An Empirical Study of Selective Overhearing in Hierarchical Organizations*

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

Abstract

Overhearing is a plan recognition approach for monitoring multi-agent systems, by listening to the routine inter-agent communications. Previous investigations of overhearing explored an extensive set of techniques for overhearing, mostly relying on the assumption that all inter-agent communications are accessible to the overhearing agent. However, in real-world settings, overhearing resources are limited, and thus the overhearing agent must be selective in carefully choosing the conversations it will overhear. This paper presents an empirical study of such selective overhearing. We focus on overhearing hierarchical organizations that are common in the real-world settings. We first present a model of the conversations expected in such organizations. We then present the results of extensive experiments with several overhearing strategies, particularly suited for such organizations. Based on these extensive experiments, we are able to isolate the parameters influencing their behavior. We reach several qualitative conclusions.

1 Introduction

Recent multi-agent systems (MAS) are often built applying an *open, distributed* design. These systems involve various challenges of monitoring geographically-distributed and independently-built multiple agents. *Monitoring by overhearing* [Kaminka *et al.*, 2002] has been found to provide a powerful monitoring approach particularly suited for open distributed MAS settings. Here, an overhearing agent monitors the exchanged communications between the system's agents. It uses these observed communications to independently assemble and infer the needed monitoring information using key-hole plan recognition.

Previous investigations of overhearing have demonstrated a range of overhearing techniques and applications. Overhearing was used to maintain situational and organizational awareness [Novick and Ward, 1993; Legras, 2002; Rossi and Busetta, 2004], for monitoring progress [Kaminka *et al.*, 2002], and for discovering opportunities for providing advice [Aiello *et al.*, 2001; Busetta *et al.*, 2002]. Formal investigations include [Gutnik and Kaminka, 2004; Platon *et al.*, 2004].

Although these previous investigations provided an extensive set of overhearing techniques, most rely on the ability of an overhearing agent to overhear *all* inter-agent communications. However, this assumption is often challenged in real-world settings, and in *large-scale* MAS. Instead, the overhearing agent has limited resources, and may only overhear a subset of conversations committed in monitored organizations. Consequently, the overhearer must be selective in choosing which conversations must be overheard.

We present an empirical study of selective overhearing in pyramidal-hierarchical organizations, often found in the real-world (e.g., many corporates). To carry out this study, we first introduce a model of the conversations in hierarchical organizations. This model specifies (i) the characteristics of conversations in such organizations; and (ii) the overhearing strategies suitable for these settings. Based on the model, we simulated thousands of conversations in hundreds of pyramidal-hierarchical organizations. We contrasted the performance of different selective overhearing strategies, in terms of the value of information that they are able to collect.

In particular, we contrast strategies that focus on largest value (i.e., most important agents), with those that focus on overhearing the largest volume (i.e., most active agents). The results lead to several qualitative conclusions as to overhearing in pyramidal organizations.

First, as communication volume increases, largest-value strategies show a surprising tailed parabolic curve in terms of value of information: With little volume, or much volume, these strategies gather more information than with moderate amounts of activity. In contrast, largest-volume strategies show linear growth as the activity increases. A second conclusion we draw is that the relative performance of these two strategies is dependent on the factor by which the information at top levels of organization is more important than at lower levels. The more important the information at the top is, the more a largest-value policy makes sense. However, there are cases where a trade-off exists between the two strategies. Finally, we have found that the number of overheard agents does *not* impact the relative performance of the two strategies, nor is it affected by the number of levels in the organizational hierarchy.

*This research was supported in part by BSF Grant #2002401.

2 Background & Motivation

Most previous work has investigated overhearing in settings where all relevant conversations were accessible to the overhearing agent. Nowick and Ward [1993] show an early use of cooperative overhearing, where pilots maintain their situational awareness not only by conversing with an air-traffic controller, but also by overhearing the conversations of other pilots. Similarly, Legras [2002] uses overhearing as a method for maintaining organizational knowledge. Here, agents overhear others to keep track of organizational memberships. Aielo et al. [2001] and Busetta et al. [2002] 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. Our work seeks to allow these tasks to be carried out even when not all communications are overheard.

