

Nearly Monotonic Problems: A Key to Effective FA/C Distributed Sensor Interpretation?

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

Victor Lesser and Robert Whitehair
Computer Science Department
University of Massachusetts
Amherst, MA 01003
(lesser or whitehair @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. A key assumption of the FA/C model has been that the agents' local solutions can substitute for the raw data in determining the global solutions. This is not the case in general, however. Does this mean that researchers' intuitions have been wrong and/or that FA/C problem solving is not likely to be effective? We suggest that some domains have a characteristic that can account for the success of exchanging mainly local solutions. We call such problems *nearly monotonic*. This concept is discussed in the context of FA/C-based distributed sensor interpretation.

Introduction

The *functionally accurate, cooperative* (FA/C) distributed problem-solving paradigm (Lesser & Corkill 1981; Lesser 1991) has been important in *cooperative distributed problem solving* (CDPS) research. Several FA/C-based research systems have been built (e.g., (Carver, Cvetanovic, & Lesser 1991; Carver & Lesser 1995a; Lesser & Corkill 1983)). However, until some recent work of ours (Carver & Lesser 1994; Carver 1995; Carver & Lesser 1995b), there had never been any formal analysis of the conditions that are necessary for an FA/C approach to be successful or of the potential performance of FA/C systems. In this paper we examine some of the assumptions behind the FA/C model and look at a problem domain characteristic that can make the FA/C approach successful.

The development of the FA/C model was motivated by the recognition that in many CDPS domains it is impractical/impossible to decompose problems and/or transfer data so that individual agents work on independent subproblems. FA/C agents are designed to produce tentative, partial solutions based on only local information (which may be incomplete, uncertain, or inaccurate). They then exchange these results with the other agents, exploiting inter-agent constraints among the subproblems to resolve uncertainties and inconsistencies due to the deficiencies in their local information.

A critical issue for the FA/C model is whether high quality global solutions can be produced without the need for "excessive" communication among the agents (when exchanging and "integrating" local results). Most FA/C work has assumed that this is the case because it has been assumed that the local partial solutions can substitute for the raw data in resolving contradictions and uncertainties. Unfortunately, this is not true in general. Does this mean that researchers' intuitions have been wrong and/or that FA/C problem solving is not likely to be effective?

In this paper we suggest that some domains have a characteristic that justifies the role of local solutions in producing global solutions (at least for approximate, satisficing problem solving). We call such problems *nearly monotonic*. Basically, while belief and/or solution membership may be nonmonotonic with increasing evidence, in nearly monotonic problems they become nearly monotonic once certain conditions (like fairly high belief) are reached.

This paper discusses the FA/C model and the concept of nearly monotonic problems in the context of distributed *sensor interpretation* (SI). We concentrate on this domain because most FA/C applications have been in distributed SI (particularly distributed vehicle monitoring), we are engaged in related research on the FA/C model for distributed SI, and we have available a new analysis tool for SI. SI is also a very complex problem and distributed SI plays an important role in many situation assessment (decision support) systems.

The next section briefly describes distributed SI. The FA/C Issues section examines the use of local agent solutions to determine global solutions in FA/C problem solving. What we mean by near monotonicity is expanded on in the Nearly Monotonic Problems section. This is followed by a section that examines the role near monotonicity can play in developing coordination strategies for FA/C-based SI. The Analytic Support section provides some support for the existence of nearly monotonic SI problems using a framework for analyzing SI domains. The paper concludes with a summary of our conclusions and future plans.

Distributed Sensor Interpretation

By sensor interpretation, we mean the determination of high-level, conceptual *explanations* of sensor and related data. For example, vehicle monitoring applications involve tracking and identifying vehicles, and possibly determining the purpose of individual vehicles and patterns of vehicles. The model of SI that we assume is essentially that described in (Carver & Lesser 1991; Carver & Lesser 1994). An *interpretation* of a data set is an explanation of what caused all of the data. Typically it will be a *composite* of a set of hypotheses whose types are from a specified subset of the abstraction types (the *explanation corpus* (Pearl 1988)), each of which explains some subset of the data, and which together explain all of the data. In general, there will be multiple possible interpretations of a data set.

A *solution* to an SI problem is an interpretation of the available data that is judged “best” according to some criteria. One possible definition of best is the *most probable explanation* (MPE) (Pearl 1988) given the available data. The problem with this definition is that for many SI problems it is impractical to compute the MPE or exact belief ratings (conditional probabilities). (Carver & Lesser 1994) contains an explanation of this issue.¹ We will simply assume here that even centralized SI systems usually must use approximate, satisficing approaches to construct solutions (so solutions are only approximations of the MPE).²

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,” so agents must somehow combine their data in order to construct a *global solution*.

