Examining Some Assumptions of the FA/C Distributed Problem-Solving Paradigm*

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

Abstract

The functionally-accurate, cooperative (FA/C) distributed problem-solving paradigm has been an important approach for organizing distributed problem solving among homogeneous, cooperating agents. They idea behind the FA/C approach is that agents can produce tentative, partial results based on only local information, and then exploit the constraints that exist among these local results to resolve uncertainties and global inconsistencies that result from the use of incomplete information. While this approach has been used in several implemented systems, there has been little formal analysis of the quality of the solutions that are produced by the approach or of the conditions that are necessary for the approach to be successful. This paper builds on work we have done to begin to formally analyze the quality of solutions that can be produced by FA/C systems, by examining some of the assumptions implicit in the approach. The analysis will be done in the context of distributed sensor interpretation.

1 Introduction

In the *functionally accurate, cooperative* (FA/C) paradigm for distributed problem solving [5, 7], agents need not have all the information necessary to completely and accurately solve their subproblems. Instead, agents produce tentative, partial results based on local information and then exchange these results with the other agents to resolve local uncertainties and global inconsistencies. The basic intuition behind this approach is that for many applications there exist (inter-agent) constraints among the subproblems, and these constraints can be exploited to resolve the inconsistencies and uncertainties that occur in local problem solving due to the lack of complete, accurate, and up-to-date information.

The FA/C approach has been important in *cooperative distributed problem solving* (CDPS) research. Several research systems that use an FA/C approach have been built (e.g., [2, 6]). However, until our recent work ([4]) there has never been any formal analysis of the conditions that are necessary for the approach to be successful or the quality of the solutions that can be produced. In that paper we presented two theorems that compared the quality of solutions produced by an FA/C distributed system to the solutions that would be produced by an equivalent centralized system, relative to certain aspects of the agent problem-solving and coordination strategies. We showed that there are conditions under

[0]*This work was supported in part 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.

which it is possible to guarantee that the distributed system produces a solution that is comparable to the centralized solution, and other conditions under which there is merely some probability of obtaining such a solution.[1]

In this paper we will examine some of the assumptions behind the FA/C model to begin to understand what problem characteristics are necessary for successful application of an FA/C approach. These assumptions or their consequences have not been made explicit in earlier work. For example, an assumption of the FA/C approach is that solutions can be produced without the need for "excessive" communication among the agents. This has certain implications, such as that "partial results" (e.g., data abstractions) from other agents can substitute for the raw data in detecting and resolving contradictions and uncertainties.

Since most FA/C applications have been in distributed sensor interpretation (e.g., distributed vehicle monitoring) this paper concentrates on that application. It is also important to point out that we are examining only some aspects of FA/C problem solving: solution quality and the amount of data that must be communicated among agents. In other words, the paper concentrates on what was referred to as *data uncertainty* in [5, 7].[2]

In the next section we describe our model of distributed sensor interpretation and the distributed problem-solving model that we assume for our analysis. Section 3 contains the discussion of the assumptions behind the FA/C model and their implications for problem characteristics. The paper concludes with a summary of our future research plans.

2 Distributed Sensor Interpretation

By *sensor interpretation* (SI), we mean the determination of high-level, conceptual explanations of sensor data. For example, in vehicle monitoring applications this involves tracking and identifying individual vehicles, and possibly determining their purpose. Our model of SI was described in [1, 4]: interpretation hypotheses are incrementally constructed via *abductive inferences*, based on a causal model that defines the relationships among the data types and abstraction types. An abductive inference identifies a possible *explanation* for a piece of data or a hypothesis, and conversely, the data/hypothesis provides *support* for the explanation hypothesis. Abductive inferences are uncertain inferences that provide *evidence* for hypotheses rather than conclusively proving them. The key source of uncertainty for any hypothesis is the possibility of alternative explanations for the data that supports the hypothesis.

[1]The theorems showed that if agents were locally doing *complete* evidence propagation (see Section 2 for an explanation of complete vs. incomplete evidence propagation) then there was a fairly simple global propagation strategy that would guarantee the distributed system would obtain an equivalent (though not necessarily) identical solution to what would be obtained by a centralized system, but that this was not the case when agents were locally doing *incomplete* evidence propagation. In the analysis, we assumed that "best" solutions were determined using a satisficing approach based on belief thresholds rather than being the optimal solution (the MPE), as mentioned in Section 2.

[2]We are currently pursuing both empirical and analytic approaches to address other issues in FA/C problem solving. For example, dealing with what was referred to as *control uncertainty* in [5, 7] in terms of coordination strategies for efficient FA/C problem solving. Another key issue in the design of FA/C systems is the role that agent architectures play in allowing a wide range of inconsistencies to be resolved without requiring excessive communication among the agents [2, 3].