Few previous investigations have touched overhearing where some conversations may not be accessible. However, these focused on lossy cases, where due to noise, some messages may not reach the overhearing agent. In contrast, we deal with the case where complete conversations are unknown to the overhearing agent. Kaminka et al. [2002] used plan recognition in overhearing a distributed team of agents. They evaluated the use of their system in lossy settings, and showed that the performance of the overhearing agent drops when messages are lost. Rossi and Busetta [2004] applied overhearing to monitor changes in MAS settings caused by transition from one state to another. They mention that lost messages can cause inconsistencies.

Our previous work [Gutnik and Kaminka, 2004] also discussed lost messages, as part of a formal approach to overhearing. Here, we defined *conversation systems*, the set of conversations generated by an organization, and developed algorithms for conversation recognition, a first key step in overhearing. These algorithms are suited for lossy settings, in that they allow for a limited set of communicative acts to be lost, while still allowing the conversation to be recognized. Our current work extends and specializes conversation systems to hierarchical organizations, based on studies of such organizations [Dewan *et al.*, 1997; Friebel and Raith, 2004]. Moreover, we tackle the case where only a limited set of complete conversations can be overheard.

3 Overhearing in Hierarchical Organizations

Overhearing extracts information from conversation systems [Gutnik and Kaminka, 2004], the set of conversations generated by an organization. Thus, conversation systems change based on the type of organization that is being overheard, and, in turn, overhearing agents must adapt their overhearing strategies to match the conversation system. This section describes the conversation systems expected of hierarchical organizations [Dewan *et al.*, 1997; Friebel and Raith, 2004], and proposes a number of general overhearing strategies for such organizations.

3.1 Modeling Conversation Systems

We define a conversation system of hierarchical organizations as a tuple (L, A, P, Λ, I, C). Some of these parameters have already been defined in [Gutnik and Kaminka, 2004], while others extend the previously proposed model. All of these are defined below. A common term we use throughout the discussion is that of the *value* of monitoring by overhearing. This value reflects the importance of the information in question to the task or organization being monitored.

Hierarchy Levels **(L)**. The notion of hierarchy levels is an extension of the previous model. It is used to determine the relative value of various organizational roles. Thus, one agent is considered to be more important (in terms of conversations it commits) than another agent if and only if its hierarchy level is higher than the level of the other agent. For each hierarchy level $l \in L$, we define a *value range* associated with it, i.e $\nu_{range} = [\nu_{min}(l), \nu_{max}(l)]$. It is used to define relation between two hierarchy levels. Thus, we will say that one hierarchy level is higher than another hierarchy level if and only if its minimum overhearing value is greater than the maximum value of the other hierarchy level, i.e. $\forall l_1, l_2 \in L,\ l_1 > l_2 \Leftrightarrow \nu_{min}(l_1) > \nu_{max}(l_2)$.

Agents **(A)**. A is the set of communicating agents in the monitored organization, each associated with a hierarchy level. The distribution of agents among hierarchy levels determines the type of hierarchical structure in organization. For instance, in pyramidal-hierarchies, discussed in this paper, the number of agents in higher hierarchal levels is always smaller than in the lower ones.

Conversation Protocols **(P)**. P indicates the set of conversation protocols used in a conversation system (see [Gutnik and Kaminka, 2004] for a detailed discussion). Intuitively, this is the set of conversation types that can occur, e.g., queries, brokering, informing, etc. Each protocol $p \in P$ has a set of roles, denoted $R(p)$. For a given conversation protocol, each role has value $\nu(r)$.

Conversation Topics **(Λ)**. Λ denotes the set of conversation topics. Each topic has a relative value indicated as $\nu(\lambda),\ \forall \lambda \in \Lambda$. This value associates each conversation topic to a corresponding hierarchy level, i.e. $\forall \lambda \in \Lambda\ \exists l \in L$ such that $\nu_{min}(l) \le \nu(\lambda) \le \nu_{max}(l)$.

Intervals **(I)**. An interval is a time period within the lifetime of a multi-agent system. Thus, we define I as follows: $I = \{[t_1, t_2] |\ t_1, t_2\ time\ stamps,\ t_1 \ge 0,\ t_2 \le lifetime,\ t_1 \le t_2\}$.

