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. FA/C agents produce tentative, partial solutions based on only local information and then exchange these results, exploiting the constraints that exist among their local subproblems to resolve uncertainties and global inconsistencies. A key assumption of the FA/C model has been that the 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 postulate that some domains have a characteristic that can account for the success of exchanging mainly local solutions. We call such problems *nearly monotonic*. The basic idea is that while belief and solution membership are nonmonotonic with increasing evidence, partial solutions with particular attributes (e.g., high belief) are unlikely to be "significantly" affected as additional evidence is considered. The concept of nearly monotonic problems is discussed in the context of FA/C-based distributed sensor interpretation.

* 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.

1 Introduction

In the *functionally accurate, cooperative* (FA/C) paradigm for distributed problem solving [10, 12], agents need not have all the information necessary to completely and accurately solve their own subproblems. Instead, agents are designed to produce tentative, partial solutions based on only local information and to then exchange these results with the other agents. 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 uncertainties and inconsistencies that occur due to the lack of complete, accurate, and up-to-date local information.

The FA/C model has been important in *cooperative distributed problem solving* (CDPS) research. Several research systems that use an FA/C approach have been built (e.g., [3, 6, 11]). However, other than some recent work ([4, 5, 7]) there has never been any formal analysis of the conditions that are necessary for the approach to be successful or the potential performance of FA/C systems. In this paper we examine some of the assumptions behind the FA/C model 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.

A key assumption of the FA/C approach is that global solutions can be produced without the need for "excessive" communication among the agents. This requires that domains have certain properties, for example that local, partial solutions from other agents can substitute for the raw data in resolving contradictions and uncertainties to produce global solutions. While this property has been assumed in most FA/C systems, it 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 postulate that some domains have a characteristic that makes justifies the role of local solutions—at least for approximate, satisficing problem solving. We call such problems *nearly monotonic*. Basically, while belief and solution membership are nonmonotonic with increasing evidence, in nearly monotonic problems, partial solutions with particular attributes (e.g., high belief) are unlikely to be "significantly" affected as additional evidence is considered. In other words, solutions are nearly monotonic with increasing evidence once they obtain a certain level of support. This can make the consistency of local partial solutions meaningful for constructing global solutions and so reduce communication requirements.

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 that is designed 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 contains a brief description of distributed SI. Section 3 examines the use of local agent solutions to efficiently determine global solutions in FA/C problem solving. The concept of nearly monotonic problems is introduced in Section 4. This is followed by a section that examines the role this concept can play in developing coordination strategies that can produce high quality global solutions. Section 6 provides some support for the existence of nearly monotonic SI problems using a framework for analyzing SI domains and problem solvers. The paper concludes with a summary of our conclusions and future plans.

2 Distributed Sensor Interpretation

By sensor interpretation, we mean the determination of high-level, conceptual *explanations* of sensor and related data. For example, in vehicle monitoring applications this involves 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 [2, 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 *interpretation* of a data set is an explanation 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* [13]), each of which explains some subset of the data, and which together explain all of the data set. In general, there will be multiple possible alternative (uncertain) interpretations.

A *solution* to an SI problem is an interpretation of the available data that is judged "best" according to some criteria. Ideally, this would usually be the *most probable explanation* (MPE) [13] given all of the available data.[1] The problem with this definition is that for many SI problems it is impractical to compute the MPE.[2] [4] contains a fairly detailed explanation of this issue (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). Further discussion is beyond the scope of this paper—we will simply assume here that even centralized SI systems usually must use approximate, satisficing approaches to construct solutions. For instance, they may assemble solutions from hypotheses whose belief ratings surpass some *acceptance threshold* and they usually do not process every piece

[1]The terms interpretation and explanation are equivalent and interchangeable.

[2][13] contains a discussion of some of the limitations of using the MPE as *the* solution even when it can be computed. An important issue for many SI systems is differing utilities of identifying (or failing to identify) particular sources—e.g., enemy vs. friendly aircraft.

of available data. These kinds of approaches result in solutions that are only approximations of the MPE.

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*. The FA/C model is based on the recognition that in many CDPS domains it is impractical/impossible to transfer data so that agents work on independent subproblems (such that the global solution is simply the union of the local subproblem solutions). Instead, agents work on possiblely interdependent local solutions using their locally available data, and then construct a global solution by exchanging and merging these local solutions. Merging of local solutions may not be straightforward when the solutions are not independent because they may be incomplete and/or highly uncertain, and they may be inconsistent because they are based on different incomplete subsets of the data.

