Towards Robust Teams with Many Agents

Gal A. Kaminka *
Computer Science Department
Bar Ilan University
galk@cs.biu.ac.il

Michael Bowling
Computer Science Department
Carnegie Mellon University
mhb@cs.cmu.edu

ABSTRACT

Agents in deployed multi-agent systems monitor other agents to coordinate and collaborate. However, as the number of agents monitored is scaled up, two key challenges arise: (i) the number of monitoring hypotheses to be considered can grow exponentially in the number of agents; and (ii) agents become physically and logically unconnected (unobservable) to their peers. This paper examines these challenges in teams of cooperating agents, focusing on a monitoring task that is of particular importance to robust teamwork: detecting disagreements among team-members. We present YOYO, a highly scalable disagreement-detection algorithm which guarantees sound detection in time linear in the number of agents despite the exponential number of hypotheses. In addition, we present new upper bounds for the number of agents that must be monitored in a team to guarantee disagreement detection. Both YOYO and the new bounds are explored analytically and empirically in thousands of monitoring problems, scaled to thousands of agents.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Coherence and coordination*

General Terms

Algorithms, Performance, Reliability

1. INTRODUCTION

Agents in realistic, complex, multi-agent domains must monitor other agents to accomplish their tasks, detect failures, coordinate, and collaborate. Indeed, the importance of agent monitoring in deployed multi-agent systems has long been recognized in theory (e.g., [4, 7, 8]), and in practice, ranging from industrial systems (e.g., [12]), to virtual environments for training and research (e.g., [22, 23]), to human-computer interaction (e.g., [17]), and multi-agent robotics (e.g., [19, 2]). Agent monitoring infrastructure is of particular importance in teams of cooperating agents, since the correct execution of teamwork mandates that team-members come to agree on the task that is jointly executed by the team, and manage interdependencies among team-members [4, 8].

As the number of agents in the system increases, a monitoring agent cannot monitor all activities of all other agents at all times [12, 7, 8]. Much of previous work has therefore explored different approaches to reducing monitoring activities to accommodate the bandwidth in the target application, by relying on focused communications [12, 21], using plan-recognition to infer agents' state based on their observable behavior [10, 17], and monitoring a shared environment [23].

However, key challenges raised by a scale-up in the number of monitored agents remain largely unaddressed. First, as bandwidth is used more selectively, less information is available about monitored agents, and thus there is some uncertainty about their state. However, in many monitoring tasks (e.g., coordination, teamwork monitoring) the computational complexity of reasoning about multiple agents and their possible interactions can increase exponentially in the number of agents, even under limited uncertainty [15]. Second, previous monitoring approaches rely, in general, on the monitoring agent to be able to communicate or observe all monitored agents. However, as the number of agents increase, agents become more physically and logically separated, and thus a monitoring agent may not be able to observe (or communicate with) the agents it is to monitor. We call this challenge *limited connectivity*.

This paper addresses these difficulties in the context of a particularly important monitoring task in robust multi-agent teams—that of *detecting disagreements among teammates*. Theoretical and empirical research on teamwork in synthetic agents [4, 12, 8, 21, 16] and in humans, (e.g, [3]) stresses agreement as a cornerstone to effective teamwork (although the literature differs in the terms used and in grounding agreement in various theoretical and practical constructs). Thus disagreements are a source of great concern in all of these different investigations (see Section 2 for details).

We present two sets of contributions. First, we present YOYO, a disagreement-detection monitoring algorithm, which navigates the (potentially exponential) space of monitoring hypotheses by representing only hypotheses in which all agents are in agreement. This allows YOYO to represent the relevant state of all monitored agents together, in a single highly scalable structure. YOYO is an example of a *Socially-Attentive* monitoring algorithm, exploiting knowledge of the social relationships in the monitored team. It is based on earlier work on visualization [14], but differs from it in many ways (see Section 6).

A second set of results tackles the challenge of limited connectivity by providing new bounds on the number of agents that must be monitored in a team to detect disagreements. Previous work has shown analytically that disagreement detection can sometimes be guaranteed if all team-members monitor all of certain key agents

*This research was carried out while the author was a Post Doctorate Fellow at Carnegie Mellon University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'02, July 15-19, 2002, Bologna, Italy.
Copyright 2002 ACM 1-58113-480-0/02/0007 ...$5.00.

in the team [15]. However, in practice, limited connectivity restricts the the usefulness of this bound, as often not all key agents can be observed or communicated with. To address this, we show analytically that sound and complete detection can be guaranteed in practice even if non-key team-members monitor only one key-agent, while key agents monitor each other. In addition, we show that unfortunately, while some potential for limiting the number of monitored agents exists in centralized settings, the worst case still requires monitoring all agents in a team.

These results are motivated by practical concerns raised in our application domains, as monitoring settings become realistic and the number of agents is scaled up. Using the techniques presented, a monitoring agent can detect disagreements in large teams, involving thousands of agents, that are guaranteed to perform specific monitoring tasks under conditions of limited connectivity. In addition to the analytic results, we present an empirical evaluation in thousands of monitoring problems, scaled to thousands of agents.

This paper is organized as follows. Section 2 presents motivating examples and background. Section 3 presents an overview of the monitoring task. Section 4 presents the YOYO algorithm. Section 5 presents new bounds on the number of agents that must be monitored. Section 6 addresses related work, and Section 7 concludes.

2. MOTIVATION AND BACKGROUND

Teamwork literature, addressing human and synthetic teams, has often emphasized the importance of team-members being in agreement on various features of their state, such as goals, plans, and beliefs[1]. Teamwork theory often defines agreement as a state of mutual belief, where agents reason to infinite recursion about their beliefs and their beliefs in others' beliefs in a proposition. For instance, *SharedPlans* theory requires team-members to mutually believe in a shared recipe [8] during the planning and execution phases of the task; the *Joint Intentions* framework emphasizes mutual belief in the team goals' selection, as well as in team-members' beliefs about the goals' achievability and relevance [4, 16]. Other investigations of agent teams have emphasized agreement on team plans to be jointly executed by team-members [12], on hierarchical team operators [21], on tasks to be executed collectively [19], etc. Investigations of human teamwork have not only emphasized agreement on the joint task, but also agreement on features of the environment that are important to the task being carried out by the team [3].