Conversations **(C)**. We define a conversation as a group of agents $g \in 2^A$ implementing a conversation protocol $p \in P$ on a conversation topic $\lambda \in \Lambda$ within a time interval $i \in I$. Thus, the C set can be formulated as $C \subseteq \{(p, g, \lambda, i) | p \in P, g \in 2^A, \lambda \in \Lambda, i \in I\}$. Using this definition, we can formulate the value of conversation for a certain communicating agent as $\nu(c, a) = \nu(\lambda) \oplus \nu(r)$ where $c = (p, g, \lambda, i),\ a \in g$ and $r \in R(p)$. Meaning that the value of conversation c for agent a (participating in it) is a function of conversation topic λ and role r (within conversation protocol p) that agent a implements. The information value of conversations distinguishes between the more important conversations and the less important ones.

Based on concepts adopted from organizational theory [Dewan *et al.*, 1997; Friebel and Raith, 2004], we can formulate the characteristics of conversations systems in hierarchical organizations. A set of conversations (C), generated in such organizations, has the following characteristics:

- **Conversations Distribution**. Conversations distribution depends on the distribution of agents among various hierarchy levels. For instance, in pyramidal-hierarchical organizations, lower levels are the "working" levels. Thus, most conversations are held between agents in lower hierarchical levels.
- **Conversation Topics**. Agents communicate on topics within their responsibility scope. Thus, agents mainly communicate on conversation topics associated with their hierarchy level or topics relatively close to it. As a result, agents of higher hierarchy levels carry out conversations on more valuable topics.
- **Conversation Groups**. Agents communicate mostly with their peers, subordinates and their close superiors. Thus, most communications are held between agents of the same hierarchy levels or between agents in relatively close hierarchy levels.
- **Conversation Roles**. Mostly, agents of higher hierarchical levels implement higher-value roles in conversation protocols.

In the experiments, generating conversation systems, we accomplish these characteristics using a set of probability functions (see discussion on experimental settings).

3.2 Overhearing Strategies.

A single overhearing agent, acting in a cooperative environment, assumes some knowledge of the monitored organization. An overhearing agent may know what agents generally communicate, which protocols are being used, which topics are discussed, etc. On the other hand, some information necessarily remains unknown. For instance, it does not necessarily know the complete list of conversations being held in organization at any given time.

We begin by describing how a single overhearing agent monitors a single target communicating agent. We assume that in such settings, the overhearing agent overhears all the conversations simultaneously carried out by the target. Of course, only conversations within the overhearing time interval, the time period in which the communicating agent is targeted, are being overheard.

The overhearing agent performs conversation recognition [Gutnik and Kaminka, 2004] for each conversation. Since conversation recognition takes time (to track and match the communicative acts being exchanged), the overhearing agent initially does not know the participants, protocol and topic associated with an overheard conversation. The overhearing agent starts overhearing assuming that the conversation protocol and topic can be any of the $p \in P$ and $\lambda \in \Lambda$ respectively. Gradually, the overhearer is able to disqualify inappropriate protocols and topics until it determines the correct protocol and topic. This information, at its different stages, can be used to determine whether to continue to overhear the current agent or to find another target.

Since a single overhearing agent can only hear a small subset of conversations in a conversation system, multiple overhearing agents can be deployed to maximize coverage of the overheard conversations. However, available overhearing resources (overhearing agents), are limited. Thus, overhearing targets should be carefully chosen in order to increase the total monitoring information produced by the overhearing group.

The systematic targeting of communicating agents by an overhearing group is called *overhearing strategy*. Various strategies can be proposed: centralized vs. distributed, assuming full vs. limited knowledge of the conversation system, given various levels of collaboration between overhearing agents, etc.

We focus on centralized overhearing strategies, leaving distributed strategies for future work. We focus on strategies with full information disclosure, where a centralized overhearing strategy has knowledge of key conversation system parameters (e.g. agents' hierarchy levels, agents active at time t, etc.). Using this information, it selects targets for the overhearing agents in the group.

4 Experiments

This section presents an empirical analysis of selective overhearing in pyramidal-hierarchical organizations. Each overhearing strategy may choose to overhear different target agents, and thus overhears different conversations. Consequently, some strategies may perform well while others perform poorly. Furthermore, the same overhearing strategy may vary in its performance, in principle, under different configurations of conversation systems and overhearing resource constraints. The experiments we report on seek to determine the scope of the strategies and the factors that influence their performance.