FA/C Issues

As we have said, a critical issue for the FA/C approach is whether high quality global solutions can be produced without the need for “excessive” communication among the agents.³ Because FA/C agents work on possibly interdependent local solutions, they must exchange and integrate these solutions to construct a global solution.⁴ Integrating local solutions may not be straightforward, however, because these solutions may

be incomplete and/or highly uncertain, and because solutions from different agents may be inconsistent (since they are based on different incomplete data subsets). Conditions under which possibly interdependent local solutions can be efficiently integrated is the main issue to be addressed in this paper.

The FA/C model can impose substantial delays over a centralized model if the determination of a global solution requires some agent(s) to have access to *all/most* of the globally available raw data. This would happen, for instance, if interrelated local solutions could be integrated only with access to their supporting raw data and nearly all solutions were interrelated. Delays would result because FA/C agents obtain data from external sensors by communicating with the agents responsible for those sensors as it becomes clear that the data is needed (they either explicitly request data from other agents or wait for the other agents to decide that the data needs to be sent).

Thus, effective FA/C-based SI requires that agents need access to limited amounts of raw data from external sensors. There are two basic ways in which this requirement may be met: (1) only a small subset of each agent’s subproblems interact with those of other agents (agents’ subproblems are largely independent of those in other agents) or (2) local solutions can be integrated with limited need for their supporting raw data. We focus on the second approach in this paper.

Subproblem independence is problematic for SI, since in many domains there is no way to determine a priori whether two pieces of data are interrelated and, in fact, virtually any two pieces of data may have a non-zero relationship.⁵ For example, in a situation analysis application for tactical air command, targets hundreds of miles apart may be interrelated since they might be acting in concert as part of some scenario/pattern. This means that even widely separated pieces of sensor data are potentially interrelated. Furthermore, even where subproblem interactions are consistently limited, it must be possible to determine what data is relevant to which other agents, and this must be able to be done in a timely manner, without excessive communication.

We are interested in understanding whether or when local solutions can substitute for the raw data in determining the global solutions (where there are inter-agent subproblem interrelationships). If agents can transmit mainly solution-level hypotheses rather than raw data, then communication can be greatly reduced. Interpre-

(Carver, Cvetanovic, & Lesser 1991; Carver & Lesser 1995b) describes how it is possible to identify interdependencies with SI applications. Agent solutions are *interdependent* whenever data (evidence) for a hypothesis is spread among multiple agents or when agent “interest areas” overlap as a result of overlapping sensor coverage.

⁵When we speak about data being interrelated and about the “strength” of this relationship, we mean evidential relationships: the presence, absence, characteristics, or interpretations of data can affect the belief in the possible interpretations of other data.

¹(Pearl 1988) contains a discussion of some of the limitations of using the MPE even when it can be computed. For example, there may be very different utilities for identifying (or failing to identify) enemy vs. friendly aircraft.

²While it is nearly always necessary to trade-off solution quality for efficiency in SI problems, approximation can be done in a variety of ways—see (Bar-Shalom & Fortmann 1988; Carver & Lesser 1995b, Cox & Leonard 1994).

³Precisely what constitutes *excessive* communication will depend on the reasons for taking a distributed approach to problem solving.

⁴FA/C agents must have some mechanism to identify interdependencies among their solutions/hypotheses.

tation hypotheses are abstractions of the sensor data and can generally be represented using a fraction of the storage that would be required for their supporting data. Communication of solution hypotheses should also require receiving agents to do less processing (interpretation/probabilistic inference) than would be required with raw data. In addition, the number of possible interpretations of a data set can be very large, focusing on the agents' local solutions can greatly reduce communications.

The developers of the FA/C paradigm certainly believed that local solutions (and other abstract hypotheses, or "results") could substitute for the raw data in determining global solutions. (Lesser & Corkill 1981) referred to "consistency checking" of the tentative local solutions with results received from other nodes as "an important part of the FA/C approach." When there were inter-agent subproblem interactions, agents would transmit their local solutions and check the consistency of the related components of these solutions. *Consistent* solution components would be "integrated" using only the abstract hypotheses (not the raw data), while inconsistencies would trigger additional communication. We will refer to this basic procedure for developing a global solution as the *consistent local solutions strategy*.