However, the literature also recognizes that achieving and maintaining agreement can be difficult. Teamwork theory recognizes that attainment of agreement by mutual belief is undecidable [9] and must therefore be approximated in practice. Such approximations frequently involve assumptions of trustworthiness of team-members, of foolproof communications [12], of team-members being able to observe each other [10], and/or of a mutually-visible environment. As is often the case with approximations, they sometimes fail in practice (e.g., due to communications failures), and therefore team-members may find themselves in disagreement with each other. Such disagreements are often catastrophic, due to the unique importance of agreement in collaboration.

It is therefore critical that teams are monitored to detect such disagreements. A monitoring agent that identifies the state of team-members can compare the state of different team-members and detect differences on state features that, by design or by selection, should have been agreed upon [15]. However, as the number of monitored agents is scaled up, two challenges arise: (i) monitoring algorithm complexity due to uncertainty about the state of agents; and (ii) difficulty to observe or communicate with all agents (*limited connectivity*). We have come to realize the need to address these challenges while working on developing robust multi-agent teams in two dynamic, complex, domains: *ModSAF*, a commercially-developed, high-fidelity virtual environment, where we have been involved in the development of synthetic helicopter pilot agents that carry out a variety of missions [22]); and *RoboCup soccer simulation*, a dynamic research-oriented simulation which requires real-time teamwork and coordination, where we have been involved in the development of both soccer-playing agents and a coach agent [23].

[1] Of course, the literature also addresses other critical features of teamwork aside from agreement. But agreement is a repeating theme in recent work.

Monitoring Attacker	Other Attacker	Scout
WAIT-FOR-SCOUT	FLY	JOIN-SCOUT
WAIT-FOR-SCOUT	FLY	HALT-ORDER
WAIT-FOR-SCOUT	JOIN-SCOUT	JOIN-SCOUT
WAIT-FOR-SCOUT	JOIN-SCOUT	HALT-ORDER

Table 1: Hypotheses for state of team in Example 1.

Monitoring Algorithm Complexity. As discussed, agents cannot continuously communicate with a monitoring agent about their state. Thus in general, the monitoring agent has uncertainty about the state of monitored agents, i.e., the the monitoring agent entertains multiple hypotheses as to the state of each monitored agent. To detect disagreements, the monitoring agent must compare the state of one agent to the state of another. Since there may be multiple hypotheses as to the state of each of the agents, the monitor must select possible combinations of the hypotheses of different agents, to serve as the basis for the decision on whether a disagreement has occurred. However, the number of combinations of individual hypotheses can grow exponentially in the number of agents. Thus it would seem that we would need to go through an exponential number of hypotheses to pick those that are useful for monitoring purposes. Consider an example, borrowed from [15]:

EXAMPLE 1. *In the ModSAF domain, three helicopters are executing the* WAIT-FOR-SCOUT *plan, in which one of them (role:* scout*) is flying towards the enemy while its two teammates (role:* attackers*), have landed. Once the scout identifies the enemy, it radios back to the attackers, causing all three agents to switch to the* JOIN-SCOUT *plan, in which the scout lands, while the attackers fly forward to join it. Due to a radio equipment malfunction, one attacker failed to receive the message from the scout, causing it to continue waiting. Thus a state of disagreement occurs among the agents. Suppose this attacker is monitoring its teammates by observing their actions and inferring their currently executing plans: The scout has landed, and so may have started the* JOIN-SCOUT *plan, or may be responding to a command to land immediately (the* HALT-ORDER *plan). The other attacker, flying towards the scout, may be executing its role in the* JOIN-SCOUT *plan, or a completely different plan (the* FLY *plan). Although there are only two hypotheses for each individual monitored agent, there are* four *hypotheses as to the overall state of the team (Table 1), and as the size of the team grows, the number of hypotheses increases exponentially.*

Limited Connectivity . As the number of agents grows, agents become more logically and physically distributed, and cannot maintain continuous contact with each other. We use the term *limited connectivity* in a general sense, to denote both limited ability to observe a particular agent's actions (e.g., because of occlusion

or physical distribution), and limited communications (e.g., due to interference, range, or reliability issues). Consider the following example:

EXAMPLE 2. *In the RoboCup domain, the 11 soccer-playing agents switch between two high-level game plans, triggered by referee messages. The* INTERRUPT *plan is called in the beginning of the game, at half-time, etc., and requires players to place themselves in pre-determined home positions, standing still. The* PLAY *plan requires players to play soccer freely. Due to failures on the agents perception skills, some players sometimes fail to hear the referee, and so fail to switch plans. They are then in disagreement with their peers who did hear the referee. Under perfect connectivity, players would be able to communicate with their teammates, or see their teammates standing still, and thus detect disagreements. However, players have communication range and bandwidth limitations, and have a limited field of view. They can therefore not normally see or communicate with all their teammates. They therefore face difficulties in detecting disagreements.*

Our own previous work [15] has shown that sometimes only certain *key agents* must be monitored to guarantee detection. And yet, every agent must monitor all of these key agents. Furthermore, in a worst case, all agents are key agents, and therefore would seem to require full connectivity. Indeed, this is the case in Example 2.

3. DISAGREEMENT DETECTION

Disagreement detection involves a key step of representing the state of monitored agents, such that the state of different agents can be compared to detect disagreements. This section briefly describes a general representation of monitored agents' states, and a basic inference algorithm which uses the representation for observation-based and communication-based monitoring. The representation, inference algorithm, and their use for disagreement detection have been discussed in detail in [15].

