From Centralized to Distributed Selective Overhearing

Gery Gutnik and **Gal A. Kaminka**
The MAVERICK Group
Computer Science Department
Bar-Ilan University, Israel
{gutnikg,galk}@cs.biu.ac.il

Abstract

Overhearing is an approach for monitoring open, distributed, multi-agent systems by listening to the routine communications taking place within them. Previous investigations of overhearing assumed that all inter-agent communications are accessible to a single overhearing agent. However, as multi-agent systems grow both in size and distribution two problems arise. First, in large-scale settings, an overhearing agent cannot monitor all agents and their conversations, and must therefore be selective in carefully choosing its targets. Second, a single overhearer would encounter difficulties overhearing agents acting in a geographically-distributed environment. This paper tackles these challenges by addressing distributed teams of overhearing agents involved in selective overhearing. Building on prior work on centralized selective overhearing, we consider the consequences of transitioning from overhearing teams working in a centrally-coordinated manner to distributed overhearing teams. In doing so, we distinguish the various factors influencing the level of distribution within these teams and determine their importance in terms of effective selective overhearing.

Introduction

One of the key challenges in multi-agent systems (MAS) has always been, and still remains, the task of monitoring. Traditional monitoring relied on reports received from monitored agents (Klein & Dellarocas 1999), but this approach is problematic in monitoring open distributed MAS (Kaminka, Pynadath, & Tambe 2002). Alternatively, monitoring can be done by listening to the routine communications of the monitored agents. Here, the monitoring agent uses overheard communications to independently assemble and infer the monitoring information about the MAS. This approach is called monitoring via *overhearing* (Novick & Ward 1993; Aiello *et al.* 2001; Legras 2002; Rossi & Busetta 2005).

In previous investigations of overhearing, a *single* overhearing agent was assumed to overhear *all* inter-agent communications of the monitored agents. However, as multi-agent systems grow in size and distribution (as in real-world settings) two problems arise. First, in large-scale settings, overhearing resources are bound to be limited. Thus, the overhearing agent might not be able to overhear all inter-agent communications, and must therefore be *selective*, carefully choosing its targets. Second, a single overhearer monitoring *geographically-distributed* agents will only be exposed to a subset of their communications. Multiple collaborating overhearing agents can be deployed to increase the coverage of overheard targets. Thus, overhearing teams, acting in a distributed manner, should be considered.

Those two challenges are not necessarily related. On one hand, a single overhearer can face the problem of selectivity. On the other hand, a distributed team of overhearing agents might overhear all inter-agent communications, facing only the problem of effectively dividing the targets between themselves. However, giving the recent interest in open MAS which are both large-scale and distributed, we tackle these two problems as one.

This paper builds and significantly extends prior work on centralized selective overhearing (Gutnik & Kaminka 2006), to address these challenges. Centrally-coordinated overhearing teams, performing selective overhearing, can be viewed as equivalent to a single overhearer facing the selectivity problem. In contrast, here, we consider distributed teams of collaborating overhearing agents involved in selective overhearing. In fact, we consider a gradual transition from centrally-coordinated teams to distributed ones by incrementally decreasing the overhearers' dependence on centralized mechanisms.

We determine the various factors influencing the distribution level of the overhearing team, and empirically evaluate their importance in terms of effective selective overhearing. Here, we found that the available visibility level–knowledge of where and when inter-agent communications take place in the monitored settings–causes a significant difference in the effectiveness of the overhearing team. On the other hand, a costly-maintained shared memory has no effect over the use of easily-maintained individual memory. In between those two, we found that collisions–situations where two or more overhearers overhear the same target–although influencing the effectiveness of selective overhearing, can be only partially solved. This surprising conclusion is important since collision avoidance is a time-consuming activity that can be problematic in real-time settings.

Background

Overhearing is fast gaining attention as a generic method for monitoring agents based on their routine communications. Kaminka et al. (2002) used plan recognition in overhearing a distributed team of agents, which are collaborating to

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

carry out a specific task. Knowing the plan of this task and its steps, the overhearer uses overheard messages as clues for inferring the state of different team-members. Aielo et al. (2001) describe collaborative settings in which the overhearing agent may act on overheard messages, to offer expert assistance to the problem-solving process of the communicating agents. In the work by Rossi and Busetta (2005), the overhearer tracks changes in organizations and issues alerts on detected inconsistencies.