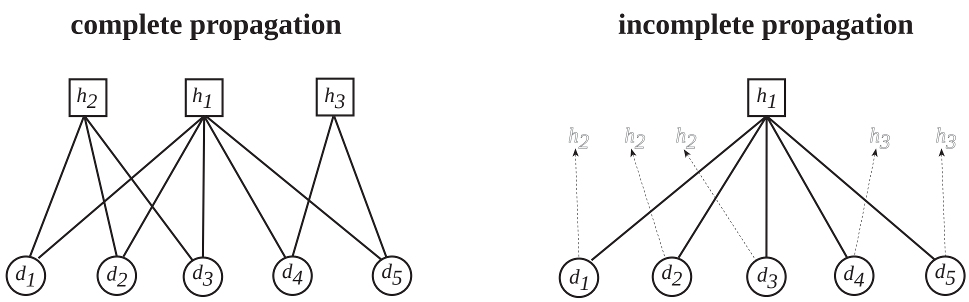

Figure 1: An example of incomplete evidence propagation.
In the complete propagation case, the system not only has created the most probable explanation, h_1, it also has created the alternative explanations, h_2 and h_3 (using the most complete support possible). This allows the system to determine the conditional probability of h_1 given the available data $\{d_1, \ldots, d_5\}$. In the incomplete propagation case, the alternative explanations for h_1 have not been created. This means that the belief computed for h_1 is only an approximation of the true conditional probability since the likelihood of the alternative explanations has not been correctly considered (h_1 is still uncertain, though, because the possibility of alternative type 2 and 3 explanations for each piece of supporting data is known, as are the a priori likelihoods of these explanations).

A *solution* to an SI problem is a set of hypotheses that explains the available data. In general, there will be multiple possible alternative (uncertain) solutions, and we want to find the "best" solution. Ideally, this should be the *most probable explanation* (MPE) [8] given *all* of the available data. The problem with this definition is that for many SI problems it is impractical to compute the MPE. This was explained in some detail in [4] (particularly as it relates to the differences between SI and the kinds of problems that are typically studied in research on abductive inference and probabilistic network inference).

The upshot of this is that SI systems usually must use heuristic, satisficing approaches to construct solutions. For instance, they may assemble solutions from hypotheses whose belief ratings surpass some *acceptance threshold* without being sure that are the most likely, and they usually do not process every piece of available data. These kinds of approaches result in solutions that are only approximations of the MPE. A key issue here is that incomplete hypothesis construction is equivalent to incomplete propagation (evaluation) of evidence. This is explained in Figure 1.

In a centralized SI system, all of the data is available to the single agent. In a distributed SI system, typically 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 and may in fact be inconsistent because they are based on different incomplete subsets of the data (see [2, 3]). Agent solutions are *interdependent* whenever data (evidence) for a hypothesis is spread among multiple agents or when agent areas of interest overlap as a result of overlapping sensor coverage.

These are exactly the kinds of distributed problems for which the FA/C approach was intended. This means that in an FA/C system, there must be some mechanism to identify interdependencies among the agents' hypotheses/solutions and cause appropriate communication between the agents. The DRESUN agent architecture [1, 2, 3] provides this capability,

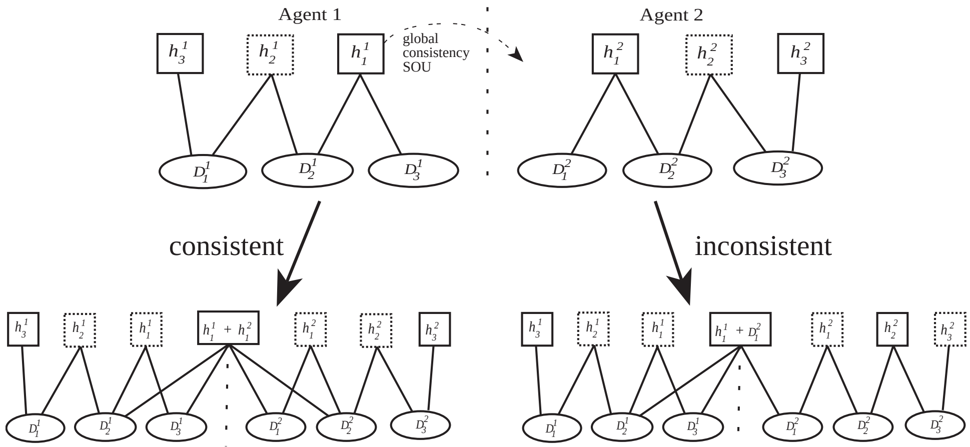

Figure 2: An example of the resolution of a global consistency SOU.
When there is a consistent explanation in the external agent, resolution of the global SOU associated with h_1^1 results in the creation of a merged hypothesis as a new alternative explanation in each agent. When the local hypothesis is inconsistent with hypotheses in the external agent, new alternatives may be created (as shown here). When the local hypothesis is inconsistent with the data in the external agent, new evidential links are created to represent the contradictory evidence.

and forms the basis for our model of the capabilities of an FA/C agent. The exchange of results and data associated with interdependent subproblems is a major component of the FA/C approach. In the DRESUN model, this is termed *resolving a global SOU*: exchanging information among the associated agents so as to effectively propagate evidence between their hypothesis (belief) networks. This is shown in Figure 2. Resolution of global SOUs is analogous to (intra-agent) evidence propagation, and as with evidence propagation there are a range of strategies that may be used to determine which global SOUs to pursue and how completely to propagate their effects.