3.1 Representation and Inference

Much of contemporary theoretical and empirical work on teamwork (collaboration), both in synthetic agents and in humans, has emphasized agreement on a hierarchical recipe, or plan, as a key to effective teamwork, (see, for instance, [12, 8, 21]). Given this emphasis, we focus on a monitoring representation that follows two key constraints: (i) representing agents in terms of their currently executing plans (and plan-steps); (ii) allowing the designer, or monitoring agent to mark plans that have to be agreed upon, so that they are executed *jointly* (together) by all members of a team (or subteam). These two constraints give rise to two structures that are used by the monitoring system: A plan-decomposition hierarchy, and a team organization hierarchy. These have been fully described in [24, 14], and so we only provide a brief overview here.

A plan-hierarchy is used to represent a monitored agent's plan. It is defined to be a directed connected graph, where vertices are plan steps, and edges signify hierarchical decomposition of a plan into sub-plans. Each vertex has at most one parent (i.e., one incoming hierarchical decomposition edge); a plan that conceptually has many parents (i.e., it is a component in the decomposition of different parent plans) is represented by different vertices in the plan-hierarchy. Multiple outgoing edges signify alternatives available to the agent, of the first subplan to be executed. The graph forms a tree along hierarchical decomposition edges, so that no plan can have itself as a descendent. A vertex with no children edges denotes an atomic step.

For example, Figure 1-a presents a portion of the plan-hierarchy used to monitor the ISIS'97 RoboCup Simulation team [23]. The

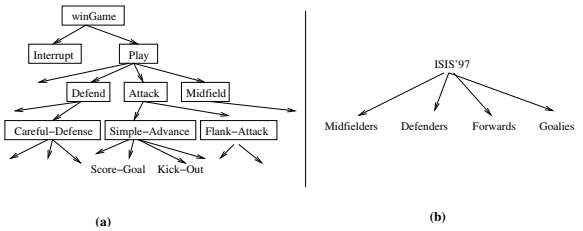

Figure 1: Plan-hierarchy (a) and team-hierarchy (b) in the RoboCup domain. Boxed plans denote team-plans, which must be agreed-upon by team-members.

top-level plan, WINGAME, is selected by all players as soon as they join a game. It has one first child, the INTERRUPT plan, which is assumed to be selected by the agent whenever the game is interrupted by the referee. WINGAME's other child, PLAY, follows INTERRUPT in order of execution, and is selected when the game is currently playing. Thus INTERRUPT and PLAY follow each other to the end of the game. In service of PLAY, players choose a plan (ATTACK, DEFEND, etc.) based on their role in the team: forwards, defenders, etc. (discussed later). This decomposition continues. For instance, at a particular given moment, a forward may be monitored to be engaged in executing the following path (from root to leave): WINGAME — PLAY — ATTACK — SIMPLE-ADVANCE — SCORE-GOAL.

To make use of the representation, a monitoring agent must have a way of associating information it senses about another agent with paths in a plan-hierarchy used to represent the monitored agent's state. An algorithm for doing this, called RESL, has been previously described in [15] and is presented here briefly: The designer of the monitoring system associates with each plan a set of *observables*, condition monitors that tie in sensor readings and received communications with particular plans in the hierarchy. When a condition monitor matches the sensor reading (e.g., when a message is received that is consistent with the plan in action, or when an action associated with a plan is observed), we tag the plan *matching*. If its observables fail to match, the plan is tagged *not-matching.* RESL infers the state of unobservable plans from their children and parents: An otherwise untagged parent with at least one successfully-matched child is tagged successfully-matching, otherwise it is tagged as failing to match. And an untagged child with a successfully-tagged parent is tagged successfully-matching, unless all of its own children are tagged as failing to match. In this way, all plans in the hierarchy are tagged as matching or not-matching the observations. Multiple matching siblings denote multiple hypotheses. The process is linear-time in the size of the plan hierarchy.

To monitor multiple agents using the above representation, we construct a separate plan-hierarchy for each monitored agent. When a specific agent is observed, its state is updated in the hierarchy which represents it. This method has been successfully used in monitoring agents deployed in ModSAF [22], RoboCup [23], the civilian evacuation simulation, and enterprise scheduling [24]. In general, it can be useful in monitoring agents whose behavior is controlled by a hierarchical process, e.g., hierarchical behaviors [2], Soar [18], or HAC [1]. Its generality is derived from its use in monitoring rather than execution and control. It avoids any details which determine how the actual decisions of agents are made (e.g., in preferring one decomposition over another) since it is only used to organize and keep track of their decisions after the fact.

3.2 Detecting Disagreements

In order to detect disagreements, the monitoring agent must first

know which plans are ideally to be agreed upon. Plans in the hierarchy must be marked as *team plans* [12, 21]. We assume that team plans are marked by the designer, or by the monitoring agent, for instance based on executable teamwork models such as STEAM [21] or GRATE* [12]). Different subteams can execute different team plans. In our experiments, team plans were known since the monitored agents had actually used STEAM, and therefore used team plans explicitly. In Figure 1-a, team plans are boxed: WINGAME, PLAY, and INTERRUPT are to be executed by the all members of the RoboCup team ISIS'97. MIDFIELD, DEFEND, etc. are to be executed jointly only by members of the corresponding subteams of ISIS'97 (*midfielders*, *defenders*, etc.).

The monitoring agent, after hypothesizing the state of agents as previously discussed, matches team plans across members of teams—if agents are in agreement about their selected team plans, then all is well. If agents are not in agreement about their team plans, then a disagreement is announced. Note that agents do not have to be in agreement about all plans—only about those plans that are marked as team plans. Furthermore, the plan-hierarchies used for different agents may themselves be different (other than in the team plans), facilitating monitoring of behaviorally-heterogeneous agents.

As previously discussed there can be in general multiple hypotheses for each individual, and an exponential number of hypotheses for the team as a whole. To evaluate hypotheses, previous work has formally defined the *coherence* of a multi-agent hypothesis as the ratio of the number of agents to the number of different selected team plans [15]. Intuitively, coherence is a measure of the Agreement in a team. For example, the hypotheses in Table 1, Rows 1, 2, & 4 have coherence level of 1. The maximally-coherent hypothesis (Row 3) has coherence level 1.5 (3/2), since unlike the other hypotheses, it has at least two agents in agreement. Optimal coherence has all N agents in agreement, and has value N.