All these previous investigations assume a single overhearing agent. This assumption can be challenged in open distributed multi-agent settings. In such settings, agents are often geographically distributed. Therefore, a single overhearer, even located in a central position, will be able to overhear only a subset of inter-agent communications held in the monitored system.

Multiple overhearing agents can be deployed to increase the coverage of overheard communications. Few previous investigations addressed the case of multiple overhearing agents. In the work by Novick and Ward (1993), pilots maintain their situational awareness not only by conversing with an air-traffic controller, but also by overhearing the conversations of other pilots. Legras (2002) uses overhearing as a method for allowing agents to keep track of organizational memberships.

However, both these works assume different overhearing agents to be non-collaborative entities, committing overhearing out of their own self interest. A non-collaborative overhearer is equivalent to a single overhearer working on its own. Accordingly, these multiple non-collaborative overhearers are still facing the problems of a single centrally-located overhearing agent.

In contrast, we examine teams of overhearing agents collaboratively monitoring multi-agent settings. Continuing our previous work on centrally-coordinated overhearing teams (Gutnik & Kaminka 2006), we explore the transition from centralized to distributed teams of overhearers.

An Empirical Approach

Our work presents an empirical study of selective overhearing. In doing so, we first model the monitored system addressing its specific characteristics, and then use this model to simulate inter-agent communications characterizing the monitored settings. Based on this simulation, we empirically evaluate and compare the effectiveness of different overhearing teams. Due to lack of space, we do not provide here a comprehensive description of our simulation environment, but refer the reader to (Gutnik & Kaminka 2006) for additional details.

However, throughout this article, we do address certain details of the simulation that are critical for understanding the paper. In this section, we explain the concept of *overhearing policy*. We assume that an overhearing policy is what controls and coordinates multiple overhearing agents within the overhearing team. Since we address the case of selective overhearing, only a portion of inter-agent conversations can be overheard. The policy determines the assignment of overhearing agents to conversations in the monitored system. We distinguish between two possible assignment types. In the first, the overhearing agent is assigned a single conversation chosen from all the conversations in the monitored system. In the second, each overhearing agent can focus on a single communicating agent (referred as its *target*), overhearing all of its conversations. Both types of scenarios exist in the real-world. In our research, we focus on the latter agent-target assignments, where all conversations simultaneously carried out by the target are overheard, as long as the overhearing agent is currently listening.

Different overhearing policies can be proposed. In this work, we distinguish between centralized and distributed policies. According to centralized policies, each overhearing agent within the team is assigned its target by a centralized decision. In contrast, distributed policies allow each overhearer to choose its target. The degree up to which a single overhearing agent within the overhearing team depends on others in this decision determines the level of distribution of the entire team.

Centralized Selective Overhearing

Centralized overhearing policies determine for each overhearing agent within the overhearing team its assigned target, assuming a perfect level of coordination between the teammates. Thus, in centralized selective overhearing, the main challenge is not the effective coordination of the overhearing agents, but finding a good assignment of targets to overhearers.

Our previous work proposed to use organizational knowledge as a basis for this assignment (Gutnik & Kaminka 2006). We focus on hierarchical organizations which are both common in multi-agent applications and real-world corporates. Using the well-defined characteristics of such organizations from social science (Best, de Valence, & Langston 2003; Dewan, Seidmann, & Sundaresan 1997; Friebel & Raith 2004; Gannon & Newman 2001; Jensen 2003), we determine effective centralized policies for selective overhearing. These characteristics are briefly summarized below:

- **Distribution Characteristic [Where do conversations take place?]**. The volume of conversations in a given hierarchy level depends on the number of agents associated with this hierarchy level (Jensen 2003). The latter, in turn, is determined by the hierarchical structure of the organization. In traditional hierarchies, the number of agents associated with each hierarchy level is smaller in higher hierarchy levels. Thus, most conversations in such organizations are held between agents in the lower hierarchical levels, simply because most agents are associated with these levels.
- **Scope Characteristic [What do agents discuss?]**. Social science studies distinguish between three types of information: strategic, tactical and operational (Best, de Valence, & Langston 2003; Gannon & Newman 2001). These different types of information are associated with different organizational hierarchy levels. The top levels handle strategic information, the middle levels are responsible for the tactical information, and the lower levels handle operational information.
- **Span Characteristic [With whom agents communicate?]**. Communications in hierarchical organizations

reflect the restricted flow of information in such organizations: either top-down or bottom-up (Dewan, Seidmann, & Sundaresan 1997; Friebel & Raith 2004; Jensen 2003). Accordingly, agents communicate mostly with their peers, subordinates and their close superiors. Meaning that most communications are held between agents of the same hierarchy levels or between agents in relatively close hierarchical levels.