This strategy has the potential to reduce communications because when local solutions are consistent they are integrated without requiring transmission of supporting raw data. Does the strategy produce high quality global solutions? To answer this question we must first consider what it should mean to integrate local solutions into a global solution. Assume that there are two agents, A_1 and A_2 , with local data sets D_1 and D_2 , respectively. Each agent's local solution would be the "best" interpretation of its own local data set (using some common definition of best). Now for the global solution, what we would ideally like is the best interpretation of the joint, global data set ($D_1 \cup D_2$), using the same definition of best interpretation. This is ideal because the distributed system would then produce the same solutions as an equivalent centralized system and solutions would not vary simply with differences in the distribution of the data among the group of agents.

Given this standard for global solutions, what can we say about the "consistent local solutions strategy?" Unfortunately, what we can say is that, in general, it provides *no guarantees at all* about the quality of the global solution. Again, consider a situation with two agents, and suppose that there are the same two alternative interpretations I_a and I_b for each of the data sets D_1 and D_2 . It is entirely possible to have $P(I_a | D_1) > P(I_b | D_1)$ and $P(I_a | D_2) > P(I_b | D_2)$, but $P(I_a | D_1, D_2) < P(I_b | D_1, D_2)$. In other words, even though interpretation I_a is the most likely (best) solution given each agent's local data set separately, it may not be the globally most likely solution even though the local solutions are consistent (here identical). Likewise, if $P(H | D_1) > threshold$ and $P(H | D_2) > threshold$, it is not necessarily the case

that $P(H | D_1, D_2) > threshold$ (where H is an interpretation hypothesis being selected for membership in the approximate solution based on its belief surpassing some acceptance threshold).

These are unavoidable consequences of the *nonmonotonicity* of domains like SI. The upshot of such observations is that integration of even consistent interrelated local solutions can require that agents recompute hypothesis beliefs and redetermine best solutions—just as with inconsistent local solutions. In some cases, this can require one agent to have complete knowledge of the other agent's raw data and evidential information (alternative interpretation hypotheses and their interrelationships).⁶

Nearly Monotonic Problems

We believe that one explanation for the apparent success of FA/C-based SI is that many SI domains have a property that makes the "consistent local solutions strategy" appropriate and effective. We have termed problems with this property *nearly monotonic*, because they nearly behave as if they are monotonic once certain conditions have been achieved. For example, while additional evidence can negatively affect the belief in a vehicle track hypothesis, once a fairly high degree-of-belief is attained, it is unlikely that the belief will change significantly and it is unlikely that the hypothesis will not be part of the best global solution. Thus, while the domain is nonmonotonic in a strict sense, the effects of additional evidence are not totally unpredictable: solution components with particular attributes (e.g., high belief) are unlikely to be affected as additional evidence is considered.

To proceed in examining near monotonicity, the following notation will be used: \mathcal{D} is the complete, globally available data set; D is some subset of \mathcal{D} that currently has been processed by an agent; $BEL(H)$ is the current belief in hypothesis H given data set D (it is $P(H | D)$); $BEL^*(H)$ is the "true" belief in hypothesis H for data set \mathcal{D} (it is $P(H | \mathcal{D})$); MPE is the current MPE solution given data set D ; and MPE^* is the "true" MPE solution for data set \mathcal{D} .

It is impossible to give a single, precise definition of what a nearly monotonic problem is. What we will present are several formulas for statistically characterizing SI problems, which can be useful in assessing and using near monotonicity. The basic approach will be to characterize ultimate properties of interpretation hypotheses if/when \mathcal{D} were processed, given current characteristics based on partial data/evidence. The ulti-

⁶A detailed explanation of the recomputation of belief and solution membership for SI is beyond the scope of this paper. In belief network terms, think of the integration of local solutions in different agents as establishing new evidential links between the agents' belief networks and then recomputing beliefs by appropriate evidence propagation. This may require complete knowledge of another agent's data because recomputation in SI problems cannot in general be done by message passing (Carver & Lesser 1995b).

mate hypothesis properties that are of interest are belief and solution membership. We have considered five possible characterizations for near monotonicity:

1. a conditional probability density function (conditional pdf) $f_{BEL^*|x}(b)$ that describes the probability that the ultimate belief in hypothesis H is b given that the current hypothesis belief is x , defined such that $\int_{p_1}^{p_2} f_{BEL^*|x}(b)db = P(p_1 \leq BEL^*(H) \leq p_2 \mid BEL(H) = x, \dots)$;
2. the probability that the ultimate belief in the hypothesis will be greater than or equal to its current belief, $P(BEL^*(H) \geq BEL(H) \mid BEL(H) = x, \dots)$;
3. the probability that the ultimate belief in the hypothesis will be greater than or equal to some specified level l , $P(BEL^*(H) \geq l \mid BEL(H) = x, \dots)$;
4. the probability that the hypothesis will ultimately be in the MPE, $P(H \in MPE^* \mid BEL(H) = x, \dots)$;
5. the probability that the hypothesis will eventually be in the MPE given that it is in the current MPE, $P(H \in MPE^* \mid BEL(H) = x, H \in MPE, \dots)$.