It was shown that as long as the monitoring process is complete (see formal definition in Section 5), the agent is guaranteed that its detection results are going to be *sound*, by selecting *maximally-coherent* hypotheses: Any detection is going to be of a real failure [15, Theorem 1], but not all failures are guaranteed to be detected. Soundness is very important to a monitoring system, since it prevents false alarms which would otherwise task the agent needlessly. However, a difficulty emerges in applying this result in monitoring an increasing number of agents. While only a single maximally-coherent hypothesis is needed, the process of finding such an hypothesis potentially still requires *going through all exponential number of hypotheses* to rank them based on coherence. Formally, the space required for monitoring multiple agents using an array of plan-hierarchies (as described above and in previous work) is $O(NL)$, where N is the number of agents, and L is the size of the plan-hierarchy. However, sorting through the hypotheses can take $O(L^N)$, as previously demonstrated (Example 1).

4. SCALING UP MONITORING

To address the challenges raised by the time and space complexity of previously known techniques, we present YOYO, an algorithm that utilizes knowledge about the team organization to carry out disagreement detection in *linear time*. The intuition behind YOYO is to represent only coherent hypotheses (of which there is a linear, not exponential, number), and then recognize disagreements as cases where the representation fails. YOYO is based upon earlier work on the YOYO* visualization algorithm [14], but differs from it in several important ways (discussed in depth in Section 6).

The key idea in YOYO is to represent all agents together, in *a single shared plan hierarchy*. The shared plan-hierarchy is fully expanded to contain the plans and transitions for all subteams, annotated so that YOYO can determine which subteam is to take which transitions, execute which plans, etc. A plan P in this hierarchy, when tagged as successfully matching, represents the hypothesis that all agents in the team associated with P are executing P. Observations about agents are then matched against the shared plan-hierarchy. Intuitively, the process of detection proceeds as follows: If some team members are executing P, while others are executing a different plan Q, and assuming the observations allow us to differentiate P, Q then both will be marked matching and not-matching at the same time, and we will know that a disagreement has occurred.

However, members of different subteams execute different plans by design. Therefore, YOYO needs to differentiate cases where members of the same team have selected different plans P, Q, and cases where members of different teams have selected P, Q. To do this, YOYO exploits knowledge of the social structure within the monitoring system, in the form of a designer-specified *team-hierarchy*, a tree-like structure which encodes knowledge about the relationships between teams, subteams, and team-members: Each node in the team-hierarchy corresponds to a monitored organizational unit. The top (root) team represents the entire monitored team. These teams are then split into several subteams, etc., until the leaves of the hierarchy contain roles of individual agents, if such different roles exist. For instance, Figure 1-b presents the team-hierarchy of the ISIS'97 RoboCup team [23], composed of a root node for the entire team, and four nodes for its four subteams.

The team-hierarchy contains pointers from each node to plan nodes in the shared plan-hierarchy. The plans pointed to are the hypothesized coherent plans of the monitored team, and thus multiple pointers are allowed from a single team-hierarchy node. The pointers in the team-hierarchy point at the lowest-level plans that are consistent with the observations, and are to be executed by the team in question. For example, suppose all RoboCup players are executing the PLAY plan together, and that members of each subteam are in agreement with their teammates on the plan chosen for the subteam. A player that observes its teammates using the team- and plan-hierarchies in Figure 1 will have pointers from the ISIS'97 node in the team hierarchy (Figure 1-b) to the PLAY plan in the plan-hierarchy (Figure 1-a). Each of the subteam nodes in the team-hierarchy will have pointers to plans in the plan-hierarchy which are executed by the different subteams. For instance, the *Forwards* subteam may have a pointer to the SIMPLE-ADVANCE plan, signifying the all members of the subteam are executing this plan.

YOYO (Algorithm 1) maintains the pointers such that the hypotheses they represent are coherent with each other at all times. If it fails, then this means that the team's state is unambiguously incoherent, i.e., a disagreement exists. YOYO operates as follows: When an observation is made about an agent (called the *focus*), we not only update the pointers for this agent, but also re-align the pointers of its parent (sub)teams, such that their own pointers point at plans that are coherent with the new hypothesized state of the focus. We then go up and down the team-hierarchy to re-align the pointers of the other agents which are either part of the focus' subteam or its siblings' subteam. This is done by moving the pointers of non-focus agents (and the subteams of which they are members) such that they point at a plan that is coherent with the plans hypothesized for the focus. If the initial set of pointers for any non-focus agent is already coherent, no re-alignment is necessary. If no plan can be found for them, or if all plans for a team are tagged both matching and not-matching at the same time, then a state of disagreement has been detected.

Algorithm 1 YOYO(**plan hierarchy** L, **team-hierarchy** H)

```
 1: for all observations O_i at time t do
 2:    for all plans X that have conditionals testing O_i do
 3:       let A_i be the agent observed in O_i
 4:       if X matches O_i and X executable by A_i then
 5:          tag X as successfully matching
 6:          create pointer from A_i node (in H) to X
 7:       else
 8:          tag X as not matching
 9: for all plans X tagged matching or not matching do
10:    T ← team(X), P ← X
11:    while parent(P) ≠ null do
12:       propagate any tags to parent(P) (match or not match)
13:       if team(parent(P)) = parent_team(T) then
14:          propgate tags down to untagged plans
15:          create pointer from parent_team(T) to parent(P)
16:          T ← parent_team(T)
17:       P ← parent(P)
18: for all teams T in H do
19:    if T only points to plans matching and not matching then
20:       a disagreement has occurred.
```

For example, suppose that the players are known to be executing the INTERRUPT plan (Figure 1). Suppose now that a *defender* observes a *forward* player running towards the opponent goal (i.e., executing the ATTACK plan, in service of PLAY). The defender tags Interrupt as matching, and Play as not matching (based on its own selected plans). However, YOYO (executed by the monitoring defender) will create pointers for observed attacker to point to the ATTACK plan (lines 1–8), and will tag as not-matching other plans, in particular INTERRUPT. It will then enter the loop on line 9. For the ATTACK plan, it will propagate its successful tag up to the PLAY plan (line 12), and the create a pointer from the root node in the team hierarchy, representing the ISIS'97 team (the root in Figure 1-b) to PLAY, since the team that executes PLAY is the ISIS'97 team (the parent of the the *Forwards* subteam). It will then climb up in both hierarchies (lines 16 and 17) and begin another iteration. Later on, the same process will be repeated for the INTERRUPT plan. Since both INTERRUPT and PLAY are pointed to from the node representing the team ISIS'97, a disagreement will be detected.