An FA/C system must have some mechanism to identify interdependencies among the agents' hypotheses/solutions and some mechanism for merging local solutions. [3, 6] describes how it is possible to identify interdependencies for SI applications, based on the SI model we have assumed above. Basically, 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. The ability to efficiently merge possibly interdependent local solutions so as to end up with an appropriate global solution is one of the main issues to be addressed in this paper.

3 FA/C Assumptions

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. Just what constitutes *excessive* communication will depend on the reasons for taking a distributed approach to problem solving.[3] One thing that is clear is that 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 merged only with access to their supporting raw data

[3]An common goal of a CDPS approach is improved performance in terms of reduced time to a solution relative to a centralized approach, in which case it is critical to limit the amount of raw data that must be communicated and processed by the agents. However, there are other possible goals of taking a distributed approach, including reduced communcation costs/bandwidth; tighter coupling of sensors and processors for faster sensor control; and increased reliability or graceful performance degradation. When these goals are primary, it may be less critical to reduce the communication (and redundant processing) of raw data.

and nearly all solutions were interrelated. Substantial delays would result because in the FA/C model, agents obtain data from external sensors by communicating with the agents responsible for those sensors as it becomes clear that the data is needed (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).

Thus, a requirement for FA/C problem solving without excessive communication is 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 merged with limited need for their supporting raw data. We will focus on the second approach in this paper. SI is particularly problematic for the first approach 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.[4] 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 interrelated if they have any possibility of being from targets of interest (though they may be only very weakly related). 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 (i.e., not very far into local problem solving), without excessive communication.

Our interest here is in understanding whether or when communication of raw data can be limited because the local solutions can substitute for the raw data in determining the global solutions (where there are inter-agent subproblem interrelationships). As we have said, interpretation hypotheses are high-level explanations of the sensor data—i.e., they are abstractions of the data. In general, solution hypotheses can be represented using a fraction of the storage that would be required for their supporting data. If agents could transmit mainly solution-level hypotheses rather than raw data then communication would be greatly reduced. Communication of solution hypotheses should also require receiving agents to do less processing (interpretation/probabilistic inference) to integrate this information than would be required to integrate raw data. In addition, the number of possible interpretations of a data set can be very large, so being able to concentrate on agents' local solutions can reduce communications and processing.

The developers of the FA/C paradigm certainly believed that communication of raw data could be limited by transmitting local solutions (and other abstract hypotheses, which

[4]When we speak about data being interrelated and about the "strength" of this relationship, we mean evidential relationships: the presence/absence/characteristics/interpretations of a piece of data can affect the belief in the possible interpretations of the other piece of data.

they called "results"). In particular, [10] refers to "consistency checking" of the tentative local solutions with results received from other nodes as "an important part of the FA/C approach." The basic procedure was that where there were subproblem interactions, agents would transmit local solutions and check the consistency of the related components of the solutions. They would then "integrate" consistent solution components and use these to improve and extend their local solutions.

If this approach is to reduce data transmission then "consistent" solutions must be able to be "integrated" with limited access to their underlying raw data and the resulting, merged solution must be appropriate. To evaluate the merits of this approach we must first consider what it should mean to merge local solutions into a global solution. Consider 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 merged, 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 as was used by each agent. This is the ideal in the sense that it would result in the distributed system producing the same solutions as an equivalent centralized system and it would eliminate variations in solutions due simply to differing distributions of a data set among a group of agents.

Given this standard for the merged, global solution, what can we say about "consistency checking" of local solutions? Unfortunately, what we can say is that the "consistency" of local solutions in general provides *no guarantees at all* about the quality of the global solution. Consider, for instance, a situation with the two agents above in which there are just 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 $P(I_a \mid D_1) > P(I_b \mid D_1)$ and $P(I_a \mid D_2) > P(I_b \mid D_2)$, but $P(I_a \mid D_1, D_2) < P(I_b \mid D_1, D_2)$. 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 they are identical). Likewise, if $P(H \mid D_1) > threshold$ and $P(H \mid D_2) > threshold$, it is not necessarily the case that $P(H \mid D_1, D_2) > threshold$ (where H is some 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 for the case of inconsistent local solutions. In the general case, this may require one agent to have complete knowledge of the other agent's raw data and evidential information (alternative interpretation hypotheses and their interrelationships).[5]

4 Nearly Monotonic Problems

Does the inherent nonmonotonicity of SI mean that researchers' intuitions about the usefulness of local solution consistency have been wrong? More importantly, does it mean that FA/C problem solving is not likely to be effective for real-world SI? We postulate that SI domains have a characteristic that accounts for the success of approaches that are based around consistency checking of local solutions and which would allow for the exchange of mainly partial solutions rather than raw data.

