A Formal Analysis of Solution Quality in FA/C Distributed Sensor Interpretation Systems

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

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

Abstract

The functionally-accurate, cooperative (FA/C) distributed problem-solving paradigm is one approach for organizing distributed problem solving among homogeneous, cooperating agents. While several FA/C systems have been implemented, there has been little formal analysis of the quality of the solutions that can be produced using the approach or of the conditions that are necessary for the approach to be effective. This paper reports on work we have done to formally analyze the FA/C model in the context of distributed sensor interpretation (SI). Results are presented that compare the quality of solutions produced by a distributed FA/C system to those produced by an equivalent centralized system. We first establish that while it is possible for an FA/C system to produce the same solution as a centralized system, this requires the use of interpretation and coordination strategies that are impractical for most SI applications. We then consider the effect of "approximate" interpretation and coordination strategies.

Introduction

In the *functionally accurate, cooperative* (FA/C) paradigm for distributed problem solving (Lesser & Corkill 1981; Lesser 1991), agents are designed to produce tentative, partial results based on only local information. They then exchange these results with the other agents, exploiting constraints among their subproblems to resolve the global inconsistencies and local uncertainties that occur due to the agents' lack of accurate, complete, and up-to-date information. While several systems that use the FA/C approach have been built (e.g., (Carver, Cvetanovic, & Lesser 1991; Carver & Lesser 1995; Lesser & Corkill 1983)), there has been little formal analysis of the potential performance of such systems or the role that domain conditions and coordination strategies play in achieving acceptable performance.

This paper reports on our efforts to date in formally characterizing the performance of FA/C-based distributed problem solving systems. Most FA/C applications have been in distributed *sensor interpretation* (SI) and our analysis focuses on this domain. We believe, however, that our basic results should be applicable to other distributed problem-solving domains in which agent solutions are interdependent and local problem solving involves approximate, satisficing search.

The aspect of performance that we address here is the quality of the (global) solutions that can be produced by FA/C systems. To analyze solution quality, we need some standard for comparison. One possible definition of the best global solution would be the *most probable explanation* (MPE) (Pearl 1988) given all of the *globally available* data. The problem with this definition is that for most real-world SI problems it is impractical to compute the (true) MPE. For this reason, and because our main interest is the effect of FA/C distributed problem solving, we have chosen to use the solutions of an *equivalent* centralized system as the standard. By this we mean the solutions that would be produced by a centralized (single-agent) system that has access to all of the data of the distributed system and uses the same interpretation strategies that are used by the individual agents of the distributed system (in processing their local data).

Two theorems are presented that compare the solutions of an FA/C-based SI system to the solutions that would be produced by an "equivalent" centralized SI system. The theorems differ in the interpretation and coordination strategies that are assumed. We first show that it is possible for an FA/C-based distributed SI system to produce exactly the same global interpretation as would be produced by a centralized system (and that this is the optimal solution). However, "exact" interpretation and coordination strategies are required.

We then consider the effect of more realistic "approximate" strategies. The key issue we explore is whether FA/C-based SI systems can focus on the agents' local solutions (i.e., the interpretations that locally look "best"). This can greatly increase the efficiency of distributed problem solving. Unfortuntely, we show that such strategies do not in general provide any guarantees about the quality of the global solutions (regardless of the local interpretation strategies being used). In (Carver, Lesser, & Whitehair 1996) we suggested that many SI problems have a characteristic that can allow such strategies to produce reasonable global solutions. We have called such problems *nearly monotonic*. As part of our discussion of "local solutions coordination strategies" we will discuss nearly monotonic problems and the interaction between the agents' local problem solving and the global termination process.

In the next section we describe our model of distributed sensor interpretation and FA/C agents. This is followed by two sections that discuss interpretation and coordination strategies. The Solution Quality Results Section contains the solution quality theorems. The paper concludes with a summary of our work and future research plans.

Distributed Sensor Interpretation

By sensor interpretation, we mean the determination of high-level, conceptual explanations of sensor data and other information. Our model of the interpretation process will be essentially that used in (Carver & Lesser 1991). Each *interpretation* of a data set is an explanation of what caused the data. In general, an interpretation will be a *composite* of a set of *hypotheses*, each of which explains some subset of the data and which together explain all of the data. A *solution* is an interpretation that is judged "best" according to some criteria (we will return to this issue below).