For each of these characterizations, being nearly monotonic would require that once an interpretation hypothesis has certain characteristics (based on only a portion of the available data) then the probabilities will be *high enough* to make it appropriate to assume the hypothesis is in the solution. For example, using formula 4, we would like something along the lines of: once a vehicle track hypothesis reaches a belief level of 0.8, the probability that it is in MPE^* is greater than 0.95. This would allow us to use the consistency of such a local hypothesis with another agent's local solution to conclude with high confidence that the track is in the global solution.

Which of the above characterizations is most appropriate will depend on: (1) the domain and its characteristics; (2) the statistical information that is available; and (3) the solution selection strategy. The probabilities in formulas 2 and 3 can be derived from the pdf of formula 1, but are included because detailed knowledge such as the pdf may not always be practical to obtain. Variations on these formulas can result from the use of approximate beliefs and solutions, rather than the exact ones used here.

We are exploring what hypothesis characteristics should be conditioning factors in the above formulas. Again, this will depend on the particular problem domain, as the predictiveness of different characteristics is likely to vary across domains, and systems vary in their solution quality requirements. From our experience, it appears that for SI problems both hypothesis belief and hypothesis *type* are important factors. Another possibility is the "quantity" of data supporting the hypothesis or the fraction of the overall data that has been processed. The RESUN SI framework (Carver & Lesser 1991) also provides detailed information about the reasons for uncertainty in beliefs, and such information may be necessary to identify hypotheses that are reliable enough to be assumed for global solutions.

Solution Quality and Coordination

Nearly monotonic problem domains are of interest for CDPS because they can make it possible to produce high quality global solutions with low communication costs. Near monotonicity means that consistency of local solutions can be *highly predictive* that the merged solution would be the best global solution. In this section we will examine the issue of solution quality when using the "consistent local solutions strategy." We will also discuss the trade-offs involved in developing FA/C coordination strategies to take advantage of near monotonicity.

MPE^* is one possible standard to use in evaluating the quality of a global solution S^G produced by an FA/C-based SI system. S^G could be compared against MPE^* in terms of $P(S^G = MPE^*)$, however this is often not the most meaningful metric for SI problems. First, S^G will tend to be incomplete ($S^G \subset MPE^*$) if it is based on incomplete data (a subset of \mathcal{D}), but these missing hypotheses are not important in SI applications if the data is selected appropriately (i.e., we care about targets, but not "noise"). Second, the likelihood of individual hypotheses being correct is more useful than the likelihood of the complete set of hypotheses being correct, because it tends to be the individual hypotheses (e.g., vehicles) which we must decide whether to respond to rather than the entire set of solution hypotheses. Because of these factors, in judging solution quality we will consider $P(H \in MPE^* \mid H \in S^G)$.

To produce solution quality results, we first need to better define what "consistency" of local solutions means and what it means to "integrate" local solutions to produce a global solution. Our definition of consistency of local solutions is an evidential one: solutions are consistent if hypotheses that comprise each of the local solutions are pairwise identical, independent, or corroborative. Two local solutions are inconsistent when any of their component hypotheses are contradictory (i.e., have a negative evidential relationship).

Hypotheses can be corroborative in either of two ways: when one is evidence for the other (one explains the other and the other supports the one), or when they are of the same type and can be *merged* into a single consistent hypothesis. Merging typically involves producing a single "more complete" hypothesis from two or more "less complete" hypotheses. For instance, two partial vehicle track hypotheses may be merged into a longer track hypothesis. While the resulting hypothesis could always be built from scratch from the combined supporting data/evidence of the component hypotheses, when we refer to the "merging" of hypotheses we will assume that this is done from the solution hypotheses, without reference to their supporting data.⁷ While

⁷The Distributed Vehicle Monitoring Testbed (DVMT) (Durfee & Lesser 1987; Lesser & Corkill 1983) had "merge" operators that did exactly this. DRESUN (Carver & Lesser 1995a) allows hypotheses to be "external evidence" for hypotheses of the same type.

this is clearly more efficient, in general beliefs for combination hypotheses cannot be precisely computed in this way (i.e., without access to the supporting data).