The overhearing strategies are evaluated in three steps. First, the optimal overhearing value, also referred as *optimum*, is calculated. Optimum is the value of most efficient overhearing possible, i.e. at each time unit t overhearing the k-best agents. Then, we calculate the evaluated strategy's *overhearing value*, which is the accumulative value of all overheard conversations using the specific overhearing strategy. Finally, the overhearing strategy is evaluated as a percentage of optimum.

4.1 Experimental Settings

The experimental settings have been defined to simulate communications in pyramidal-hierarchical organizations. The number of communicating agents, i.e. $|A|$, was set to 50 simulating relatively small organizations. In these simulated settings, we examined organizations with various number of hierarchy levels. The value range for each hierarchy level was calculated as a relative portion of [1,100], which was divided equally between the levels. The number of topics, i.e. $|\Lambda|$, has been set to 80. This value reflects our intuition that each agent has at least one conversation topic under its direct responsibility. The additional topics are generally common to all communicating agents. Each topic has been randomly given a value between 1 and 100 associating it with a hierarchy level, as described above.

The number of protocols was defined as 25 simulating a diversity of interactions that are possible in organization. The duration of each protocol has been randomly set to a value within {5,10,15,20,25}. For each protocol, two roles have been defined. Their values were randomly set to one of the following combinations: {50,50}, {67,33}, {75,25} and {99,1}. In this manner, we simulate differences in the importance of roles within the conversation. Finally, the conversation value is calculated using an accumulative function, i.e. $\nu(c, a) = \nu(\lambda) + \nu(r)$. Thus, conversation values range from 2 to 199.

To simulate a pyramidal organization, agents were distributed among different hierarchy levels according to a Zipf-like hyperbolic distribution. The probability of an agent to be associated with hierarchy level l, $1 \leq l \leq |L|$ was set to $1/l$ (normalized). Accordingly, the number of agents assigned to each hierarchy level becomes smaller as the hierarchy levels get higher.

In the experiments below, we generated conversation systems and simulated their dynamic execution, in a manner consistent with the characteristics of hierarchical organizations, described earlier. At the beginning of each simulation run, $|C_t|$–the number of conversations at time t–new conversations are generated using the procedure below. Then, each time a simulated conversation ends, a new conversation is generated. Thus, a constant *level of conversation activity* is maintained throughout the $lifetime$ of the conversation system (fixed at 1000).

Generating conversations.

The procedure for generating a single conversation at time t followed the steps below. We assume that each simulated conversation involves two communicating agents. However, this procedure can easily be extended to support larger conversation groups. First, we choose a level l_1 according to the distribution above. We then arbitrarily select an agent a_1, associated with l_1, to initiate the conversation.

Next, a conversation topic λ is chosen using the conditional probability $Pr(\lambda|l_1)$, calculated according to Bayes' rule as $[Pr(l_1|\lambda) \cdot Pr(\lambda)]/Pr(l_1)$. $Pr(l_1)$ is known from the hyperbolic distribution above. $Pr(\lambda)$ is assumed to be taken from the uniform distribution over Λ.

The calculation of $Pr(l_1|\lambda)$ requires some explanation. We remind the reader that a topic λ has a value v in the range [1,100], which determines its associated hierarchy level. We define a normal distribution with mean $\mu = v$ and standard deviation $\sigma = 0.5$. The value $Pr(l_1|\lambda)$ is given by this distribution. Intuitively, this translates into ensuring that agents usually carry out conversations on topics associated with their hierarchy level or relatively close to it.

The next step is to determine the level of the other agent. Here the process is reversed. We sample the topic's normal distribution to determine a new overhearing value, and its associated level l_2. We again arbitrarily select an agent a_2 associated with this level. Thus, $a_1 \text{and } a_2$ are likely to be associated with the same hierarchy level, or close.

Finally, a conversation protocol p is randomly chosen from the uniform distribution over P. We assign roles to the two agents from $R(p)$. To reflect the intuition that agents of higher hierarchies take the more important roles, we constrain the assignments such that 80% of assignments give the agent of higher hierarchy, the role of higher value.