3 Assumptions behind the FA/C Model

A key assumption of the FA/C approach is that a global solution can be produced without the need for "excessive" communication among the agents. It has never been specified exactly what "excessive" communication means, but one goal of a CDPS approach is improved performance in terms of reduced time to a solution relative to a centralized approach, and focusing on this goal allows us to draw some conclusions about what constitutes excessive communication.[3] One thing that is clear is that the FA/C model will impose substantial delays over a centralized model if agents require access to *all* of the globally available raw data in order to arrive at a global solution. This is the case because in this model agents can obtain data from external sensors only by communicating with the agents responsible for those sensors (either by explicitly requesting the data from those other agents or by waiting for the other agents to decide that this data needs to be sent).

[3]Note that there are other possible goals of a distributed approach, including reduced communcation costs/bandwidth; tighter coupling of sensors and processors; and increased reliability or graceful performance degradation.

Thus, a requirement for FA/C problem solving without excessive communication is that agents need access to only limited amounts of raw data from external sensors. This will be the case if just a small subset of each agent's subproblems interact with those of other agents—i.e., if agents' subproblems are largely independent. Of course, the amount of interdependency may vary substantially from situation to situation, and some level of communication may be required to determine the degree of interaction. However, even if subproblem interactions are consistently limited, it must still be possible to determine exactly what data is relevant to which other agents or from which other agents. Furthermore, this must be able to be done in a timely manner (i.e., not very far into local problem solving) and, again, without excessive communication.

Another way that this requirement for limited communication can be fulfilled is if the tentative, partial local results (e.g. abstractions of the data) can substitute for the raw data in determining the global solutions. This is certainly the view of the developers of the FA/C paradigm. For example, [5] refers to *consistency checking* of the tentative local solutions with results received from other nodes as "an important part of the FA/C approach." However, we have come to recognize that the processing of results/abstractions may have to be considerably more sophisticated than suggested above. For instance, "consistency" of local solutions may not provide any guarantees about the quality of the merged, global solution. It is entirely possible to have $P(H_a \mid D_1) > P(H_b \mid D_1)$ and $P(H_a \mid D_2) > P(H_b \mid D_2)$, but $P(H_a \mid D_1, D_2) < P(H_b \mid D_1, D_2)$ (where H_a and H_b are competing hypotheses and D_1 and D_2 are data sets in different agents). In other words, the solution that is locally most likely in both agents (H_a) may not be the globally most likely solution—even though the local solutions are surely consistent.

Thus, an important element of understanding the applicability of the FA/C paradigm is to determine when we can communicate (mainly) the partial solutions and still obtain the globally most likely solution (or at least a satisfactory approximation of the optimal solution), and to be clear on how exchanged solutions must be processed to achieve this result. One reason that the situation is not quite as difficult as is suggested above is that many real world domains have characteristics that allow this approach to succeed. For example, in vehicle monitoring, a considerable amount of evidence is required to achieve a vehicle track hypothesis with high belief (probability). While it is certainly possible for additional evidence to negatively affect the belief in this track, it will take substantial additional evidence to have a major effect on the degree of belief. In other words, while these beliefs are nonmonotonic with increasing evidence, in some situations the effects of additional evidence are not totally unpredictable.

4 Conclusions

In this paper we have discussed the fact that the successful application of the FA/C paradigm for CDPS relies on several assumptions that have not previously been made explicit nor subjected to formal analysis. In particular, we have shown that successful application of an FA/C approach requires the ability to make use of the local partial solutions (data abstractions) in assessing global solution quality, since communication of raw data among agents must be fairly limited for reasonable performance (relative to a centralized system).

We are currently working to provide a more formal analysis of the conditions under which data abstractions can be useful in FA/C applications and to be able to characterize the average/expected amount of information that must be transmitted among FA/C agents. We are also currently extending our theorems from [4] by deriving specific probabilities in connection with the theorems as a function of domain models and acceptance thresholds; understanding how to formalize the effects of different evidence propagation and coordination strategies; extending the analyses to deal with processing of only subsets of the available data; and generalizing to CDPS tasks other than sensor interpretation. This paper represents another preliminary step in formally analyzing the FA/C paradigm for CDPS.

Acknowledgements: I wish to thank Victor Lesser for discussions about the work reported here and other work related to formal analysis of the FA/C paradigm.

References

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

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

[3] Norman Carver, Victor Lesser, and Qiegang Long, "Resolving Global Inconsistency in Distributed Sensor Interpretation: Modeling Agent Interpretations in DRESUN," *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, May, 1993.

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

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

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

[7] 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.

[8] Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.