Taking these characteristics into account and assuming that it is more important to overhear strategic information rather than the operational one, we discovered two effective centralized policies. On one hand, a *value* policy targets agents of higher hierarchy levels. These agents are involved in small amount of highly valuable conversations. On the other hand, a *volume* policy targets agents involved in higher amount of conversations. These are usually agents in lower hierarchy levels, that involved in conversations of a low information value.

The experiments on these centralized policies showed a classical value-volume trade-off. This trade-off was found to be surprisingly robust to many characteristics of the monitored organizations. Further studying centralized policies, we have come to another surprising conclusion: Combining the two types of policies (such that some agents follow the value type policy, while others follow the volume type policy) improves their individual performance. Moreover, the combined policies have been found to be effective unrelated to any characteristic of the monitored organization (even those that influence each policy separately).

Distributed Selective Overhearing

Using the combined value-volume policy as a baseline, we now consider the transition from centralized to distributed policies. In this paper, we gradually decrease the inter-dependence of each overhearing agent on its teammates. In doing so, we evaluate and compare the effectiveness of overhearing policies, thus examining the effects of moving from centralized to distributed selective overhearing. Our goal is to determine what are the factors of teammates' inter-dependency that influence the effectiveness of distributed selective overhearing, and what are the factors whose effect can be ignored.

Centralized vs. Distributed Policies

The baseline centralized policy, i.e. the combined value-volume policy, relies on the three following assumptions: (i) full visibility; (ii) shared memory and (iii) collision avoidance. We explain those in details below.

Visibility is defined as knowledge of where and when conversations take place (without knowing their content). Full visibility assumes knowledge of all conversations in the monitored system.

Knowing the number of conversations that each potential target is involved in, this centralized policy assigns targets to overhearing agents based on a value-volume decision. For each overhearing agent a decision is made whether it is better to overhear x conversations by $agent_i$–the target of value policy–or y conversations by $agent_j$–the target of volume policy. Remember that due to the value-volume trade-off y will usually be greater than x, but the value of conversations by $agent_i$ will be higher than the ones by $agent_j$. Thus, the dilemma.

This value-volume decision is made based on monitored agents' past performance–the average value of conversations in which the agents were overheard earlier. In centralized policies, this memory is assumed to be shared.

Finally, centralized policies assume collision avoidance. Since overhearing agents are centrally coordinated, collisions, i.e. situations where a target is overheard by two or more overhearers, can easily be avoided.

However, the assumptions of full visibility, shared memory and collision avoidance are only possible in centralized selective overhearing. Moving towards distributed policies means breaking these assumptions. Thus, transitioning from centralized to distributed policies, we incrementally decrease the inter-dependence of overhearing agents on their teammates along these three dimensions. This transition is summarized in Table 1.

Dimensions	**Centralized**	**Distributed**
Memory	shared	individual
Visibility	full	group/agent
Collision Avoidance	yes	no

Table 1: **Centralized vs. Distributed Policies**

Addressing the memory dimension, we consider the use of an individual memory instead of a shared one. With respect to visibility dimension, group and agent visibilities are examined. Group visibility assumes that overhearing agents are only aware of conversations that are carried out by their targets, while agent visibility assumes that overhearing agent is only aware of conversations committed by its target (but not by the targets overheard by other overhearers). As for collision avoidance, we distinguish between two extremes. On one hand, the case where all collisions are avoided, and, on the other hand, the case where collisions are allowed and there is no collision avoidance mechanism applied.

Obviously, the fully distributed policy assumes agent visibility, individual memory and no collision avoidance. In contrast, we remind the reader that fully centralized policy assumes full visibility, shared memory and collision avoidance. Any other combination is considered to be partially centralized and partially distributed.

Memory Dimension

Exploring overhearing policies with respect to memory dimension, we first assume full visibility and collision avoidance. Later on, we will challenge this assumption.