In a centralized SI system, all of the sensor data is available to the single agent. In a distributed system, each agent has (direct) access to data from only a subset of the sensors and each sensor is associated with a single agent. As a result, each agent monitors only a portion of the overall area of interest, and agents' local solutions must be "combined" in order to construct a global solution. This may not be straightforward, however, because the local solutions are often not independent. In fact, they may be inconsistent because they are based on different data sets. Agent solutions are *interdependent* whenever data (evidence) for an interpretation hypothesis is spread among multiple agents or when agent areas of interest overlap as a result of overlapping sensor coverage.

In an FA/C system, there must be some mechanism to drive interactions 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 allowed agents to understand where there are constraints among their subproblems, so that information interchange could be highly directed. DRESUN (Carver & Lesser 1991; Carver, Cvetanovic, & Lesser 1991; Carver & Lesser 1995) provides this capability, and it will form the basis for our model of the capabilities of an FA/C agent. DRESUN agents create symbolic *source of uncertainty* statements (SOUs) to represent the reasons why their hypotheses are uncertain. Whenever an inter-agent subproblem interaction (constraint) is detected, DRESUN agents create a *global consistency SOU* (GSOU). Each GSOU denotes that another agent(s) can provide evidence for/against the agent's associated local hypothesis.

Interpretation Strategies

By *interpretation strategy* we mean the algorithm that a distributed or centralized SI agent uses to process its locally available data to arrive at the "best" interpretation of that data. As we have said, one standard definition of the best interpretation of a data set is the MPE, but this is impractical to compute for most SI problems.[1] It is not the goal of this paper to investigate interpretation strategies in detail. The only distinction we will make is between "exact" and "approximate" strategies. We will use the name *exact local interpretation strategy* (ELIS) to refer to a strategy in which an agent's local solution (ignoring global interactions) would be the true MPE of the local data set. The term *approximate local interpretation strategy* (ALIS) will be used to refer to any strategy that is not exact.

We will not be specific about what constitutes an ALIS, since the subject is complicated and strategies and their effects will be largely domain specific. The key point is that a solution produced by an ALIS will be an approximation of the MPE: it may include hypotheses that are not in the most probable composite interpretation and/or it may be missing some that are. We will consider both exact and approximate interpretation strategies. It is important to keep in mind, though, that most real-world SI systems—whether centralized or distributed—will be forced to use *approximate, satisficing* interpretation strategies.

Coordination Strategies

While DRESUN agents have a representation of all inter-agent interrelationships (based on the locally created hypotheses), we must now consider how this information should be used to develop global solutions. In other words, what are appropriate *coordination strategies*: what agents should each agent communicate with, what information should they communicate, when should they communicate, and so forth. The aspects of coordination that we focus on are what are covered by the notion of "resolving a GSOU" in the DRESUN model. Each GSOU denotes that another agent(s) can provide evidence about a local interpretation hypothesis. In order to "resolve a GSOU," information must

[1] A thorough discussion of this issue is beyond the scope of the paper. Determination of the MPE has been shown to be NP-hard, but there are efficient algorithms for some classes of belief networks. The key source of complexity for SI is what is known in the target tracking literature as the *data association problem* (Bar-Shalom 1988): which target should data be associated with? The data association problem and the possibility of an indeterminate number of top-level causes can lead to an exponential growth in the number of possible interpretations. Factor in the existence of massive amounts of data in many SI systems, and SI problems are very different from the "diagnosis problems" typically studied within the belief network research community. Despite these differences, it can sometimes be useful to think of the networks of data and hypotheses that SI systems construct as being similar to belief networks.

