions must be coordinated based on (non-local) subproblem interactions to take maximum advantage of a distributed framework. Subproblem interactions are not as important in a centralized (uniprocessor) system because the issue of relative timing of concurrent actions does not arise. Another issue that is primarily of concern in a distributed system is minimizing the communication necessary to reach and confirm a state of global consistency. Here again, uncertainty about subproblem interactions is important since it affects the ability to judge the probability of consistency and so may distort judgements about coordination activities. See, for example, figure 6.

5 Issues in Representing External Evidence

While the architecture as described so far provides the basis for flexible agent interactions, communication of incomplete information or information at different levels of detail complicates the representation of external evidence. Our initial experimentation 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 information from another agent. This section will try to make the differences between representing local evidence and external evidence clear 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 hypothesis extensions 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 extensions 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 RESUN.⁹

⁹RESUN 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 tracks hypotheses can be judged without requiring access to their evidence. The reasons for inconsistency are represented as a set of discrepancy statements (similar to SOUs) and are associated with any negative evidence inference that may result. For instance, the discrepancies could denote that the tracks include different positions at identical times—i.e., that the tracks imply that the vehicle is in two places at once; that merging the tracks would violate vehicle movement constraints; and/or because the vehicle identities (vehicle types) are inconsistent.




Figure 7: Examples of the differences in representing local and external evidence.

When consistent local hypotheses are merged, agents have complete evidence so they can actually construct a new hypothesis (or extension). For example, in the consistent local evidence example in Figure 7, 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 still create a new hypothesis extension T_3 , which has the same attributes (i.e., positions and vehicle ID) as the T_3 created from purely local evidence. However, without access to the evidence for the external T_2 , the belief in 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 can 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 interrelationships among hypotheses. In the case of local evidence only, these relationships can at least be properly evaluated. For example, in the inconsistent local evidence example in Figure 7, 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 results in alternative extensions of V_3 and V_4 with *alternative-extension* SOUs (this is not shown in the figure, but is analogous to what is shown in Figure 2). This representation 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 makes it impossible to properly 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 does not now 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 (i.e., are alternative explanations for some of T_1 ’s support). These alternatives may also be inconsistent with the external evidence or they may be consistent. This will not be discovered automatically and will lead to great uncertainty in the belief in T_1 .

This brief example shows that communication of incomplete hypothesis information can lead to uncertainty about the effect of external evidence on local hypotheses (from the DCSP perspective, this is because hypotheses are abstract representations of constraint information). Communicating incomplete information can still be useful, however, since in many situations the uncertainty may not have to be resolved to satisfy the termination criteria. DRESUN provides the flexibility to communicate incomplete information because it represents any resulting uncertainty using SOUs. These SOUs are used both to decide when further information is required and drive directed communications to obtain this information.

A variety of related representation issues have arisen in DRESUN research using incomplete information and information at multiple levels of detail. We will summarize the key issues here and show examples of several in the next section. To handle unrestricted transmission of incomplete or abstract information among agents, DRESUN has to provide the ability to:

- Link multiple, incomplete views of the evidence for a hypothesis. Acquiring a complete view of the evidence for an external hypothesis is expensive and may not be necessary to meet the termination goals, but transmission of raw data can require expensive redundant interpretation processing. Thus, agents should be able to (incrementally) exchange evidential information at only the abstraction levels that are appropriate for the particular situation (i.e., provide the appropriate level of detail about the interpretation constraints). Among

the consequences, this requires that the evidence evaluation routines have to be able to handle local representations of external hypotheses that have incomplete abstractions of these hypotheses supporting data. This issue is illustrated in the example in the next section.

- Link multiple, incomplete hypothesis extensions (versions). This issue is analogous to the previous issue: it is often costly and unnecessary to acquire all the extensions of an external hypothesis, but local reconstruction can be expensive and redundant, and may be impossible with the locally available information. Thus, agents should (incrementally) exchange extension information as needed, integrate it appropriately into the local representation of the external hypothesis, and make use of the information when relevant to evaluating hypothesis belief. An example of where this issue arises is when the local agent wants to explore alternative local interpretations that are consistent with only alternative or intermediate extensions of an external hypothesis (e.g., alternative extensions or partial segments of a single external vehicle track hypothesis).
- Locally create alternative or intermediate extensions of external hypotheses. This issue is related to the previous one and is illustrated in the example in the next section. In order to pursue alternative interpretations (e.g., because of global inconsistency) an agent may need information about extensions of an external hypothesis that it has not yet acquired. If the uncertainties that result from locally constructing these extensions with incomplete knowledge can be represented, the agent may be able to get some useful evidence and it can identify the information that it needs to obtain from the external agent.
- Reformulate hypotheses for more efficient exploration and representation of alternatives. As result of exchanging evidence among agents, it may become clear that it is necessary to have some particular extension of an agent's hypothesis that has not been constructed by the agent. This will require the agent to reformulate its hypothesis by creating a new intermediate extension or a new alternative extension. It may also require the agent to redefine "the hypothesis" by changing its "anchor" evidence or to create a new alternative hypothesis.¹⁰ The need for reformulation is mentioned in the example in the next section.
- Communicate back results of integrating information that was sent by other agents. When an agent has gathered enough evidence from one or more agents to produce a satisfactory, consistent interpretation, it must transmit back sufficient information about the consistent