We call problems with this characteristic *nearly monotonic*. The basic idea is that while SI solutions are nonmonotonic with increasing evidence, in nearly monotonic problems the effects of additional evidence are not totally unpredictable: (partial) solutions with particular attributes (e.g., high belief) are unlikely to be significantly affected as additional evidence is considered. In other words, these problems are nearly or approximately monotonic once certain conditions have been achieved. For example, in vehicle monitoring, a significant amount of evidence is required to achieve a vehicle track hypothesis with high belief (conditional probability). While additional evidence can negatively affect the belief in this track, it will typically take substantial additional evidence to have a "major" effect on the degree of belief. In particular, once a track achieves high enough belief, it is unlikely that it will not be part of the global solution (the solution in which all of the available data is considered).

The practical consequence of this property is that consistency of appropriate local solution interpretations is now *highly predictive* that the merged, global solution will be correct, in the sense of being globally "best." This makes it possible to use a strategy in which local solutions are exchanged and when these solutions are consistent (and have appropriate characteristics) they can be easily (cheaply) merged while still ending up with a global solution that is very likely to be correct. Notice, though, that we are not talking about "exact" determination of the best global solution under this strategy—even when an approximate (non-MPE) definition of the best interpretation is being used. However, this type of ap-

[5]A detailed explanation of the recomputation of belief and solution membership for SI is beyond the scope of this paper. In probabilistic network terms, we can 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. To help understand why this may require complete knowledge of another agent, consider that probabilistic networks for SI problems are virtually always multiply connected and sensor data is *not* conditionally independent. Thus, recomputation cannot in general be done simply by message passing. A discussion of related issues in belief computation for SI and the use of local solutions can be found in [7].

proximattion is entirely reasonable given the necessity of approximation for virtually all SI problems, distributed or centralized.[6]

To proceed in examining this concept, we must present formal defintions. We will use the following notation: $\mathcal{D}$ is the complete, globally available data set, D is some subset of $\mathcal{D}$, $BEL(H)$ is the current belief in hypothesis H for data set D that has been processed (it is $P(H \mid D)$), $BEL^*(H)$ is the true, ultimate belief in hypothesis H for data set $\mathcal{D}$ (it is $P(H \mid \mathcal{D})$), MPE is the current MPE solution for data set D that has been processed, and MPE^* is the true, ultimate MPE solution for data set $\mathcal{D}$. The basic approach is to statistically characterize certain eventual properties of interpretation hypotheses given particular characteristics based on partial data/evidence. That is, given an interpretation hypothesis H, with particular characteristics (e.g., belief) based on the interpretation of some subset D of $\mathcal{D}$, we will statistically characterize particular properties of H if/when $\mathcal{D}$ were processed. The ultimate hypothesis properties that are of interest are belief and solution membership.

There are five possible definitions that we have considered:

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

Definitions 1 through 3 deal with the belief in individual hypotheses, while 4 and 5 deal with solution membership. The probabilities in definitions 2 and 3 can be derived from the pdf of defintion 1, but are included because we expect that detailed knowledge such as the pdf may not always be practical to obtain. Variations on these definitions can result from the use of approximate beliefs and/or solutions, rather than the exact ones used here. For example, definitions 1 through 3 could deal with current approximate beliefs relative to ultimate approximate or exact beliefs, and definitions 4 and 5 could instead involve solutions that are not the MPE but some other definition of best interpretation (this might be either

[6]It is nearly always necessary to trade-off solution quality for efficiency in SI problems, though, approximations can be done in a variety of ways (see [1, 7, 8]).

an approximation of the MPE or it might take into account the goals of the SI system by way of the differing utilities of identifying or failing to identify different types of signal sources). Which of these defintions is most appropriate will depend on: 1.) the domain and its characteristics; 2.) the solution selection strategy; and 3.) the knowledge that is available or can be reasonably collected.