We now compare several overhearing strategies using their evaluation values (as a percentage of optimum) in different configurations of activity levels, number of hierarchies, overhearing agents and importance of the hierarchy (measured as the ratio between the average value of conversations in the highest and the lowest hierarchy levels). Each evaluation is performed based on an average of 50 independent experiments with the same parameters. Thus, in the figures below, each data point corresponds to 50 trials.

4.2 Results

Static vs. Active Overhearing Strategies.

Our initial hypothesis has been that the most successful overhearing in pyramidal-hierarchical organizations (under the restriction of selectivity) would be achieved by overhearing conversations of the most important agents. The main intuition behind this hypothesis is that most important agents carry out the most valuable conversations. We refer to this type of strategy as *Largest Value*.

Several largest-value strategies are possible. In our first such overhearing strategy, called *MostImportantStatic*, k overhearing agents were set to overhear the k most important agents (in terms of their hierarchy level). So as to never miss a conversation carried out by these agents, the strategy committed to monitoring them regardless of whether they are currently communicating or not.

To evaluate this strategy, we define baseline overhearing strategy, called *RandomStatic*. Here, k overhearers were set to target k random agents chosen at the beginning of the experiment. Just as *MostImportantStatic* does not switch targets, neither does *RandomStatic*.

A potential drawback of these strategies is that their overhearing targets are determined statically. In cases where the overheard agent is idle, overhearing it has zero value. We thus contrast these static strategies with *active* strategies, in which the selection of targets is made out of those agents that are communicating at the moment of selection (though the agent may not know at what stage in the conversation they may be). The *RandomActive* chooses k target agents, similarly to *RandomStatic*. However, each time a target is idle, an alternative target is randomly chosen. The *MostImportantActive* strategy improves on *MostImportantStatic* by choosing the k most important agents from those that are currently active.

Figures 1-a,b compare these strategies. The values on the X-axis show the activity levels of the examined conversation systems, i.e. the ratio between the number of conversations at time t ($|C_t|$) and the number of communicating agents ($|A|$) (note that each agent may engage in more than one conversation in parallel). The Y-axis measures performance as percentage of the optimum. The overhearing coverage, defined as the ratio between the number of overhearers and the number of communicating agents—$k/|A|$, was set to 30% and the number of hierarchy levels was set to 7. A comparison between Figure 1-a and 1-b shows that the two *Largest Value* strategies outperform the random strategies in most activity levels.

However, more importantly, there is a qualitative difference in the behavior of the active and static strategies. In low activity levels, the likelihood of a given agent being idle is relatively high. In such settings, active strategies outperform static

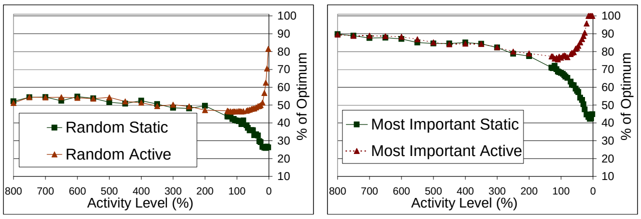

(a) Random (b) MostImportant

Figure 1: Static and Active Strategies

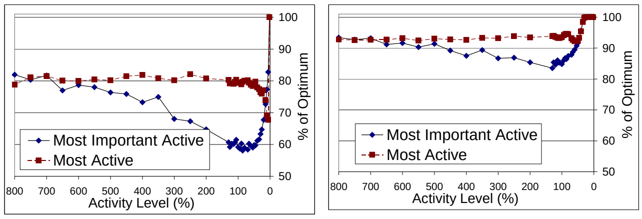

(a) Coverage: 5% (b) Coverage: 50%

Figure 2: Effect of Coverage

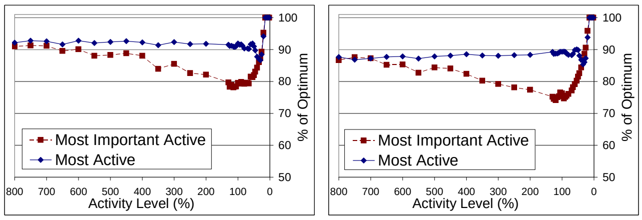

(a) 3 Levels (b) 11 Levels

Figure 3: Effect of Hierarchy Levels

strategies. However, as the activity level grows, the probability of an agent to be idle reduces. Thus, static overhearing strategies monotonically rise as the activity level grows until the probability of an agent to be idle is close (or equal) to 0. Indeed, in high activity settings, the difference between static and active policies is insignificant.