We introduce two policies–*FullVis-ShrdMem-CollAvd* and *FullVis-IndMem-CollAvd*. The first is in fact the baseline centralized policy discussed earlier. As explained above, it uses shared memory. In contrast, the *FullVis-IndMem-CollAvd* policy uses individual memory. Each overhearing agent has only the memory of conversations it overheard in the past, without any knowledge on conversations overheard by other overhearers. Thus, the basis for the value-volume decision for each overhearing agent lies in its own past experience.

The distinction between the two types of memory is important mainly due to the requirements of their maintenance.

Maintaining a consistent shared memory requires broadcasting information value of each overheard conversation to all overhearing agents, whereas individual memory requires no communications between overhearing agents.

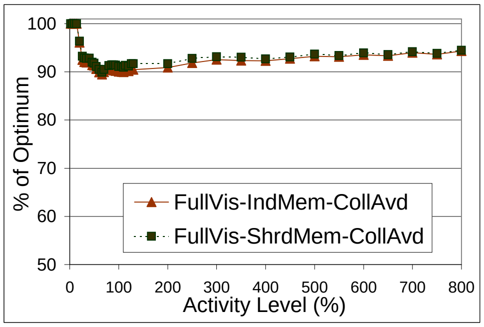

Figure 1: **Full Visibility: Memory Effect**

Figure 1 shows the performance of the *FullVis-ShrdMem-CollAvd* and the *FullVis-IndMem-CollAvd* policies. The values on the X-axes show the activity level, i.e. the number of conversations taking place in the monitored system as a ratio to the number of potential targets. The Y-axes measures performance as percentage of the theoretical optimal performance. Each data point in the graphs corresponds to an average of 50 independent experiments. We refer the reader to (Gutnik & Kaminka 2006) for additional details on the experimental settings and the calculation of optimum.

The results in Figure 1 show that overhearing teams perform better using shared memory than in the case where an individual memory is applied. However, the gap between the performance of the two policies appears to be insignificant.

Visibility Dimension

We are now going to confront the assumption of full visibility. Full visibility is usually unobtainable. Thus, we now assume group visibility. According to this visibility, overhearing agents are only aware of the conversations carried out by agents that are being overheard. Still, each overhearing agent is aware of both conversations by its target and of conversations by targets overheard by other overhearers. Group visibility restricts the visibility of the monitored system to the limits of an overhearing group. However, it still assumes that overhearing agents collaborate informing each other on overheard conversations.

Figure 2 shows the performance of the two policies, discussed in previous section, with respect to group visibility. These policies are called *GroupVis-ShrdMem-CollAvd* and *GroupVis-IndMem-CollAvd* respectively. Comparing these policies to the ones relying on full visibility, we can see a degradation in performance. This conclusion is straightforward to some extent since the transition from full to group visibility reduces the knowledge on the monitored system.

However, a more surprising result comes from contrasting the *GroupVis-ShrdMem-CollAvd* and the *GroupVis-IndMem-CollAvd* policies. In contrast to the performance of the two policies with respect to full visibility, it is the policy using individual memory that outperforms the policy using shared memory in group-visibility settings.

Group visibility depends on the number of overhearing agents. In the following set of experiments, we tested the

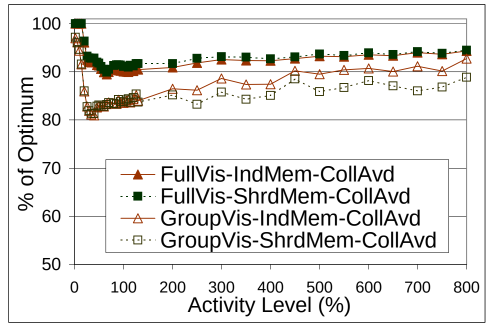

Figure 2: **Group Visibility: Memory Effect**

surprising behavior of the *GroupVis-ShrdMem-CollAvd* and the *GroupVis-IndMem-CollAvd* policies with respect to different number of overhearers–measured as coverage %, i.e. the ratio of the number of overhearing agents in the overhearing team to the number of potential targets.

Figure 3-b shows the performance of the two group-visibility policies with respect to the default overhearing coverage of 30%, while Figures 3-a and c show the performance results under condition of two additional coverage ratios (5% and 50% respectively). Figures 3-a,b and c all show that the surprising behavior of individual memory outperforming the shared memory, found in group-visibility settings, is consistent in all levels of overhearing coverage.