For all of these definitions, the concept of a nearly monotonic problem means that once an interpretation hypothesis has certain characteristics (based on only a portion of the available data) then these probabilities will be *high enough* to make it appropriate to assume the hypothesis is in the solution. For example, 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 the 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 (merged) track is in the global solution.

We are currently exploring the set of hypothesis characteristics which the belief and solution probabilities described above should be a function of. This will certainly 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 with SI problems it appears that both current hypothesis belief and hypothesis type are important factors. There are other possibilities as well, though. For example, an evaluation of the "quantity" of data supporting the hypothesis or the fraction of the overall data that has been processed—possibly relative to the different data sources involved. The RESUN SI framework [2] also provides more detailed information about belief, such as whether uncertainty is due simply to the possibility of alternative explanations for the data or whether it is due to contradictory data. Such information may be necessary to be able to identify hypotheses that are reliable enough to be assumed for global solutions.

5 Nearly Montonic Problems, Solution Quality, and Coordination Strategies

Nearly monotonic problem domains are of interest because they provide the possibility of producing high quality global solutions with greatly reduced communication costs. In this section we will examine the issue of solution quality when using the "consistent local solutions strategy" for merging local solutions We will also discuss the numerous trade-offs involved in developing FA/C coordination strategies to take advantage of near monotonicity.

One standard for judging solution quality is the MPE solution given all of the data, MPE^*. In other words, evaluate a global solution $\mathcal{S}^G$ produced by an FA/C-based SI system against MPE^*. $\mathcal{S}^G$ could be compared against MPE^* in terms of $P(\mathcal{S}^G = MPE^*)$, however this is often not the most meaningful metric for SI problems. First, $\mathcal{S}^G$ will be tend to be incomplete ($\mathcal{S}^G \subset MPE^*$) 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, we care more about the likelihood of individual hypotheses being correct than we do about 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 $P(H \in MPE^* \mid H \in \mathcal{S}^G)$.

In order 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 *merged* solution. Our definition of (global) 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 hypotheses of the same type that can be *combined* into a single consistent hypothesis of that type. Combination typically involves producing a single "more complete" hypothesis from two or more "less complete" hypotheses. For instance, two partial vehicle track hypotheses may be combined into a longer track hypothesis. While a combination hypothesis could always be built from scratch from the combined supporting data/evidence of the component hypotheses, when we refer to the "merging" of such hypotheses we will assume that combination is done from the component hypotheses, without reference to their supporting data.[7] While 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 formal definition of what we mean by the "consistent local solutions strategy" for developing global solutions:[8]

1. each agent uses their 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;

[7] The Distributed Vehicle Monitoring Testbed (DVMT) [9, 11] had "merge" operators that did exactly this while DRESUN [6] allows hypotheses to directly be "external evidence" for hypotheses of the same type.

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

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. now the agent proceeds to merge its solution, one-by-one, with each of the solutions it has received and may receive prior to termination;.
5. processing terminates when all agents have transmitted their solutions according to steps 2 and 3, and have merged their solution with all the received solutions;
6. the global solution is simply the union of all the final individual agent merged solutions.

If two agents local solutions are consistent when they are exchanged then the merged 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 combined. 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, based on the joint data each has processed.

Now to derive some results about solution quality using this strategy, we will assume that we have available the probability of a hypothesis being in the ultimate MPE solution, as in definition 4 above in Section4. Let us further assume that this probability is well correlated with hypothesis type and hypothesis belief (so no additional conditioning factors are needed). Thus, we assume that we have available $P(H \in MPE^* \mid type(H), BEL(H))$, which we will refer to as $P_{MPE^*}(H)$.

Solution Quality Theorem: Assume that the agents use the "consistent local solutions strategy" described above for merging their local solutions into a global solution. Assume also that agents compute $BEL(H) = P(H \mid D)$ for the subset D of their own local data that they process before transmitting their solutions and that in the case of inconsistent local solutions the agents compute the MPE joint partial solution (based on the data they have jointly process).[9] Under these conditions, what we can say about the quality of the resulting global solution is that $\forall H : P(H \in MPE^* \mid H \in \mathcal{S}^G) \geq P_{MPE^*}(h_{max})$. h_{max} is simply H unless H resulted from the *combination* of consistent hypotheses of the same type or identical hypotheses. In this case, h_{max} is the component hypothesis of H with the *maximum* belief out of all the component hypotheses (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 this 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 joint solution (i.e.,

[9]This 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 due compute proper conditional probabilities. If the assumption is relaxed this results in additional approximations in the global solution generated.