The last thing that needs to be done to produce solution quality results is to provide a more complete definition of what we mean by the “consistent local solutions strategy” for developing global solutions:⁸

1. each agent first uses only its own local data to develop a (local) solution;
2. upon satisfying some solution criteria, an agent communicates its solution’s abstract hypotheses to all agents with which it has recognized subproblem interactions;
3. the agent also sends its solution to any agents from which it has received solutions that were not included in step 2 and it continues to do so as solutions are received from any such additional agents;
4. the agent now proceeds to integrate its solution, one-by-one, with each of the solutions it has received and may yet receive prior to termination;
5. processing terminates when all agents have transmitted their solutions according to steps 2 and 3, and have integrated their solution with all received solutions;
6. the global solution is simply the union of all the final, integrated agent solutions.

If two agents’ local solutions are consistent when they are exchanged, then the integrated solution will be as described above in our discussion of consistency and merging: independent hypotheses will be added to the joint solution and corroborative hypotheses will be linked or merged. If agents’ local solutions are inconsistent when they are exchanged then the agents will be forced to engage in further communication (possibly involving raw data) to resolve the contradictions and determine the “best” joint solution.

Solution Quality Theorem: To derive some results about solution quality, we will make the following assumptions:

- We have available statistical information as in formula 4 in the previous section, and this probability is well correlated with hypothesis type and belief (so no additional conditioning factors are needed)
- Thus, we have $P(H \in MPE^* \mid type(H), BEL(H))$, which we will refer to as $P_{MPE^*}(H)$.
- Agents use the “consistent local solutions strategy” described above.
- Agents compute $BEL(H) = P(H \mid D)$ for the subset D of their own local data that they process before transmitting their solutions.
- In the case of inconsistent local solutions, the agents

involved compute the MPE joint partial solution (based on the data they have jointly processed).⁹

Under these conditions, what we can say about the quality of the resulting global solution is that $\forall H : P(H \in MPE^ \mid H \in S^G) \geq P_{MPE^*}(h_{max})$. h_{max} is simply H , unless H resulted from the merging of consistent hypotheses (identical or of the same type). In this second case, h_{max} is the hypothesis with the maximum belief out of all the hypotheses merged to produce H (e.g., if H resulted from the combination of h_1 from A_1 and h_2 from A_2 , and $BEL(h_1) \geq BEL(h_2)$, then $h_{max} = h_1$).*

Proof: Under the specified strategy, approximations will occur only when agents compute their local solutions, exchange them, and they are consistent. If they are inconsistent then the agents will engage in further communication to find the MPE solution to their joint data sets. When solutions are consistent, no further exchange of information will take place, and the joint solution will be the “merge” of the consistent solutions. Suppose that agent A_1 ’s solution is S_1 and agent A_2 ’s solution is S_2 , and they are consistent. If hypothesis H is in S_1 then either it is (1) independent of every hypothesis $H_i \in S_2$; (2) identical to some hypothesis $H_i \in S_2$; or (3) corroborative with one or more hypotheses, say $\{H_j\} \subset S_2$ (of the same or of different types as H). If H is independent of all $H_i \in S_2$ then $BEL^{A_1, A_2}(H) = BEL^{A_1}(H)$, so $P_{MPE^*}(H)$ is based on the local belief $BEL^{A_1}(H)$ computed by A_1 .¹⁰ If H is identical with some $H_i \in S_2$, $BEL^{A_1, A_2}(H) \geq maximum(BEL^{A_1}(H), BEL^{A_2}(H_i))$. Since $P_{MPE^*}(H)$ will be *monotonically nondecreasing* with increasing hypothesis belief, $P_{MPE^*}(H)$ following the merge must be greater than or equal to its value from either agent’s local data. If H is corroborative with hypotheses in S_2 then either (1) it is supporting or explanatory for these hypotheses, or (2) it can be merged with a hypothesis of the same type. In the first case, $BEL^{A_1, A_2}(H) \geq BEL^{A_1}(H)$ by the definition of being corroborative. Thus, $P_{MPE^*}(H)$ following the merge must be greater than or equal to its value from A_1 ’s local data. Now we must deal with corroborative hypotheses of the same type. Suppose that H is the result of merging h_1 from A_1 and h_2 from A_2 . We must have $BEL(H) \geq BEL(h_1)$, $BEL(H) \geq BEL(h_2)$, and so $BEL(H) \geq maximum(BEL(h_1), BEL(h_2))$, by our definition of corroborative hypotheses. Thus, we would have $P_{MPE^*}(H) \geq maximum(P_{MPE^*}(h_1), P_{MPE^*}(h_2))$.

What this theorem tells us is that we can use the “consistent local solutions strategy” and *potentially* get a global solution whose components are as likely to be

⁸Our description of the strategy is intended to be clear for the analysis—it is not intended to represent the way one would actually implement the strategy. It does not worry about agents duplicating work when integrating solutions, nor how to efficiently produce a final global solution, and so forth. Also, it does not worry about trade-offs in FA/C problem solving, discussed below.