Moving further towards distributed selective overhearing, we now examine agent visibility. As opposed to group visibility, here, an overhearing agent is only aware of conversations committed by its target, but not by targets overheard by other overhearers. However, the overhearing agent has a limited knowledge on conversations held by other communicating agents: it is aware of conversations that these agents hold with its target.

Figures 4-a,b,c compare selective overhearing policies in group- and agent-visibility settings. Considering that (as shown in Figures 3-a,b,c) individual memory outperforms shared memory, we only evaluate policies based on individual memory. This step also supports our desired goal– distributed overhearing teams where overhearers depend as little as possible on their teammates.

Again, we see that transition from group visibility to agent visibility results in poorer performance. In all coverage ratios (Figures 4-a,b,c), the *GroupVis-IndMem-CollAvd* policy outperforms the *AgentVis-IndMem-CollAvd* policy. This result is similar to the transition from full to group visibility. In both cases, the available information on the monitored system becomes smaller with each transition.

Collision Avoidance Dimension

The final subject is the influence of collision avoidance on performance of overhearing teams. Here, we assume agent visibility and the use of individual memory.

Collision is defined as a state where two or more overhearing agents target the same communicating agent at the same time. In centralized settings, such collisions can easily be avoided since all agents are coordinated by a single centralized authority. In contrast, overhearing agents operating in distributed settings must handle collisions in a distributed fashion. Therefore, collision avoidance might be a