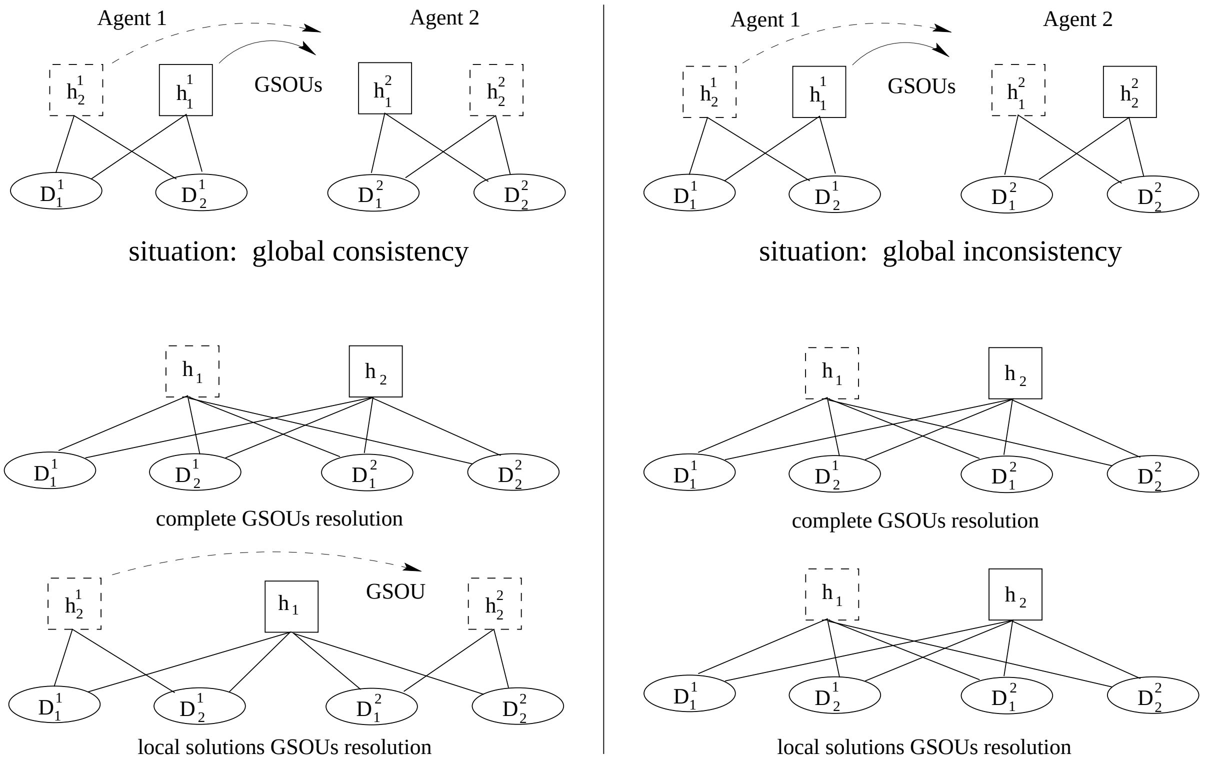

Figure 1: Example results from the different GSOU resolution strategies.
Agent A_1 is resolving a GSOU associated with its solution hypothesis h_1^1, with agent A_2 (note that hypotheses that are deemed part of the solutions are shown with solid boxes while those that are not are shown with dotted boxes). In one case, the agent solutions are globally consistent: A_2 has consistent solution h_1^2. In the other case, the agent solutions are *in*consistent: A_2 has inconsistent solution h_2^2 (A_2 judged h_2^2 "better" than h_1^2). The *complete GSOUs resolution strategy* produces the same result in both cases: both possible interpretations of the globally available data, h_1 and h_2, are constructed and evaluated (and h_2 is here shown to be judged better). The *local solutions GSOUs resolution strategy* produces different results in the two cases. When the local solutions are consistent, the merged solution h_1 is created, without considering the complete set of global interpretations (i.e., without constructing h_2), and the GSOU associated with h_2^1 is never resolved. Inconsistent local solutions can, however, cause the agents to do the additional propagation. Here the same solution as with the CRS is developed.

be exchanged between the interrelated agents so as to propagate the effects of each agent's (interrelated) local evidence between the agents. Resolution of a GSOU is analogous to local (intra-agent) evidence propagation when additional data is interpreted. Conceptually, evidential links are established among the associated agents' "belief networks" (the networks of interpretation hypotheses and sensor data).