YOYO requires a slightly modified inference procedure than previously described (Section 3). First, in following children transitions, YOYO is careful to only take paths legal to the team in question (i.e., plans and transitions that the team is allowed to execute in its role): It thus makes the assumption that transitions in the plan hierarchy are marked for the subteams that are allowed to take them. Second, YOYO must use a time-stamp to tag plans, so that observations that arrive simultaneously (but processed serially) will cause a detection of disagreements (if one exists), instead of overwriting the effects of each other. Per the example above, if the monitoring defender observes another forward to be executing INTERRUPT, while the first forward is executing PLAY, then the inference process for the two observations would tag these plans as both matching and not-matching at the same time, and a disagreement would be detected.

Evaluation. YOYO's first part matches plans against all observations (lines 1–8), and thus takes $O(NL)$, where N is the number of agents (observations), and L the size of the plan-hierarchy. The nested loops potentially traverse the entire plan-hierarchy $O(L)$ for each team. Since the team-hierarchy grows with N, we use that to denote its size; a traversal of the team-hierarchy is $O(N)$. The propagation down in line 14 may still traverse the entire $O(L)$ plan-hierarchy (in a theoretical worst case). The process thus takes $O(NL^2)$. Finally, the disagreement detection goes from every team in the team-hierarchy to every plan (in the worst case), thus $O(LN)$ again. Overall, YOYO's complexity is $O(NL^2)$. The key to this complexity is that YOYO only maintains coherent hypotheses. If it cannot, then a disagreement has occurred—but YOYO does not represent the underlying incoherent hypothesis. This time complexity should be contrasted with RESL's $O(L^N)$. YOYO's space complexity also compares well with RESL: With each additional agent, YOYO's space requirements grow by one node which represents the agent in H. In contrast, RESL uses an additional copy of the entire plan hierarchy for every additional agent.

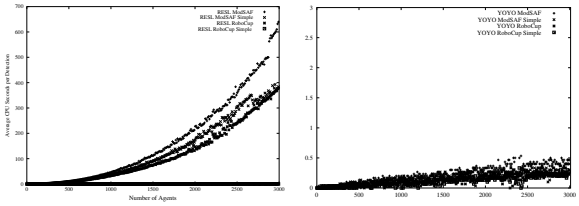

(a) RESL (b) YOYO

Figure 2: Comparing YOYO to RESL. The X axis marks the number of agents monitored; the Y axis denotes computation time in CPU seconds. ***Note (a) and (b) have different ranges on the Y axis.***

To provide empirical evaluation, the run-time performance of RESL and YOYO was compared as the number of monitored agents is scaled up. Trials were carried out in 4 different domains: ModSAF, RoboCup (both previously described), ModSAF-simple (uses smaller plan-hierarchy than ModSAF), and RoboCup-simple. RoboCup-simple involved no uncertainty in monitoring: It is not intended to be realistic, but allows exploration of the performance boundaries of RESL and YOYO. In each trial, the number of monitored agents was fixed, and then a monitoring problem (given by the observables available to the monitoring agent) was randomly generated. The same monitoring problem was given as input to RESL and to YOYO, and their execution time (matching observations, inference, and disagreement detection) was recorded. 30 trials were done for each fixed number of agents.

Figure 2 summarizes the results of the comparison, projecting average computation time (per trial) against the number of agents. Figure 2-a shows that the computation time for RESL in the ModSAF, ModSAF-simple, and RoboCup domain is clearly non-linear as predicted. A monitoring problem of 3000 agents takes, on average, 400–650 CPU seconds in these domains. However, in the RoboCup-simple domain the curve is linear since there is no uncertainty in the domain. As a result, the number of hypotheses to be considered does not grow exponentially, and the only factor in run-time complexity is observation matching and inference.

In contrast to RESL, YOYO's complexity curve (Figure 2-b) is roughly linear (due to the randomness of the monitoring problems, YOYO solves some problems quicker than others of the same size, thus the averages do not form a perfect linear curve), and its execution time is measured in *milliseconds*, rather than seconds. Note that even when monitoring 3000 agents, YOYO takes less than half a second to complete its task (compared to 400–650 seconds).

The difference in performance is not just a function of YOYO's maintenance of a single structure (rather than than the N structures maintained by RESL), but the fact that it considers only coherent hypotheses, of which only a linear number (in the size of the plan hierarchy) exist. However, YOYO's scalability comes at the ex-

pense of the ability to represent failure hypotheses: When a disagreement is detected, YOYO knows that it has occurred, but cannot identify what agents are involved, or the extent of the disagreement. Thus for more advanced monitoring tasks, such as diagnosis [15], YOYO has to be augmented by mechanisms that allow the monitoring agent to reconstruct the hypotheses underlying the disagreement. In contrast, RESL facilitates explicit reasoning about each agent separately from its teammates. It is also the case that YOYO's run-time complexity is still dominated by a key factor—the number of observed monitored agents: As long as simultaneous observations are coming in about agents, YOYO still needs to process all of them, much like the full array approach. We address ways to reduce the number of agents which require observations in the next section.