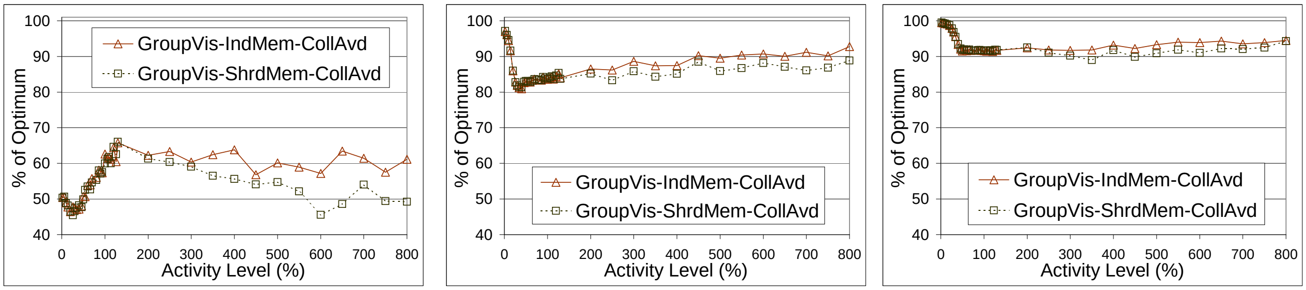

(a) Coverage: 5% (b) Coverage: 30% (c) Coverage: 50%

Figure 3: **Group Visibility - Shared vs. Individual Memory - Coverage Effect**

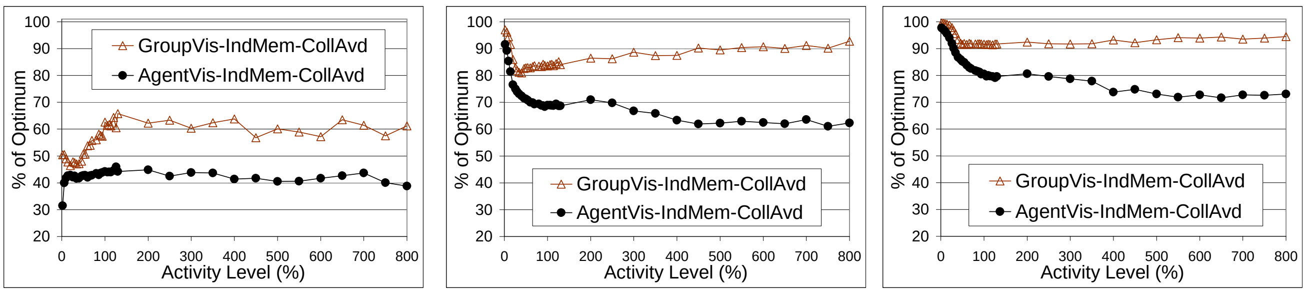

(a) Coverage: 5% (b) Coverage: 30% (c) Coverage: 50%

Figure 4: **Group vs. Agent Visibility - Individual Memory - Coverage Effect**

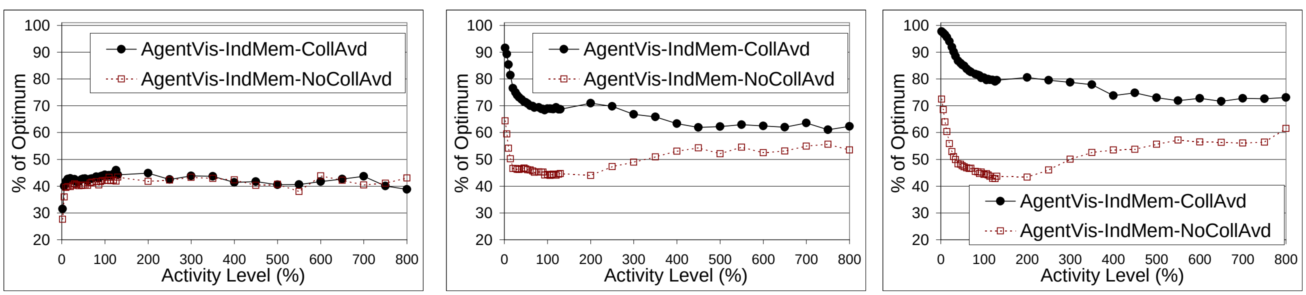

(a) Coverage: 5% (b) Coverage: 30% (c) Coverage: 50%

Figure 5: **Agent Visibility - Collision Avoidance Effect**

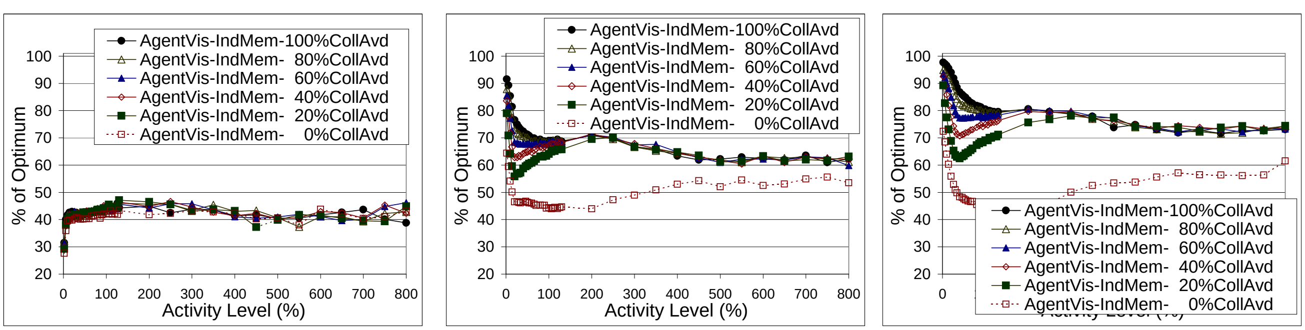

(a) Coverage: 5% (b) Coverage: 30% (c) Coverage: 50%

Figure 6: **Agent Visibility - Collision Avoidance % Effect**

time-consuming activity for an overhearing agent.

Figures 5-a,b,c show the performance of selective overhearing policies with respect to the two extremes: one where collisions are always avoided, i.e. *AgentVis-IndMem-CollAvd*, and other where collisions are allowed–*AgentVis-IndMem-NoCollAvd*. Moreover, in the *AgentVis-IndMem-NoCollAvoid* policy, even if a collision is detected, i.e. overhearing agent chooses a target already overheard by another overhearer, the overhearer does not change its selected target. On the other hand, in the *AgentVis-IndMem-CollAvoid* policy, in case of collision, the overhearing agent chooses a second best target, then the third-best target and so on in case of additional collisions. Finally, if it has no more desired targets, it simply chooses a random one.

Since collisions depend on the number of overhearers, Figures 5-a,b,c show the performance of these two overhearing policies with respect to different overhearing coverage percentage (5%, 30% and 50% respectively). Indeed, it can be seen that collision avoidance has a significant effect on the performance of overhearing policies. Furthermore, this effect becomes more significant in overhearing teams containing more overhearing agents. Although in settings with low overhearing coverage no difference is witnessed (with or without collision avoidance), it becomes highly significant as the number of overhearing agents increases due to the higher probability of collisions in such settings.