$P(H \in MPE \mid D_1, D_2) = 1.0)$. When solutions are consistent, no further exchange of information will take place, and the joint solution will be the merge of the consistent solutions. Supposed 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 $H_i \in S_2$, or 2.) corroborative with one or more hypotheses, say $\{H_j\} \subset S_2$ (they may be of the same or of different types), or 3.) identical to a hypothesis $H_i \in S_2$. 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. If H is corroborative with hypotheses in S_2 in the sense or being supporting or explanatory then $BEL^{A_1,A_2}(H) \geq BEL^{A_1}(H)$. Since the function $P_{MPE^*}(H)$ will be monotonically nondecreasing with increasing hypothesis belief, the joint value for $P_{MPE^*}(H)$ must be greater than or equal to its value from A_1's local data. Now we must deal with hypotheses in $\mathcal{S}^G$ that have resulted from combining hypotheses of the same type from different local solutions. Suppose that H is the result of merging h_1 from A_1 and h_2 from A_2. We must have $BEL(H) \geq BEL(h_1)$, $BEL(H) \geq BEL(h_1)$, and so $BEL(H) \geq maximum(BEL(h_1), BEL(h_2))$, by our definition of consistency. Thus, $P_{MPE^*}(H)$ here is based on the component hypothesis with the maximum belief prior to combination. The situation in which H is identical with some $H_i \in S_2$ is analogous since $BEL^{A_1,A_2}(H)$ should be $maximum(BEL^{A_1}(H), BEL^{A_2}(H_i))$.

What this result tells us is that we can use the "consistent local solutions strategy" and *potentially* get a global solution that is as close as possible to the MPE global solution as we desire. Of course whether we can get a solution of a certain quality or not depends on the ability of the agents to produce local solutions whose hypotheses have high enough degree-of-belief ratings to give us the desired solution quality. This suggests that an important criteria for effective use of this strategy is structuring the distributed system so that agents have sufficient data to achieve appropriate local solutions. If this is not possible, then the system may be forced to communicate considerably more information/data among the agents in order to produce a global solution of the desired quality. Another factor that affects the power of this strategy is the likelihood that local solutions will be consistent. Here too, the distributed system should be structured so as to make this sufficiently likely.

For effective FA/C-based SI, with limited communications requirements, 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 some agents fail to produce nearly monotonic solution hypotheses or some agents have inconsistent solutions then raw data will have to be transferred and processed by some agents, in general. 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.

6 Analytic Support for Nearly Monotonic Problems

While we have demonstrated that nearly monotonic problems have the potential to support efficient FA/C-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, and 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 domans 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 [14, 15]. 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) [9, 11]. A complete definition of a vehicle tracking domain (that is more complex than any DVMT grammars) and an associated problem solver is presented in [15]. Some statistics from this grammar will be presented at the end of this section, but first we will use a very simple vehicle tracking grammar to illustrate a nearly monotonic SI domain.

Consider the problem domain defined by the grammar rules in Figure1 In this domain, there are two kinds of vehicles, 1 (V1) and 2 (V2). In the grammar, the nonterminals Track1 and Track2 correspond to vehicle tracks of type V1 or V2. The terminal symbols in this grammar, S2, S3, S4, and S5, correspond to actual signal 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

P.1.1. Start $\rightarrow$ Track1_t (p = 0.7)
P.1.2. Start $\rightarrow$ Track2_t (p = 0.3)
P.2.0. Track1_t $\rightarrow$ V1_t N1_t Track1_{t+1} (p = 1.0)
P.3.0. Track2_t $\rightarrow$ V2_t N2_t 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 $\rightarrow$ lambda (p = 0.2)

Figure 1: IDP grammar for simple vehicle tracking domain.

are the potential solutions for problem instances generated with this grammar. We refer to grammars of this form as *Generational IDP Grammars*, (IDP$_g$) [15]. 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 signal data "S2 S4" is observed. This data is ambiguous, meaning it could have been generated by either a Track1 or a Track2. A Track1 could have generated an "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 partial results V1 and V2 generated from an "S2 S4," and given that $BEL(H)$ is the problem solver's belief in the partial result H, the values $P(V1 \in MPE^* \mid BEL(V1))$ and $P(V2 \in MPE^* \mid BEL(V2))$ are approximately equal for any values of $BEL(V1)$ and $BEL(V2)$. This means that it is not possible to use the beliefs of V1 and V2 to differentiate between a Track1 and a Track2—i.e., the domain is nonmonotonic. The belief in V1 could be very high, and it would not matter because the probability that the "S2 S4" was generated by a Track2 would be equally high.