However, active strategies may not always be preferable, since they require qualitatively different knowledge about the monitored organization. Active strategies rely on the ability to detect agents that are conversing at any given time, unlike static strategies. Thus, a trade-off exists between the need to improve the results of overhearing, and the additional costs that may be required in detecting activity of potential targets.

Value or Volume.

As moving from static to active strategies increased the overall volume of overheard conversations (and thus the total derived value), a second overhearing approach—*Largest Volume*—suggests itself. We implemented a version of it, called *MostActive*, which targets the k most active agents, i.e. the k agents that are carrying out the highest number of conversation at time t. Since the overhearing agent overhears all conversations committed by its target, the idea is that this strategy will be more productive due to the greater quantity of overheard conversations.

Figures 2-a,b show the performance of the *MostActive* and *MostImportant* strategies, under conditions of different overhearing coverage—the ratio of the number of overhearing agents to the number of potential targets. It shows that overhearing volume can in fact be a successful strategy, using less knowledge about the monitored organization. The *MostActive* strategy does not require knowledge of the organizational role of the targets. This result is surprising given that in pyramidal-hierarchical organizations, most conversations are held between agents of lower hierarchy levels. Thus, in fact, *MostActive* targets the less important agents.

Indeed, Figures 2-a,b examine selective overhearing directly. They show both strategies become more efficient with higher overhearing coverage. Clearly, this conclusion is to some extent straightforward. However, an additional, less-trivial conclusion is that relative performance of these strategies does not change with selectivity. While increased coverage (reduced selectivity) increases the performance of both strategies, the *MostActive* strategy remains on top. It can be seen that the parabolic curve of *MostImportantActive* graph becomes less pronounced. In large overhearing groups, this effect can be explained by a significant overlap in overhearing targets for both strategies.

The Height of the Hierarchy.

We now turn to analyzing the effect of hierarchies on the performance of the proposed strategies. Figures 3-a,b show the strategies' behavior in organizations with 3 and 11 hierarchy levels, respectively (these are just a subset of results). The figures show that no significant change occurs in the performance of both strategies. Instead, only a slight performance decrease occurs when the number of hierarchy levels is larger.

This lack of change is caused by the two strategies essentially marking two extremes in the space of strategies in overhearing hierarchical organizations; they tend to prefer the top and bottom levels. The *MostImportantActive* strategy tends to always prefer agents in the top level. The *MostActive* strategy tends to prefer the bottom level (where there is more activity). Thus the middle levels in the organizations tend to be ignored by these strategies, regardless of the number of such middle levels.

The Importance of Importance.

It would seem that the relative performance of the two strategies, is qualitatively unaffected by selectivity level, nor by the height of the hierarchy (measured in number of levels). Yet hierarchical organizations are not characterized solely by their height. Rather, it is the difference in the importance of the different levels that is significant.

In a final set of experiments we changed the *importance ratio* between the low-value and the high-value conversations, i.e. the ratio between the average values of conversations in the bottom and top hierarchy levels, respectively. In the previous experiments, the value of conversations ranged from 2 to 199. On average, conversations committed by agents of lowest hierarchy level were valued close to 50, while conversations of highest-level agents were valued around 150 (ratio of 1:3). In these experiments, we examine the two strategies with additional ratios.