As with the local interpretation process, there are a range of possible GSOU resolution strategies. These strategies differ in how accurately global interactions are considered. They make different choices about which global SOUs to pursue and how completely to assess the evidential effects of external agent data. The most comprehensive resolution strategy is for *all* the GSOUs in all the agents to be *completely* resolved: for each GSOU the entire set of possible interpretation hypotheses constructed by the relevant external agents are examined and used to create (and evaluate) all possible joint interpretation hypotheses. We will refer to this as the *complete GSOUs resolution strategy* (CRS).

A key criterion that we impose on coordination strategies for FA/C agents is that the agents' local solutions must be "consistent" at termination.[2] This allows the local solutions to be merged into a global solution without the need for additional problem solving. The CRS meets this criteria as long as all agents associated with a GSOU are somehow made aware of the result of resolving the GSOU. An "incomplete" coordination strategy that meets the global consistency criterion is the *local solutions GSOUs resolution strategy* (LSRS): each agent to resolves only those GSOUs that are associated with its local solution ("best" interpretation) and pursue evidence propagation using only the local solution hypotheses of the relevant external agents. The *consistent local solutions strategy* described in (Carver

[2] Our definition of global consistency of local solutions is an evidential one: solutions are consistent if their hypotheses are pairwise identical, independent, or corroborative. Local solutions are *in*consistent when any of their hypotheses are contradictory—i.e., have a negative evidential relationship.

& Lesser 1996) is an instance of a LSRS.

Figure 1 shows examples of GSOU resolution using both the CRS and the LSRS, when local solutions are globally consistent and inconsistent. While the CRS produces the same global solution in both cases, the LSRS does not. This is because the effects of the agents evidence are not fully propagated with the LSRS when the agents local solutions are consistent. The figure shows one possible outcome for the example when the local solutions are inconsistent, but we have deliberately left unspecified, the behavior of an LSRS when local solutions are inconsistent (because it is not required for our results). See (Carver & Lesser 1996) for a more complete description of one possible LSRS.

Finally, the quality of the global solutions that result from these coordination strategies is clearly influenced by the local interpretation strategy being used by the agents. In DRESUN, GSOUs are created based on only the interpretation hypotheses that are actually constructed. Thus, since an ALIS may not construct all possible interpretations, a DRESUN system may fail to detect the existence of some global evidential interactions (affecting inter-agent propagation and the global solutions that are chosen).

Solution Quality Results

In this section we will examine the quality of solutions that can be produced using an FA/C approach, relative to those that could be produced using an equivalent centralized approach. The following notation will used in presenting the results and their proofs:

- $\mathcal{A}$ is the set of agents $\{A_1, A_2, \ldots\}$, with their interest area specifications.
- $\mathcal{D}$ is the complete data set available to the (distributed or centralized) system.
- D_i denotes some subset of $\mathcal{D}$ ($D_i \subseteq \mathcal{D}$).
- $\mathcal{D}^i$ is the complete data set available (directly) to only agent A_i—i.e., the complete data set that is available from agent A_i's own sensors.
- $\mathcal{D}^G$ refers to the complete, globally available data set—i.e., the combined data from all of the agents, $\mathcal{D} = \bigcup_{A_i \in \mathcal{A}} \mathcal{D}^i$. $\mathcal{D}^G = \mathcal{D}$ for an equivalent centralized system.
- $\mathcal{I}(D_i)$ is the true MPE of the data set D_i.
- $\hat{\mathcal{I}}(D_i)$ is the "best interpretation" solution of the data subset D_i, given the local interpretation strategy being used (which may be approximate).
- $\mathcal{I}^i$ is the MPE for agent A_i's data (and any external evidence known to A_i).
- $\hat{\mathcal{I}}^i$ is agent A_i's "best interpretation" solution—given the local interpretation and coordination strategies being used.
- $\hat{\mathcal{I}}^G$ denotes the combined, global "best interpretation" solution. That is, $\hat{\mathcal{I}}^G = \bigcup_{A_i \in \mathcal{A}} \hat{\mathcal{I}}^i$. $\hat{\mathcal{I}}^G$ is defined only when all the $\hat{\mathcal{I}}^i$ are consistent, as in a final context (at termination).
- $\hat{\mathcal{I}}^C$ denotes the "best interpretation" solution for an equivalent centralized system.
- $BEL(h, D_i)$ denotes the true degree-of-belief in hypothesis h based on D_i—i.e., $P(h \mid D_i)$. This may not be the belief rating actually associated with h.
- $BEL^i(h, D_j)$ denotes the true degree-of-belief in hypothesis h computed by agent A_i based on D_j—i.e., $P(h \mid D_i)$.