5. LIMITED CONNECTIVITY

As the number of monitored team-members increases, it becomes increasingly difficult to monitor all of them (Section 2). Furthermore, as we have seen, the run-time of monitoring algorithms can be dominated by the number of monitored agents, even with efficient algorithms such as YOYO. Thus a key question is how to guarantee detection results while limiting the number of agents that must be monitored. This section provide new bounds on the number of agents that must be monitored to guarantee disagreement detection.

We begin with a short overview of *key agents*, first introduced in [15] for their important role in disagreement detection. We show that centralized monitoring requires monitoring only (but all) key agents to guarantee sound detection. We also re-examine the role of key agents in distributed monitoring settings. Previous work has shown that sound and complete detection can take place if all agents monitor all key agents. We provide a stricter bound and show that non-key agents must monitor only a single key agent. However, key agents must still monitor each other.

Key Agents for Disagreement Detection. We abstract the underlying monitoring algorithm, such as YOYO, by providing the following notation when discussing agent A's hypotheses as to the state of an agent B: Suppose B's state is P (for instance, P is a plan selected by B). We denote by $M(A, B/P)$ the set of agent-monitoring hypotheses that A constructs based on communications from B, or inference based on B's observable behavior. In other words, $M(A, B/P)$ is the set of all A's hypotheses as to B's state, when B's state (e.g., selected plan) is P. Note that when A monitors itself, it has direct access to its own state and so $M(A, A/P) = \{P\}$.

We make the following definitions which ground our assumptions about the underlying monitoring process that implements M:

DEFINITION 1. *Given a monitoring agent A, and a monitored agent B, we say that A's monitoring of B is* complete *if for any plan P that may be executed by B, $P \in M(A, B/P)$. If A is monitoring a team of agents $B_1, \cdots, B_n$, we say that A's* team-monitoring *of the team is* complete *if A's monitoring of each of $B_1, \cdots, B_n$ is complete.*

Monitoring completeness is commonly assumed (in its individual form) in plan-recognition work, (e.g., [20, 6, 11]), and generally holds in our own applications. It means that the set $M(A, B/P)$ includes the correct hypothesis P, but will typically include other matching hypotheses besides P. Using this notation, we can now formally explore the role of key agents in disagreement detection.

Key agents have the property that their behavior when executing two given plans is sufficiently unambiguous, such that any agent that monitors them and is executing either one of the two plans can identify with certainty whether a disagreement exists between it and the key agents. Thus these agents play an important role in limiting the number of agents that must be observed to guarantee disagreement detection. We repeat here the formal definition of key agents from [15].

DEFINITION 2. *Let P_1, P_2 be two team plans. Suppose an agent A is monitoring an agent B. If $M(A, B/P_1) \cap M(A, B/P_2) = \emptyset$ for any agent A, we say that B has observably-different roles in P_1 and P_2, and call B a* key agent *in P_1, P_2. We assume symmetry so that if two plans are not observably different, then $M(A, B/P_1) \cap M(A, B/P_2) \supseteq \{P_1, P_2\}$.*

For instance, both attackers and the scout have observably-different roles in the plans executed in Example 1: The attackers land in WAIT-FOR-SCOUT, but fly in JOIN-SCOUT. The scout flies in WAIT-FOR-SCOUT, but lands in JOIN-SCOUT. Note that observably-different behavior does not imply complete disambiguation. For instance, if the scout is observed to be flying, that does not allow the hypothesis that it is executing LANDS-SCOUT, but permits two other hypotheses: WAIT-FOR-SCOUT and JOIN-SCOUT.

The key-agent is the basis for the conditions under which a team-member A_1 will detect a disagreement with a team-member A_2 using a maximal-coherence hypothesis. A_1 (executing a team-plan P_1) will detect a disagreement with a team-member A_2 (executing different team-plan P_2) if A_2 is a key agent for the plans P_1, P_2 [15, Lemma 1]. A_1 knows that it is executing P_1. If A_2 is executing P_2, and is a key-agent in P_1 and P_2, then A_1 is guaranteed to notice that a disagreement exists between itself and A_2, since A_2 is acting observably different than it would if it had been executing P_1. A_1 can now alert its teammate, diagnose the failure, etc.

Bounding the Number of Observed Agents.. As the previous section demonstrates, agents who wish to detect disagreement must focus their attention on the key agents in a team. Indeed, in centralized monitoring settings, a team-member which is monitoring itself and its teammates using maximal-coherence can focus its attention only on key-agents, since it cannot expect to detect disagreements with non-key agents, as their behavior is ambiguous. Such monitoring is guaranteed to be sound for centralized settings [15, Theorem 1]. However, in such centralized settings, all key agents must be monitored.

We now consider the case of distributed monitoring settings, where team-members monitor each other. Previous work has shown that if at least a single key agent exists for every pair of plans (i.e., the team employs an *observably-partitioned set of team plans*), and if all team-members monitor *all* key agents, then detection is not only sound, but also complete [15, Theorem 4]: At least one team-member will detect a disagreement if one occurs, and no false detections will take place. This result is of particular interest to building practical robust teams, and fortunately the conditions for it are often easy to satisfy: Teams are very often composed such that not all agents have the same role in the same plan, and in general, roles do have observable differences between them (e.g., Examples 1 & 2). Often, the set $M(A, B/P)$ can be computed offline, in advance; this allows the designer to identify the key agents in a team prior to deployment. Furthermore, any agent can become a key-agent simply by communicating its state to the monitoring agent and therefore eliminating ambiguity; thus a team can use highly-focused communications to guarantee detection.

However, the requirement that *all* key-agents be monitored prohibits deployment of scaled-up applications: First, as the size of

the team grows, limited connectivity becomes more common, since agents become more physically and logically distributed. Thus not all agents, and in particular key agents, are going to be visible. Second, the run-time of monitoring, even using the efficient YOYO algorithm presented earlier, is dominated by the need to process observations of each agent. Thus reducing the number of observed agents can improve monitoring run-time in practice.