As a final step of our research, we examined different levels of collision avoidance. Here, we compare overhearing policies where overhearing agent seeks to avoid only a certain amount of occurring collisions.

Figures 6-a,b,c show the corresponding results. The *AgentVis-IndMem-p%Coll Avoid* policy represents a selective overhearing policy according to which the overhearing agent chooses to avoid only p percent of occurring collisions. This p percent ranges from 0% to 100% (with a 20% hop). The *AgentVis-IndMem-0%CollAvoid* and the *AgentVis-IndMem-100%CollAvoid* policies correspond to the *AgentVis-IndMem-NoCollAvoid* and the *AgentVis-IndMem-CollAvoid* policies above.

Again, it can be seen that in settings with low overhearing coverage collision avoidance has no significant effect (Figures 6-a). Similarly, different levels of collision avoidance lose their impact as activity levels rise (Figures 6-b and c). However, in low and medium activity levels, collision avoidance and its percentage are important. Still, it can be seen that the performance boost, achieved due to the higher levels of collision avoidance, is not always significant.

Summary & Future Work

Transitioning from centralized to distributed policies, we incrementally decreased inter-dependency of overhearing agents on their teammates along three dimensions. With respect to memory, we considered the transition from *shared* to *individual* memory. With respect to visibility, we explored the transition from *full*, to *group*, to *agent* visibility. Finally, we considered cases where *all*, *some*, or *none* of the overhearing collisions were solved.

Our experiments studied the changes in performance of overhearing policies caused by the changes in degree of overhearing agents' inter-dependence along these three dimensions. We come to conclusion that some of the changes have greater effect than others. For instance, the transition from shared to individual memory does not influence the performance of selective overhearing. This conclusion is important since maintaining shared memory causes frequent communications between overhearing agents, and thus might burden the communication network.

In addition, we show that it is sufficient to solve only some of the collisions and not all of them. Again, this conclusion is important since collision avoidance is a time-consuming activity that might be costly in real-time settings. Still, the transition from centralized to distributed policies causes a significant decrease in performance when these is a decrease in visibility. Given these conclusions, it is clear that most efforts in collaborative overhearing should be on mechanisms that facilitate greater visibility.

Acknowledgments. This research was supported by ISF Grant #1211/04. Thanks to K. Ushi.

References

Aiello, M.; Busetta, P.; Dona, A.; and Serafini, L. 2001. Ontological overhearing. In *Proceedings of ATAL-01*.

Best, R.; de Valence, G.; and Langston, C. 2003. *Workplace Strategies and Facilities Management*. Elsevier.

Dewan, R.; Seidmann, A.; and Sundaresan, S. 1997. Communications in hierarchical organizations and standards policies for information technology. *International Journal of Electronic Commerce* 1(3):43–64.

Friebel, G., and Raith, M. 2004. Abuse of authority and hierarchical communications. *RAND Journal of Economics* 35(2):224–244.

Gannon, M. J., and Newman, K., eds. 2001. *The Blackwell Handbook of Cross-Cultural Management*. Blackwell Publishing.

Gutnik, G., and Kaminka, G. 2006. Experiments in selective overhearing of hierarchical organizations. In van Eijk, R. M.; Flores, R.; and Huget, M. P., eds., *Agent Communications, LNAI Book*. Springer-Verlag. In Press.

Jensen, M. 2003. Organizational communication: A review. R&D Report 1/2003, Agderforskning, Norway.

Kaminka, G.; Pynadath, D.; and Tambe, M. 2002. Monitoring teams by overhearing: A multi-agent plan-recognition approach. *JAIR* 17:83–135.

Klein, M., and Dellarocas, C. 1999. Exception handling in agent systems. In *Proceedings of the Third International Conference on Autonomous Agents (Agents-99)*.

Legras, F. 2002. Using overhearing for local group formation. In *Proceedings of AAMAS-02*.

Novick, D., and Ward, K. 1993. Mutual beliefs of multiple conversants: A computational model of collaboration in air traffic control. In *Proceedings of AAAI-93*, 196–201.

Rossi, S., and Busetta, P. 2005. With a little help from a friend: Applying overhearing to teamwork. In *Proceedings of MOO-05*.