Exact, Complete Strategies

In this section we will examine the quality of FA/C global solutions when the centralized system and the distributed agents use the ELIS to process their local data and the distributed agents use the CRS for coodination. To rephrase the main question we are interested in addressing in terms of our notation: what is the relationship between $\hat{\mathcal{I}}^G$ and $\hat{\mathcal{I}}^C$ given these particular interpretation and coordination strategies?

Theorem 1: *Given, a centralized system that uses the ELIS and a set of agents $\mathcal{A}$ each of which uses the ELIS and the CRS. Then $\forall \mathcal{D}$ and $\forall \{\mathcal{D}^i\}$ where $\mathcal{D}^G = \mathcal{D}$ (i.e., for all possible divisions of the data among the agents), $\hat{\mathcal{I}}^G = \mathcal{I}(\mathcal{D}) = \hat{\mathcal{I}}^C$.* In other words, when agents use the ELIS and the CRS, the FA/C distributed system produces the exact same interpretation as the equivalent centralized system, and this is the true MPE of the globally available data.

Proof: Clearly $\hat{\mathcal{I}}^C = \mathcal{I}(\mathcal{D})$ by the definition of the ELIS and the fact that the centralized system (agent) has access to the complete data set $\mathcal{D}$. To show that $\hat{\mathcal{I}}^G = \mathcal{I}(\mathcal{D})$ we must show that $h \in \hat{\mathcal{I}}^G$ implies $h \in \mathcal{I}(\mathcal{D})$ and that $h \in \mathcal{I}(\mathcal{D})$ implies $h \in \hat{\mathcal{I}}^G$. To prove these implications, we will have to specify two properties for the ELIS and CRS: (1) when using the ELIS, each agent constructs *all* possible interpretations of its available data and evidence[3]; (2) using the ELIS and CRS in conjunction, enough evidential information must be transferred among the agents when "resolving a GSOU" so that the local agent solutions together represent the MPE given the agents' joint data/evidence.[4] Now, suppose that $h \in \hat{\mathcal{I}}^G$. This means that there is at least one $A_i \in \mathcal{A}$ such that $h \in \hat{\mathcal{I}}^i$ at termination and h is globally consistent. Because of our specification of the

[3]This simply means that agents cannot "magically" guarantee that they will arrive at the optimal, MPE solution without constructing the complete "belief network."

[4]This expands on what it means to "completely resolve a GSOU." While we previously stated that all possible joint interpretations be constructed and evaluated, we did not require that this evaluation be "exact." Here, we specify that when agents are locally using the ELIS, solution membership of the joint interpretation hypotheses will be correct. We leave open whether belief ratings will be the true condition probabilities See (Pearl 1988) for a discussion of the differences between computing beliefs (belief updating) and computing MPE solution membership (belief revision) in probabilistic belief networks.

behavior of the ELIS and CRS, A_i must know enough about the subset of $\mathcal{D}$ D, that is relevant to h[5] such that it can determine if $h \in \mathcal{I}(D)$. Since each agent is using the ELIS, if $h \in \hat{\mathcal{I}}^i$ then $h \in \mathcal{I}(D)$ and so $h \in \mathcal{I}(\mathcal{D})$. The proof of the converse implication is similar. Since $h \in \mathcal{I}(\mathcal{D})$, h must have been produced by at least one agent A_i.[6] If h is an interpretation hypothesis known to A_i, then by our specification of behavior of the ELIS and CRS we can again see that A_i must know enough about the subset of $\mathcal{D}$ D, that is relevant to h that it can determine if $h \in \mathcal{I}(D)$. If $h \in \mathcal{I}(\mathcal{D})$ then $h \in \mathcal{I}(D)$, and since A_i is using the ELIS it must have $h \in \hat{\mathcal{I}}^i$. With the CRS, all agents with data relevant to h will reach consistent decisions about h's solution membership. Thus, it must be that $h \in \hat{\mathcal{I}}^G$.