¹⁰Each interpretation hypothesis is represented as a set of alternative extensions based around "anchor" evidence (the evidence initially used to create the hypothesis). Anchor evidence is typically selected on the basis of its providing strong support for the hypothesis. However, as further evidence is gathered, it can become apparent that the anchor data is not that strong and that alternative hypothesis versions could be pursued more efficiently if the anchor data was changed or if a new alternative hypothesis was created. Hypothesis reformulation is relevant to single-agent problem solving, but the incomplete view of hypotheses as a result of external evidence makes it more critical in distributed systems.

interpretation to “convince” the other agents that a consistent interpretation has been developed given the currently available evidence. In addition, the level of detail should be sufficient so that it is unlikely for the other agents to rapidly demand further detail as they continue to process their data.

- Identify when shared information should be updated. Since interpretation hypotheses and evidence are virtually always uncertain, as agents continue to gather evidence the level of belief in their hypotheses may change. When these hypotheses have been shared with other agents, it may or may not be important to update the information in the other agents. This depends on the magnitude of the change (e.g., a change in status from believed to disbelieved versus a change in degree of belief from 0.98 to 0.97) and on the role that the evidence played in the external agent (thus, this issue is related to the previous one).
- Avoid circular reasoning when exchanging evidence among agents. Since reaching the termination goals often requires multiple stages of evidence exchange and since evidence relevant to any hypothesis may continue to accumulate, it is important to identify the original source of the belief in hypotheses that are passed between agents. In general, these sources can only be approximately specified (without passing complete evidential structures), but such information can serve to identify when additional detail is required for accurate evidence assessment.

6 Resolving Global Inconsistency: An Example

In this section, we will return to the example of Figure 1 to explore in more detail the role that the representation of external evidence plays in driving the communication of information to resolve global inconsistency in DRESUN. 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 is reached (based on the local evidence) and that agent A is the first to communicate about its SOU.

When agent A initiates a dialog (with agent B) to resolve its “overlap” SOU, there are two options depending on whether the bulk of the processing to check consistency should be done by agent A 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




Figure 8: Agent A's representation of evidence for the example.

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 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 8. 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 Section 5), the ability of agent A to pursue other sources of uncertainty in T_a (to locally increase/decrease its belief 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 an alternative extension of the local track that is consistent with a portion of the external evidence. Exactly which of the possible alternatives is initially created depends on the information that is available about the relative credibility of the various portions 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 pursue an alternative track extension using local evidence from times 5 through 11. Here, the figure shows that this extension, $T_a(5 - 11)$, is an intermediate extension of the hypothesis that had already been created (so agent A does not have to reformulate hypothesis T_a).

The third stage of Figure 8 shows that agent A next creates T_{a+b} , which is an extension of $T_a(5 - 11)$ based on the external track extension $T_b(1 - 6)$. $T_b(1 - 6)$ was selected because agent A determined that it was the maximal portion of $T_b(1 - 10)$ that was consistent with $T_a(5 - 11)$. $T_b(1 - 6)$ is a special kind of external extension because it 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 . 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, agent A at least had the overall rating for $T_b(1 - 10)$). The reasons for this uncertainty are represented by SOUs that are posted in T_{a+b} .

Assuming that agent A decides to pursue T_{a+b} further since it is a credible globally consistent solution, the evaluation SOUs drive the selection of a plan that requests agent B to communicate additional information. The necessary information could be gathered in any of several different ways. For example, agent B may explicitly construct and evaluate extension $T_b(1 - 6)$, and then transmit its belief summary to agent A. This may require agent B to *reformulate* its view of T_b depending on how it gathered the evidence to construct $T_b(1 - 10)$.¹¹ One problem with this approach is that there will still be uncertainty about T_{a+b} because agent A will not have detailed information about the overlapping support in the extensions $T_a(5 - 11)$ and $T_b(1 - 6)$. Resolving this

¹¹For example, if $T_b(1 - 10)$ was constructed with the time 1 data as its anchor, then Agent B will simply have to create the $T_b(1 - 6)$ intermediate extension (if it does not already have it) so that it can develop an alternative extension of the hypothesis. If $T_b(1 - 10)$ was constructed with the time 10 data as its anchor, though, the anchor point of the hypothesis will have to be changed to allow construction and evaluation of the the alternatives $T_b(1 - 10)$ and $T_b(1 - 6)$.

uncertainty requires gathering the support belief summaries for $T_b(1 - 6)$.¹² 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 8. With this level of information, agent A can evaluate the likelihood of T_{a+b} (based on the evidence gathered so far by itself and agent B). Depending on the results of the 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 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 several 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 —it does not need to know the details of its supporting 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 construct alternative hypotheses based on its local data and an incomplete view of T_b , it was able to further 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.