⁹This last assumption is included to simplify the theorem, but it can easily be relaxed to allow the agents to select a non-MPE joint solution, as long as they do compute proper conditional probabilities. If the assumption is relaxed, this results in additional approximations in the global solution.

¹⁰The superscripts A_1 and A_2 denote which agent’s belief we mean. The superscript A_1, A_2 denotes the merged result.

in the MPE global solution as we desire (by selecting appropriate criteria that local solutions must meet prior to being exchanged). This is a very useful result even though we are not guaranteed to produce the best global solution under this strategy, since some type of approximation is required for most SI problems.

Of course, being able to use this strategy to efficiently achieve a desired likelihood, depends on two things being true: (1) agents can produce local solutions whose hypotheses have high enough belief, and (2) local solutions are largely consistent. This suggests that effective use of the “consistent local solutions strategy” requires appropriate structuring of the distributed system.¹¹ Usually at least one agent must have sufficient data to produce a high belief version of each solution hypothesis. Agents must also have enough overlap in their data coverage that it is unlikely that they will produce inconsistent solutions. When these conditions are not met, the agents may be forced to communicate considerably more information/data among themselves in order to produce a global solution of the desired quality.

For FA/C-based SI with limited communication, it is clearly advantageous to understand whether the domain is nearly monotonic or not, and if it is to design coordination strategies to capitalize on this property. Still, the design of a coordination strategy must consider numerous trade-offs. For instance, to take maximum advantage of a problem being nearly monotonic, agents should try to produce appropriate (nearly monotonic) interpretation hypotheses based on their local data and only then (or when it is found that this cannot be done) exchange them with other agents. The problem with this approach is that while it will minimize the communication of raw data among the agents, it may not produce the best performance in terms of time to reach a solution. This is because agents may not be able to produce nearly monotonic solution hypotheses from their data and their local solutions may not be consistent even if they can. Should agents fail to produce nearly monotonic solution hypotheses and/or produce inconsistent solutions then raw data generally will have to be communicated and processed by some agents. If the need to do this is discovered only after a significant amount of processing time, then production of the ultimate solution will be delayed. In this type of situation, where agents require “constraint information” from other agents, it is advantageous to receive this information as early in processing as possible.

Analytic Support

While we have demonstrated that nearly monotonic problems have the potential to support efficient FA/C-

¹¹These basic requirements were noted in (Lesser 1991): “some qualitative intuitions on...requirements for the use of the FA/C paradigm: local partial solutions are valid sufficiently often to seed system-wide problem solving with enough correct constraints....”

- P.1.1. $\text{Start} \rightarrow \text{Track1}_t$ ($p = 0.7$)
- P.1.2. $\text{Start} \rightarrow \text{Track2}_t$ ($p = 0.3$)
- P.2.0. $\text{Track1}_t \rightarrow V1_t N1_t \text{Track1}_{t+1}$ ($p = 1.0$)
- P.3.0. $\text{Track2}_t \rightarrow V2_t N2_t \text{Track2}_{t+1}$ ($p = 1.0$)
- P.4.1. $V1_t \rightarrow S2_t$ ($p = 0.5$)
- P.4.2. $V1_t \rightarrow S3_t S5_t$ ($p = 0.5$)
- P.5.1. $V2_t \rightarrow S3_t$ ($p = 0.5$)
- P.5.2. $V2_t \rightarrow S2_t S4_t$ ($p = 0.5$)
- P.6.1. $N_t \rightarrow S4_t$ ($p = 0.5$)
- P.6.2. $N_t \rightarrow S5_t$ ($p = 0.3$)
- P.6.3. $N_t \rightarrow \text{lambda}$ ($p = 0.2$)

Figure 1: Simple vehicle tracking IDP grammar.

based distributed SI, we have not yet shown that real-world SI problems are indeed nearly monotonic. This could be done by taking data sets, determining the correct global solutions in a centralized fashion, and then collecting the necessary statistics by selecting random subsets of these data sets, interpreting them, and analyzing the (partial) solutions relative to the global solutions. As a first step toward this eventual goal, we have instead made use of a recently developed framework for analyzing SI domains and problem solvers to provide some support for the concept of nearly monotonic problems.

Complex SI problems can be represented and analyzed using the *Interpretation Decision Problem* (IDP) formalism (Whitehair & Lesser 1993; Whitehair 1996). In the IDP formalism, the structure of both problem domains and problem solvers is represented in terms of context free attribute grammars and functions associated with the production rules of the grammars. The formalism has been used to analyze a variety of simulated SI domains and SI problem-solving architectures. For example, grammars have been constructed that represent the SI domains and goal-directed blackboard architecture used in the Distributed Vehicle Monitoring Testbed (DVMT) (Durfee & Lesser 1987; Lesser & Corkill 1983).