This result is what we would expect given the definitions of the ELIS and CRS, and given the assumptions inherent in the DRESUN model (that the GSOUs represent all possible interactions for the set of created hypotheses). However, it serves to formally establish the theoretical capabilities of a DRESUN-based FA/C distributed SI system. Other FA/C-based distributed SI architectures have not had this property because they lacked a representation of all the inter-agent interactions and a mechanism for controlling their resolution. The result shows that an incremental, distributed approach to SI need not produce poorer quality solutions than a centralized system (when both systems have access to exactly the same data). Furthermore, the result is not trivially obvious. We had to specify that agents create all possible interpretation hypotheses and we had to be careful in how we defined the interactions between the local interpretation strategy and the GSOUs resolution strategy for the result to hold.

Unfortunately, the practical implications of this result are limited since it is rarely practical to do exact local interpretation and complete GSOUs resolution. The only time this might be possible would be if enough agents are used so that each agent's local processing requirements are adequate to completely and exactly process its local data and if the agents' solutions are largely independent. Since even centralized SI systems will typically have to settle for approximate, sub-optimal solutions, what we are really interested in determining is whether/when FA/C distributed systems can produce solutions that are *comparable* to those produced by approximate centralized systems—and whether they can do it *efficiently*.[7]

[5] D is the data subset "relevant" to h if D is the smallest subset of $\mathcal{D}$ such that $BEL(h, D) = BEL(h, \mathcal{D})$.

[6] Actually, no single agent may have produced *exactly* h. Each agent may locally have only a subset of the data in $\mathcal{D}$ that supports h, and because interpretation problems are not "propositional," additional evidence both modifies a hypothesis' belief and refines the values of its attributes—creating new versions of the hypothesis (what (Carver &Lesser 1991) terms *extensions*). Thus, what we really mean is that a precursor version of h would have been created by at least one of the agents.

Local Solutions Coordination Strategies

In this section we undertake some evaluation of the effects of the LSRS. This is an important strategy because it can be very efficient when local agent solutions are frequently consistent. In addition, the developers of the FA/C paradigm clearly saw this as a useful coordination strategy ((Lesser & Corkill 1981) refers to *consistency checking* of the tentative local solutions with results received from other nodes as "an important part of the FA/C approach") and it has been used in implemented FA/C systems.

Theorem 2: *Given, a centralized system using an arbitrary (possibly approximate) local interpretation strategy and a set of agents $\mathcal{A}$ each of which uses this same local interpretation strategy along with the LSRS, and the domain is not* solution monotonic.[8] *Then $\exists \mathcal{D}$ and $\exists \{\mathcal{D}^i\}$ with $\mathcal{D}^G = \mathcal{D}$ (i.e., there exists some data and some division of the data among the agents), such that $\hat{\mathcal{I}}^G \neq \hat{\mathcal{I}}^C$ and $\hat{\mathcal{I}}^G \neq \mathcal{I}(\mathcal{D})$.* In other words, the LSRS does not guarantee that an FA/C distributed system produces the same exact interpretation as the equivalent centralized system, or the true MPE of the globally available data.

Proof: Regardless of the local interpretation strategy being used, there must exist some data set $\mathcal{D}$ and some division of that data among the agents $\{\mathcal{D}^i\}$, such that there is a hypothesis h where $h \in \hat{\mathcal{I}}^i$ and $h \in \hat{\mathcal{I}}^j$ (for two distinct agents A_i and A_j), but it is not the case that $h \in \hat{\mathcal{I}}^C$. Consider a situation in which h_a and h_b are alternative interpretation hypotheses for data sets D_1 and for D_2, and these hypotheses are independent of all other data. It is possible to have $BEL(h_a, D_1) > BEL(h_b, D_1)$ and $BEL(h_a, D_2) > BEL(h_b, D_2)$, but $BEL(h_b, D_1 \cup D_2) > BEL(h_a, D_1 \cup D_2)$. If $\mathcal{D}^i = D_1$