To formally address this challenge, we define the monitoring graph of a team as follows:

DEFINITION 3. *A* monitoring graph *of a team T is a directed (possibly cyclic) graph in which nodes correspond to team-members of T, and edges correspond to monitoring conditions: If an agent A is able to monitor an agent B (either visually or by communicating with it), then an edge (A, B) exists in the graph. We say that the monitoring graph is connected, if its underlying undirected graph is connected.*

If the monitoring graph of a team is not connected, then there is an agent which is not monitored by any agent, and is not monitoring any agent. Obviously, a disagreement can go undetected in such a team: If the isolated agent chooses a plan different from its peers, it would go undetected. However, the fact that a monitoring graph is connected does not guarantee detection, since an agent that is monitoring a non-key agent would not necessarily detect a disagreement with the non-key agent.

Previous work has shown that if all agents monitor all key agents, then sound and complete detection is guaranteed [15]. However, in the case of Examples 1 & 2, this translates to an unrealistic requirement that all key agents are monitored at all times, by everyone in the team. This is difficult to guarantee in practice, for instance due communication range restrictions or limited ability to observe all key agents. Certainly, for the case of Example 2, our experiments shown that on average, a player can see at most 2–3 team-members at any given time.

The theorem below takes a step towards addressing this issue by providing more relaxed conditions on the connected nature of the monitoring graph, in particular with respect to the connectiveness of the nodes representing key agents. These conditions are: (i) every non-key agent executing a plan P monitors *a single key agent* for each possible pair of plans involving P (i.e., for each pair of plans, where one of the plans is P); and (ii) the monitoring graph for all key agents in a given pair of plans is a clique (i.e., key agents are fully connected between themselves).

THEOREM 1. *Let T be a team in which: (i) Each team-member A, executing a plan P_1, who is not a key agent for P_1, P_2 monitors a key agent for P_1, P_2; (ii) all key agents for a pair of plans X, Z monitor all other key agents for X, Z (forming a bidirectional clique in the underlying monitoring graph); (iii) the team employs an observably-partitioned set of plans; and (iv) all monitoring carried out is complete, and uses maximal-coherence. Then disagreement detection in T is sound and complete.*

PROOF. By induction on the number of agents in T. The full proof stretches over a number of pages, and is therefore presented in [13]. □

This theorem allows teams to overcome significant connectivity limitations, without sacrificing detection quality. The theorem translates into significant freedom for the designer or the agents in choosing whom (if any) to monitor; when a monitored agent is unobservable, an agent may choose to monitor another: Non-key agents need monitor only a single key agent, rather than all key agents (for every pair of plans). The upper-bound the theorem provides is more general than may seem at first glance. First, while we refer to *plans* in the theorem, there is in fact nothing in the result that uses the representation we discuss in the previous section. The theorem holds for any state feature of interest—beliefs about a shared environment, goals, etc.; it is up to the designer to pick a monitoring technique that acquires the needed information for constructing the monitoring hypotheses. Second, the theorem does not depend at all on the method by which monitoring occurs, whether by communications or by observations. Thus the connectivity of a monitoring graph does not have to be maintained visually. Some or all of the edges in the monitoring graph may actually correspond to communication links between agents.

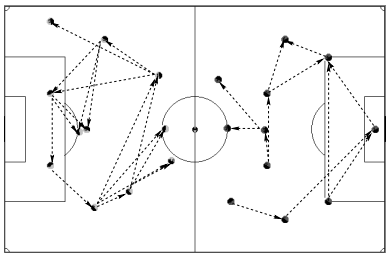

Figure 3: Monitoring graphs in actual game situation.

Though this theorem represents a significant advance in lowering the bound on the number of agents that must be monitored, all key agents must still monitor each other. This is a critical constraint in practice: We have reconstructed the visual monitoring graph in *thousands* of RoboCup game situations, to find that even with this new bound, sound and complete disagreement detection would have been possible without communications only in small percentage (approximately 5%). Typically, each RoboCup player can only see 2–3 key agents. To illustrate, Figure 3 shows the monitoring graph of two teams overlayed on a screen-shot of an actual game situation. The monitoring graphs do not guarantee sound and complete disagreement detection under the known bound.

This empirical constraint raises the bar on the challenge to find a lower-bound on the number of agents that must be monitored to guarantee detection. We have strong evidence that leads us to believe that in fact a much lower bound than that which is described above may be possible, but as of now, it remains unproven. We believe that it may be possible to guarantee sound and complete detection in all cases where each key agent is either monitored or is monitoring a single other key agent (rather than all of them). If so, this would translate to guaranteeing failure detection in over 70% of the thousands of RoboCup monitoring cases we have examined.

6. RELATED WORK

Most closely related to YOYO is our own earlier work on the YOYO* visualization algorithm [14], deployed in different domains. YOYO differs from its predecessor in several important ways: (i) YOYO utilizes a non-probabilistic representation scheme, while YOYO* utilizes a probabilistic representation; (ii) YOYO is focused on disagreement detection, and in fact does quite poorly in visualization tasks on which YOYO* performs well; and (iii) YOYO includes explicit checks for failure detection. Also, while YOYO* has revealed a potential tradeoff between expressivity and scalability in visualization algorithms, YOYO provides the same detection results as the full array approach.

$RESC_{team}$ [20] is a multi-agent plan-recognition scheme which implicitly uses coherence as a key constraint in representation. $RESC_{team}$ represents only a single coherent hy-

potheses, while YOYO represents all coherent hypotheses. However, $RESC_{team}$ can reason about the assignment of agents to roles/subteams, while YOYO assumes this knowledge is given apriori. Intille and Bobick [11] rely entirely on coordination constraints among football players to recognize team-tactics. Devaney and Ram [6] use pattern-matching to recognize team-tactics in military maneuvers. Huber [10] proposes using automatically built plan-recognition Bayesian networks to allow agents to coordinate without communications. In contrast to our work, all of these approaches utilize probabilistic representations, do not detect failures, and have not explicitly addressed scalability.