We will first use a very simple vehicle tracking grammar to illustrate a nearly monotonic SI domain. In the problem domain defined by the grammar rules in Figure 1, there are two kinds of vehicles, V1 and V2. The nonterminals Track1 and Track2 correspond to vehicle tracks of these two types, respectively. The terminal symbols in this grammar, S2, S3, S4, and S5, correspond to actual sensor data generated by the moving vehicles. The nonterminal N represents random noise in the environment. The terminal “lambda” that appears in rule P.6.3. corresponds to an empty string. The subscripts, $t+n$, correspond to the time in which an event occurs. The nonterminals Track1 and Track2 are the potential solutions for problem instances generated with this grammar. Grammars of this form are referred to as *Generational IDP Grammars*, (IDP_g) (Whitehair 1996). An IDP_g generates a specific problem instance using the probabilities that are shown in parentheses

with each rule. For example, given a nonterminal $V1$, IDP_g will generate an $S2$ with probability 0.5, or it will generate an $S3$ and $S5$ with probability 0.5. We refer to these probabilities as the grammar's distribution function.

This example grammar is important because it illustrates the relationships that can lead to nearly monotonic, complex domains. For example, consider a situation where the sensor data "S2 S4" is observed. This data is ambiguous because it could have been generated by either a $Track1$ or a $Track2$. A $Track1$ could have generated "S2 S4" by generating a $V1$ and an N , which would have generated an $S2$ and an $S4$ respectively. The probability of a $V1$ generating an $S2$ is 0.5 and the probability of an N generating an $S4$ is 0.5. A $Track2$ could have generated this data by generating a $V2$, which would have generated "S2 S4" with probability 0.5.

Thus, given the possible interpretations $V1$ and $V2$ for "S2 S4," and given that $BEL(H)$ is the problem solver's belief in interpretation hypothesis H , the values $P(V1 \in MPE^* | BEL(V1))$ and $P(V2 \in MPE^* | BEL(V2))$ are approximately equal for any values of $BEL(V1)$ and $BEL(V2)$. This means that it is not possible to use the beliefs in $V1$ and $V2$ to differentiate between a $Track1$ and a $Track2$ interpretation of "S2 S4." Thus, the domain is nonmonotonic.

On the other hand, consider the sequence of vehicle (position) hypotheses, " $V2_t V2_{t+1} V2_{t+2} V2_{t+3}$." Each of these $V2$ hypotheses would be associated with either an "S2 S4" or an $S3$. If a $V2$ explains an $S3$, then the probability that the partial results are from a $Track2$ is very high (actually it is a certainty). However, if a $V2$ explains an "S2 S4" it is possible that the sensor data was really generated by " $V1 \rightarrow S2$ " and " $N \rightarrow S4$." However, as the length of the *track* (vehicle position sequence) " $V2_t V2_{t+1} V2_{t+2} V2_{t+3} \dots$ " increases, it becomes more and more likely that the MPE of the data is a $Track2$.¹² In other words, as the length of the track increases, $P(track \in MPE | BEL(track))$ becomes very high.

This observation is intended to illustrate the following point. In complex, real-world domains, it is often the case that certain "equivalence classes" of partial results exhibit behavior that is nearly monotonic. In the vehicle tracking domain, this occurs for the equivalence class of partial tracks when the partial tracks extend over a significant number of time periods. More than likely, this phenomenon holds in other domains as well. For example, as the length of an interpreted fragment of speech increases, it is likely that the associated

¹²As the length of the partial track of $V2$ hypotheses increases, the probability that it was generated by a $Track1$ and noise decreases. If the probability of a single "S2 S4" being generated by a $V1$ (i.e., by a $Track1$) is 0.25 (as in the example grammar), then the probability of two "S2 S4" events occurring sequentially is $0.25 * 0.25$. The probability of three such events occurring sequentially is $0.25 * 0.25 * 0.25$. And so forth.

equivalence class of partial results will exhibit nearly monotonic properties.

As a further demonstration that this phenomena occurs in complex domains, we used the IDP framework to collect statistics for a more complex vehicle tracking problem domain (Whitehair 1996). This grammar modeled all of the phenomena that have been studied in the $DVMT$, plus some additional factors. In this domain, we defined three different interpretation types: group level (GL), vehicle level (VL), and partial tracks of length 4 (PT). We then accumulated statistics using the grammar. Problem instances were repeatedly generated and their MPE^* interpretations determined (i.e., the problem was solved). For each such cycle, we recorded the *credibilities*¹³ of any instances (hypotheses) of each of the interpretation types. However, we also divided each of the resulting sets into those hypotheses that were subsequently used in MPE^* and those that were not.