[7] To clarify this point, a distributed SI system does not necessarily do more "work" than the centralized system does to get the MPE solution (the complexity of a particular SI problem instance is inherently determined by the interrelatedness/independence of the resulting interpretation hypotheses). Completely resolving all GSOUs results in the distributed system performing precisely the same evidential propagation as would have to be performed by a centralized system doing exact interpretation. However, doing this propagation via inter-agent communication will be more costly both in terms of processing and elapsed time than it would be in a centralized system. This same issue arises in developing approximate coordination strategies: the agents in the distributed system will typically have more limited views of the globally avaiable data than would a centralized system, so approximation strategies that are appropriate for centralized SI systems may not be efficient for distributed systems.

[8] By *solution monotonic* we mean a domain in which $\forall D_i \subset \mathcal{D}$, if $h \in \mathcal{I}(D_i)$ then $h \in \mathcal{I}(\mathcal{D})$. In other words, if a hypothesis is in the MPE solution to any subset of the data then it is guaranteed to be in the MPE solution to the complete data set.

and $\mathcal{D}^j = D_2$, the LSRS would produce $h_a \in \hat{\mathcal{I}}^G$, but $h_a \notin \mathcal{I}(\mathcal{D})$ so $\hat{\mathcal{I}}^G \neq \mathcal{I}(\mathcal{D})$. The only situation in which this would *not* be a possibility would be a domain that was solution monotonic. A similar argument involving possible distributions of data among the agents can be used to show that it is possible that $\hat{\mathcal{I}}^G \neq \hat{\mathcal{I}}^C$.

This result is simply a consequence of the nonmonotonicity inherent in most domains, when reasoning is based on incomplete information. Because "consistency" of local solutions does not in general guarantee that the merged, global solution is the best solution, the LSRS may not produce optimal global solutions or even solutions that are identical to an equivalent (approximate) centralized system.[9] This is unfortunate since the strategy can be very efficient, and it also conflicts with the intuitions of ourselves and other FA/C researchers that this is a reasonable approximate FA/C coordination strategy.

In (Carver & Lesser 1996) we proposed that one possible explanation for this apparent contradiction is that many SI domains have characteristics that allow this strategy to work quite well. We have termed such problems *nearly monotonic*. The basic idea is that while belief and solution membership are nonmonotonic with increasing evidence, once certain conditions (like fairly high belief) are achieved then solutions become nearly monotonic in behavior. What this observation means for the LSRS is that if a hypothesis is part of a local solution, has "fairly high" belief, and is globally consistent, then it is likely to be correct. To formalize this notion we have used probabilistic models of the likelihood of a hypothesis being correct given its current belief based on local agent data only. (Carver & Lesser 1996) contains a theorem that shows that in nearly monotonic domains the LSRS can *potentially* produce a global solution whose components are as likely to be in the MPE global solution as desired (by selecting appropriate criteria that local solutions must meet prior to being exchanged).

Recognizing the role that near monotonicity plays in the successful application of the LSRS has led to several important conclusions about the design of coordination strategies for FA/C-based distributed SI. For example, the LSRS can be extended by recognizing that consistency of solution components is not the only criterion that should be used in deciding whether to terminate processing. While the basic LSRS simply accepts globally consistent agent solution components, this should happen only if those components meet certain properties, such as sufficiently high local belief or sufficient consideration of competing, alternative explanations. Otherwise, local problem solving should be "reactivated," to generate and evaluate alternative explanations (just as in the case of global inconsistency). From this perspective, the concept of near monotonicity is extended from dealing with single hypotheses to collections of interpretations that have been evaluated based on their global consistency. Another conclusion from this work is that certain local solution components are more useful than others in developing the global solution. This confirms intuitions of researchers that communication of highly uncertain local solution components is not usually advantageous.