Figures 4-a,b,c show the performance of *MostActive* and *MostImportantActive* for importance ratios of 1:3, 1:5 and 1:8 ratios (where overhearing coverage is set to 20% and the number of hierarchies is 7). It can clearly be seen that as the ratio of conversations value increases, the *MostImportantActive*

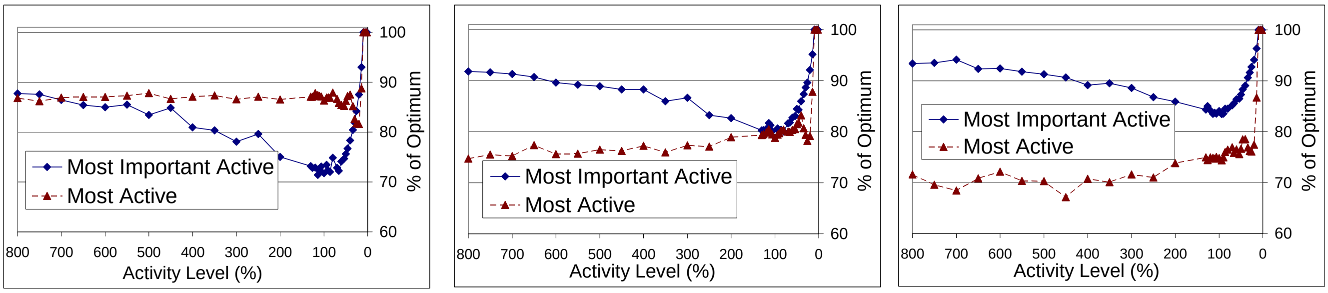

(a) Conversations Value Ratio 1:3 (b) Conversations Value Ratio 1:5 (c) Conversations Value Ratio 1:8

Figure 4: Overhearing Strategies Comparison with Respect to Conversations Value.

strategy improves (in all activity levels), while the *MostActive* strategy deteriorates. At some point (Figure 4-b), the two strategies shift relative places, and the *MostImportantActive* strategy dominates.

Thus, in case the difference between high-level and low-level conversation values is significant, it is more efficient to target highly important agents than to overhear low-level, highly-communicative ones.

5 Conclusions & Future Work

Lately, overhearing has gained interest as an approach for monitoring multi-agent systems. Previous investigations proposed an extensive set of techniques and practices using overhearing. However, the problem of selective overhearing, under the restriction of limited overhearing resources, has not been addressed so far.

In this paper, we present an empirical study of selective overhearing for hierarchically-structured organizations. Our work provides a model addressing both the characteristics of conversations in such organizations and the overhearing strategies appropriate for these settings.

Based on this model, we performed an extensive set of experiments simulating conversations in hierarchical organizations and examined some of several overhearing strategies for them. Based on these experiments, several important conclusions can be made:

- **Active Strategies**. Knowledge of which agents are active, facilitating active strategies, is crucial in an organizations with low conversation activity. However, as activity rises, the advantage of active strategies disappears.
- **Volume versus Value**. Strategies based on overhearing the most active agents tend to target agents at lowest organizational levels. Strategies based on the monitoring value of targets tend to target agents at highest levels. The key to deciding which strategy to use lies in the importance ratio, measuring the ratio in average value of conversations in the top and bottom levels. As this ratio increases, largest-value strategies do better, while largest-volume strategies do worse.
- **Volume versus Value Trade-off is Surprisingly Robust**. The height of an organization makes little difference, if any, on the monitoring information produced by a strategy, and does not change the Volume versus Value trade-off. Reduction in the level of selectivity (by increasing the number of overhearing agents) improves overhearing, but without changing the trade-off, either.

Currently, only centralized overhearing strategies have been concerned. In the future, we would like to examine the behavior of distributed and other overhearing strategies in these settings.

References

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

[Busetta *et al.*, 2002] P. Busetta, A. Dona, and M. Nori. Channelled multicast for group communications. In *Proceedings of AAMAS-02*, 2002.

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

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

[Gutnik and Kaminka, 2004] G. Gutnik and G.A. Kaminka. Towards a formal approach to overhearing: Algorithms for conversation identification. In *Proceedings of AAMAS-04*, 2004.

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

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

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

[Platon *et al.*, 2004] E. Platon, N. Sabouret, and S. Honiden. T-compound: An agent-specific design pattern and its environment. In *Proceedings of 3rd international workshop on Agent Oriented Methodologies at OOPSLA 2004*, pages 63–74, 2004.

[Rossi and Busetta, 2004] S. Rossi and P. Busetta. Towards monitoring of group interactions and social roles via overhearing. In *Proceedings of CIA-04*, pages 47–61, Erfurt, Germany, 2004.