7 Conclusion

In this paper, we have examined some of the agent modeling issues that arise when resolving global inconsistency in DSA systems. For example, we showed how achieving the ability to communicate in a very directed manner, by sending incomplete information or information at different levels of detail, complicates the representation of external evidence. The paper also discusses the conceptual basis of our approach and justifies the set of SOUs that were developed to extend the RESUN model to distributed interpretation. We believe that the DRESUN control architecture in conjunction with the extended evidential representation will provide the flexibility to investigate a wide range of coordination strategies for DSA problems—including strategies for *real-time* DSA. Initial experimentation suggests that the DRESUN framework can support complex agent interaction protocols, based on the evolving SOUs and driven overall by the need to resolve global consistency SOUs. We have recently finished upgrading the representation of external evidence to be able to handle incomplete and uncertain external information as described here, and are assessing whether this representation is sufficient. Because DRESUN supports a range of methods for resolving interpretation uncertainty and global inconsistency, coordination

¹²This 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.

strategies must consider a variety of questions about whether/when/how to pursue interpretations and SOUs (the example of Section 6 mentioned a number of options faced by the agents). We are pursuing both analytical and experimental approaches to determine appropriate coordination strategies [Decker & Lesser 1992], 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

- [Bridgeland & Huhns 1990] David Bridgeland and Michael Huhns, "Distributed Truth Maintenance," *Proceedings of AAAI-90*, 72-77, 1990.
- [Carver & Lesser 1991a] Norman Carver and Victor Lesser, "A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty," *Proceedings of AAAI-91*, 724-731, 1991.
- [Carver & Lesser 1991b] Norman Carver and Victor Lesser, "Blackboard-based Sensor Interpretation Using a Symbolic Model of the Sources of Uncertainty in Abductive Inferences," *Proceedings of AAAI Workshop: Towards Domain-Independent Strategies for Abduction*, 1991.
- [Carver, Cvetanovic, & Lesser 1991] Norman Carver, Zarko Cvetanovic, and Victor Lesser, "Sophisticated Cooperation in FA/C Distributed Problem Solving Systems," *Proceedings of AAAI-91*, 191-198, 1991 (also available as Technical Report 91-23, Computer and Information Science Department, University of Massachusetts).
- [Carver & Lesser 1993] Norman Carver and Victor Lesser, "A Planner for the Control of Problem Solving Systems," to appear in *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Planning, Scheduling and Control, vol. 23, no. 6, 1993.
- [Conry, Kuwabara, Lesser, & Meyer 1991] Susan Conry, Kazuhiro Kuwabara, Victor Lesser, and Robert Meyer, "Multistage Negotiation for Distributed Constraint Satisfaction," *IEEE Transactions on Systems, Man, and Cybernetics*, special issue on Distributed Artificial Intelligence, vol. 21, no. 6, 1462-1477, 1991.
- [Courand 1990] Gregory Courand, "Cooperation via Consensus Formation," *Proceedings of the 10th International Workshop on Distributed Artificial Intelligence*, 1990.
- [Decker & Lesser 1992] Keith Decker, and Victor Lesser, "Formalizing the Analysis of Quantitative Coordination Relationships," *11th International Workshop on Distributed Artificial Intelligence*, 1992.
- [Durfee & Lesser 1991] 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.
- [Huhns & Bridgeland 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.

- [Lesser & Corkill 1981] Victor Lesser and Daniel Corkill, "Functionally Accurate, Cooperative Distributed Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 1, 81-96, 1981.
- [Lesser & Corkill 1983] 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 Englemore and Tony Morgan, editors, Addison-Wesley, 1988).
- [Lesser 1991] 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.
- [Mackworth 1992] Alan Mackworth, "Constraint Satisfaction," in *Encyclopedia of Artificial Intelligence*, second edition, Stuart Shapiro, editor, Wiley-Interscience, 285-293, 1992.
- [Whitehair & Lesser 1993] Robert Whitehair and Victor Lesser, *A Framework for the Analysis of Sophisticated Control in Interpretation Systems*, Technical Report 93-53, Computer Science Department, University of Massachusetts, 1993.
- [Yokoo, Durfee, Ishida, & Kuwabara 1992] Makoto Yokoo, Edmund Durfee, Toru Ishida, and Kazuhiro Kuwabara, "Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving," *Proceedings of the Twelfth International Conference on Distributed Computing Systems*, 614-621, 1992.