[9] Just because solutions from such a distributed system are not identical to those of an equivalent centralized system does not necessarily mean that they will always be worse (relative to the MPE solution). For example, inconsistency of local solutions could drive a distributed system to do more evidence propagation than an approximate centralized system, resulting in the distributed system having a better solution (at the expense of the communication to do the propagation). On the other hand, in a centralized system an approximate interpretation strategy can consider all of the available data—at least in an abstract manner. Unless the agents subproblems are largely independent, we would therefore expect a centralized system with a good approximate interpretation strategy to perform better on average than a distributed system with limited inter-agent communication.

Conclusions

This paper contains some basic results about the quality of solutions that can be produced by FA/C-based distributed SI systems. First, while it is possible for an FA/C system to produce the same solution as a centralized system (and the true MPE of the available data) this requires strategies that are typically impossible for real-world SI. Second, while coordination that focuses only on local solutions has been popular in FA/C systems and can be efficient, it is not guaranteed to produce optimal or centralized-equivalent solutions. However, we have shown that under certain conditions in the domain it is at least possible for an FA/C-based SI system to efficiently produce results that are "comparable" to those of a centralized system.

In our current research, we are concentrating on pursuing the concept of nearly monotonic problems: designing coordination strategies that are efficient and that have predictable performance in terms of the probability of the global solutions being optimal. We believe the notion of near monotonicity will support some of the intuitions that people have had about approximate FA/C-based SI: certain types of hypotheses are better to construct and communicate, and it is useful to achieve a certain local level of belief in hypotheses before communicating them. In each of these cases we see that when approximate strategies are being used, the likelihood of global solutions being correct will be improved.

We should note that in our discussion of coordination strategies, we largely ignored the question of *when* information should be communicated, as well as the possible communication of meta-level information. Different coordination strategies for how and when to communicate the information necessary to resolve a GSOU can greatly affect the efficiency of the process. This issue is related to the question of local and global strategy interactions. For example, suppose that an agent wants to resolve a certain GSOU, but the appropriate external agent has not yet processed (interpreted) the data and created the hypotheses necessary to do this. Should the initiating agent wait for the external agent to get around to doing this? Should it instead be sent the raw data and do with it what it wants, even if that might produce a different result than would have been produced in the first scenario? Since our focus in this paper is on solution quality that can attained by FA/C distributed SI systems, we tried to separate the effects of the local and global strategies as much as possible. In practice, however, their interactions can have a profound effect on efficiency and solutions. Thus, while we have examined conditions that are *necessary* for effective FA/C problem solving, these conditions are not sufficient to guarantee it. Coordination strategies for *efficient* FA/C problem solving remain an important area of research.[10]

Acknowledgements

This work was supported in part by the Department of the Navy, Office of the Chief of Naval Research, under contract N00014-95-1-1198. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

Bar-Shalom, Y., and Fortmann, T. 1988. *Tracking and Data Association.* Academic Press.

Carver, N., and Lesser, V. 1991. A New Framework for Sensor Interpretation: Planning to Resolve Sources of Uncertainty. In Proceedings of AAAI-91, 724–731.

Carver, N., Cvetanovic, Z., and Lesser, V. 1991. Sophisticated Cooperation in FA/C Distributed Problem Solving Systems. In Proceedings of AAAI-91, 191–198.

Carver, N., and Lesser, V. 1994. A First Step Toward the Formal Analysis of Solution Quality in FA/C Distributed Interpretation Systems. In Proceedings of the 13th International Workshop on Distributed Artificial Intelligence (also available as Technical Report 94-37, Department of Computer Science, University of Massachusetts).

Carver, N., and Lesser, V. 1995. The DRESUN Testbed for Research in FA/C Distributed Situation Assessment: Extensions to the Model of External Evidence. In Proceedings of the International Conference on Multiagent Systems (ICMAS), 33–40.

Carver, N., Lesser, V., and Whitehair, R. 1996. "Nearly Monotonic Problems: A Key to Effective FA/C Distributed Sensor Interpretation?," In Proceedings of AAAI-96, 88–95.

[10]See the discussion in (Lesser 1991) on *solution uncertainty* and *control uncertainty* for more information about coordination to achieve efficient, coherent global problem solving.

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

Lesser, V., and Corkill, D. 1983. The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem Solving Networks. *AI Magazine*, vol. 4, no. 3, 15–33.

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

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.