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

More specifically, as the length of the partial track of V2 results 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 (which it is given the example grammar), then the probability of two "S2 S4" events occuring sequentially is 0.25 * 0.25. The probability of three such events occuring sequentially is 0.25 * 0.25 *0.25. And so forth.

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 phenomena holds in other domains as well. For example, as the length of an interpreted fragment of speech increases, it is likely that the associated equivalence class of partial results will exhibit nearly monotonic properties.

As a further demonstration of that this phenomena occurs in complex domains, we used the IDP framework to collect statistics for a more complex vehicle tracking problem domain. 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 equivalence classes, "group level (gl)" partial results, "vehicle level (vl)" partial results, and "partial tracks of length 4 (pt)." We accumulated statistics as follows. We repeatedly generated problem instances and determined their MPE^* interpretations (i.e., solved the problem). For each such cycle, we recorded the beliefs (or credibilities) of the elements of each of the equivalence classes. However, we divided each of the equivalence classes into two subsets, those partial results that were subsequently used in MPE^* and those that were not.

We found that gl and vl equivalence classes were nonmonotonic with respect to belief (really credibility[10]). In other words, the belief of a partial result from classes gl and vl was not a good indicator of whether or not the partial result was an element of MPE^*. Specifically, the distribution of $BEL(H)$ was approximately the same for all partial results in the equivalence classes gl and vl regardless of whether or not they were actually part of MPE^* and $P(H \in MPE^* \mid type(H) = vl, credibility(H) = 0.7)$=0.32. On the other hand, for partial results in the equivalence class pt, there was a strong correspondance between $BEL(()$H) and membership in MPE^*. Thus, if a partial track of length 4 had a high belief, it was very likely to be part of MPE^*, making the equivalence class pt nearly monotonic.

[10]The IDP analysis tools currently compute a credibility rating for each hypothesis, but this is not what we have termed belief—i.e., the conditional probability of the hypothesis. $Credibility(H) \geq BEL(H)$.

For example, $P(H \in MPE^* \mid type(H) = pt, credibility(H) = 0.7)$=0.92. On the other hand, $P(H \in MPE^* \mid type(H) = pt, credibility(H) = 0.55)$=0.5.

7 Conclusion

In this paper, we have shown that while consistency checking of local 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.

While 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. For example, the importance of SI domains having the property of being nearly monotonic in allowing for efficient FA/C-based SI remains an open issue. Near monotonicity can support efficient FA/C problem solving, but the property alone does not 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 merged with limited communication of raw data—particularly if we need only approximate global solutions. The focus of our future research in connection with the work is assessing the occurrence of near montonicity in real-world SI domains, determining its importance to efficient FA/C-based SI, and understanding under what other conditions communication can be limited.

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 montonicity/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.

References

[1] Yaakov Bar-Shalom and Thomas Fortmann, *Tracking and Data Association*, Academic Press, 1988.

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

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

[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] Norman Carver, "Examining Some Assumptions of the FA/C Distributed Problem-Solving Paradigm," *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Society Conference*, 37–42, April, 1995.

[6] Norman Carver, Victor Lesser, "The DRESUN Testbed for Research in FA/C Distributed Situation Assessment: Extensions to the Model of External Evidence," *Proceedings of the International Conference on Multiagent Systems*, June, 1995.

[7] Norman Carver and Victor Lesser, *A Formal Analysis of Solution Quality in FA/C Distributed Sensor Interpretation Systems*, Technical Report 95-05, Computer Science Department, Southern Illinois University, 1995.

[8] Ingemar Cox and John Leonard, "Modeling a Dynamic Environment using a Bayesian Multiple Hypothesis Approach," *Artificial Intelligence*, vol. 66, 311–344, 1994.

[9] Edmund Durfee, *A Unified Approach to Dynamic Coordination: Planning Actions and Interactions in a Distributed Problem Solving Network*, Ph.D. Thesis and Technical Report 87-84, Computer Science Department, University of Massachusetts, 1987.

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

[11] 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).

[12] 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.

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

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

[15] Robert Whitehair, *A Framework for the Analysis of Sophisticated Control*, Ph.D. Thesis, Computer Science Department, University of Massachusetts, 1996.