What we found was that the GL and VL types were nonmonotonic with respect to credibility. In other words, the credibility in a GL or VL hypothesis was not a good indicator of whether or not the hypothesis was an element of MPE^* . The distribution of *credibility*(H) was approximately the same for all GL and VL hypotheses, regardless of whether or not they were actually part of MPE^* and $P(H \in MPE^* | type(H) = vl, credibility(H) = 0.7)$ was only 0.32. On the other hand, for PT hypotheses there was a strong correlation between credibility and membership in MPE^* : if a partial track of length 4 had a fairly high credibility it was very likely to be part of MPE^* . For example, $P(H \in MPE^* | type(H) = PT, credibility(H) = 0.7) = 0.92$, while $P(H \in MPE^* | type(H) = PT, credibility(H) = 0.55) = 0.5$. Thus, PT hypotheses were nearly monotonic (in terms of solution membership) if they achieved a reasonable credibility.

Conclusion

In this paper, we have shown that while consistency checking of local agent solutions has been used in previous FA/C -based SI systems, this strategy cannot necessarily produce high quality global solutions. However, we have also shown that in certain domains the strategy can be used to efficiently find approximate, satisficing solutions. In particular, problems that we call nearly monotonic allow for consistency checking of local solutions to produce reasonable global solutions—when the local solutions meet certain criteria.

This work furthers our understanding of when the FA/C model is appropriate for distributed SI and what appropriate coordination strategies are. Much remains to be done, however. The importance of a problem

¹³The IDP analysis tools currently compute a *credibility* rating for each hypothesis. This is not exactly what we have termed belief (the conditional probability of the hypothesis). $Credibility(H) \geq BEL(H)$.

being nearly monotonic remains an open issue. Near monotonicity can support efficient FA/C problem solving, but does not alone guarantee efficiency, and FA/C-based SI can be efficient even without the property: local solutions may frequently be inconsistent or local data may provide insufficient belief to make use of the property, and in some cases (even inconsistent) local solutions can be integrated with limited communication of raw data—particularly if we need only approximate global solutions. The focus of our future research on near monotonicity will be assessing whether real-world SI domains are nearly monotonic and determining how important this is for efficient FA/C-based SI. In particular, we want to understand what other properties might make it possible to detect and resolve inconsistencies while still limiting the need to communicate raw sensor data among agents.

The concept of nearly monotonic problems should be of interest beyond the distributed AI and FA/C communities. For example, this characteristic would support efficient satisficing problem solving in complex centralized SI problems. Here the issue is not limiting communications among agents, but simply limiting the amount of data that is processed to make the problem tractable or real-time. Problem solving must be approximate and satisficing, and a statistical characterization of the monotonicity/nonmonotonicity in the domain would make it possible to evaluate the reliability of approximate solutions based on incomplete data processing. The lesson is that while nonmonotonicity is a fact of life in SI and many other domains, it is often not completely arbitrary or unpredictable, and models of its characteristics might yield important benefits.

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.
- Carver, N. 1995. Examining Some Assumptions of the FA/C Distributed Problem-Solving Paradigm. In Proceedings of the Midwest Artificial Intelligence and Cognitive Science Society Conference, 37–42.
- Carver, N., and Lesser, V. 1995a. 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, 33–40.
- Carver, N., and Lesser, V. 1995b. A Formal Analysis of Solution Quality in FA/C Distributed Sensor Interpretation Systems. Technical Report, 95-05, Computer Science Department, Southern Illinois University. (Can be obtained from: “<http://www.cs.siu.edu/~carver>”)
- Cox, I., and Leonard, J. 1994. Modeling a Dynamic Environment using a Bayesian Multiple Hypothesis Approach. *Artificial Intelligence*, vol. 66, 311–344.
- Durfee, E. 1987. A Unified Approach to Dynamic Coordination: Planning Actions and Interactions in a Distributed Problem Solving Network. Ph.D. diss., Computer Science Department, University of Massachusetts.
- 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.
- Whitehair, R., and Lesser, V. 1993. A Framework for the Analysis of Sophisticated Control in Interpretation Systems. Technical Report, 93-53, Computer Science Department, University of Massachusetts. (Can be obtained from: “<http://dis.cs.umass.edu/>”)
- Whitehair, R. 1996. A Framework for the Analysis of Sophisticated Control. Ph.D. diss., Computer Science Department, University of Massachusetts.