Durfee [7] discussed decision-theoretic and heuristic methods for reducing the amount of knowledge that agents consider in coordinating. The methods include pruning nested (recursive) models, using communications to alleviate uncertainty, using hierarchies and abstractions, etc. Our work complements Durfee's in several way: (i) we focus on monitoring in teams of cooperating (rather than self-interested) agents, allowing exploitation of socially-attentive means; (ii) we provide bounds on the number of agents that must be monitored; (iii) Durfee's work focuses on reducing computational loads in monitoring each single agent, while our work on YOYO demonstrates significant savings that are achieved when considering all monitored agents together.

Dellarocas and Klein [5] present complementary failure detection techniques for contract-net protocol interactions, using a centralized monitoring scheme that monitors all agents using pre-built fault-models. While their techniques complement ours, they do not address limited connectivity, or scalability concerns.

7. DISCUSSION AND FUTURE WORK

Multi-agent literature has often emphasized that an agent must monitor other agents in order to carry out its tasks. However, as the numbers of agents in deployed teams is scaled up, two key challenges are raised: (i) the number of hypotheses a monitoring agent has of its peers increases exponentially in the number of agents, and thus large teams cannot be monitored effectively; and (ii) agents become more physically and logically separated, and so are less visible (less connected) to the monitoring agent. Thus not all agents can be monitored.

This paper has begun to address these challenges, in the context of a critical monitoring task—detection of critical disagreements between teammates. We have presented YOYO, a linear-time sound disagreement detection algorithm, exploiting the organizational structure of the monitored team to monitor any number of agents in a single structure. We addressed the challenge of limited connectivity of agents by showing new bounds on the number of agents that must be monitored in a team to guarantee sound and complete disagreement detection in distributed and centralized monitoring settings. This new result in the distributed case shows that in fact high-quality detection can be guaranteed despite situations where many agents do not monitor others, or are not monitored by others. Our empirical and analytical evaluation demonstrates that the results indeed constitute significant advancement to the practice of monitoring large teams.

Acknowledgements. We thank Milind Tambe, Jeff Rickel, George Bekey, Victor Lesser, David Pynadath, Orna Raz and Dan Peleg for useful discussions. As always, thanks to K. Ushi.

8. REFERENCES

[1] M. S. Atkin, G. W. King, and D. L. Westbrook. Hierarchical agent control: A framework for defining agent behavior. In *Proceedings of the International Conference on Autonomous Agents*, 2001.

[2] T. Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, Georgia Institute of Technology, 1998.

[3] J. J. Burns, E. Salas, and J. A. Cannon-Bowers. Team training, mental models, and the team model trainer. In *Advancements in Integrated Delivery Technologies*, Denver, CO, 1993.

[4] P. R. Cohen and H. J. Levesque. Teamwork. *Nous*, 35, 1991.

[5] C. Dellarocas and M. Klein. An experimental evaluation of domain-independent fault-handling services in open multi-agent systems. In *Proceedings of the International Conference on Multiagent Systems*, pages 95–102, Boston, MA, 2000. IEEE Computer Society.

[6] M. Devaney and A. Ram. Needles in a haystack: Plan recognition in large spatial domains involving multiple agents. In *Proceedings of the National Conference on Artificial Intelligence*, pages 942–947, Madison, WI, 1998.

[7] E. H. Durfee. Blissful ignorance: Knowing just enough to coordinate well. In *Proceedings of the International Conference on Multiagent Systems*, pages 406–413, 1995.

[8] B. J. Grosz and S. Kraus. Collaborative plans for complex group actions. *Artificial Intelligence*, 86:269–358, 1996.

[9] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *distributed computing*, 37(3):549–587, 1990.

[10] M. J. Huber and E. H. Durfee. Deciding when to commit to action during observation-based coordination. In *Proceedings of the International Conference on Multiagent Systems*, pages 163–170, 1995.

[11] S. S. Intille and A. F. Bobick. A framework for recognizing multi-agent action from visual evidence. In *Proceedings of the National Conference on Artificial Intelligence*, pages 518–525. AAAI Press, July 1999.

[12] N. R. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75(2):195–240, 1995.

[13] G. A. Kaminka and M. Bowling. Towards robust teams with many agents. Technical Report CMU-CS-01-159, Carnegie Mellon University, 2001.

[14] G. A. Kaminka, D. V. Pynadath, and M. Tambe. Monitoring deployed agent teams. In *Proceedings of the International Conference on Autonomous Agents*, pages 308–315, 2001.

[15] G. A. Kaminka and M. Tambe. Robust multi-agent teams via socially-attentive monitoring. *Journal of Artificial Intelligence Research*, 12:105–147, 2000.

[16] S. Kumar, P. R. Cohen, and H. J. Levesque. The adaptive agent architecture: Achieving fault-tolerance using persistent broker teams. In *Proceedings of the International Conference on Multiagent Systems*, pages 159–166, Boston, MA, 2000. IEEE Computer Society.

[17] N. Lesh, C. Rich, and C. L. Sidner. Using plan recognition in human-computer collaboration. In *Proceedings of the International Conference on User Modelling (UM-99)*, Banff, Canada, 1999.

[18] A. Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1990.

[19] L. E. Parker. ALLIANCE: An architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, April 1998.

[20] M. Tambe. Tracking dynamic team activity. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, August 1996.

[21] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

[22] M. Tambe, W. L. Johnson, R. Jones, F. Koss, J. E. Laird, P. S. Rosenbloom, and K. Schwamb. Intelligent agents for interactive simulation environments. *AI Magazine*, 16(1), Spring 1995.

[23] M. Tambe, G. A. Kaminka, S. C. Marsella, I. Muslea, and T. Raines. Two fielded teams and two experts: A robocup challenge response from the trenches. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 1, pages 276–281, August 1999.

[24] M. Tambe, D. V. Pynadath, N. Chauvat, A. Das, and G. A. Kaminka. Adaptive agent integration architectures for heterogeneous team members. In *Proceedings of the International Conference on Multiagent Systems*, pages 301–308, Boston, MA, 2000.