

A Case Study of Organizational Effects in a Distributed Sensor Network *

Bryan Horling[†], Roger Mailler and Victor Lesser
Multi-Agent Systems Lab
University of Massachusetts, Amherst, MA
{bhorling,mailler,lesser}@cs.umass.edu

Abstract

We describe how a system employing different types of organizational techniques addresses the challenges posed by a large-scale distributed sensor network environment. The high-level multi-agent architecture of real-world system is given in detail, and empirical and analytic results are provided showing the various effects that organizational characteristics have on the system's performance. We show how partitioning of the environment can lead to better locality and more constrained communication, as well as disproportionate load on individuals or increased load on the population as a whole. The presence of such tradeoffs motivates the need for a better understanding of organizational effects.

1. Introduction

Distributed vehicle monitoring as an example application of distributed situation assessment and more generally distributed resource allocation has been studied in the MAS community since its infancy [6, 3]. This environment is particularly interesting when investigating issues of scale, because practical scenarios can be envisioned employing distributed sensor networks that are arbitrarily large both in number and geographic size, making purely centralized control inefficient. Each network member would have some type of data producing or interpretation capabilities, resulting in a potentially overwhelming amount of information requiring analysis. Shared resources and potentially conflicting goals add further complications. These challenges make it an ideal candidate for multi-agent techniques.

We propose using organizational structures to address these problems, which can appear in different forms in a variety of domains. This belief is based on our experiences working with a large-scale, realistic distributed sensor network over the past four years, both in detailed simulations and on real hardware [2]. Rather than employing a single organizational scheme, we have found that exploiting the strengths of a collection of heterogeneous organizational styles can be quite effective. By varying just one aspect of such an organization, we will show that the performance of the system can be greatly influenced by the organization's design parameters. We will present examples of such effects, and the methods used to discover and analyze them.

The goal of a distributed sensor network is most generally to employ a population of sensors to obtain information about an environment. In this paper, we will focus on using such a network to track one or more moving targets, although they are also commonly used to monitor weather conditions, traffic patterns and computer networks. In this work, we use detailed sensor models based on a three-head, MTI Doppler radar system [2]. No individual sensor is capable of solving the goal by itself, or else there would be little need for coordination. Instead, the sensors, each of which is under the control of an agent, must collaborate in some way to achieve their common goal. In our target tracking example, the sensors' measurements consist of only simple amplitude and frequency values, so no one sensor has the ability to precisely determine the location of a target by itself. The sensors must therefore be organized and coordinated in a manner that permits their measurements to be used for triangulation, and geographically distinct groups of such coordinated sensors used to produce a continuous track. More measurements, and particularly more measurements that are taken in groups at approximately the same time, will lead to better triangulation and a higher resolution track. Additional hurdles include a lack of reliable communication, the need to scale to hundreds or thousands of sensor platforms, and the ability to operate within a real time, uncertain environment. This environment is covered in detail in [2].

The notion of "organizational design" is used in many

* Effort sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreements F30602-99-2-0525 and DOD DABT63-99-1-0004, and by the National Science Foundation under Grant No. IIS-9812755. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency, Air Force Research Laboratory or the U.S. Government.

[†] Primary author is a student.

different fields, and generally refers to how entities in a society act and relate with one another. This is true of multi-agent systems, where the organizational design of a system can include a description of what types of agents exist in the environment, what roles they take on, and how they interact with one another. The objectives of a particular design will depend on the desired solution characteristics, so for different problems one might specify organizations which aim towards scalability, reliability, speed, or efficiency, among other things. To date, relatively little work has been done in the multi-agent community analyzing the characteristics and tradeoffs of different organizational types. We will provide quantitative results of our design to address this.

The organizational design used in this solution primarily attempts to address the scalability problem, by exploiting locality of reference and organizational constraints to impose limits on how far classes of both control and data messages propagate. The environment’s most limiting resource is the shared wireless communication medium, and we will therefore use this resource to describe the effects of the organization. Our design uses environmental partitioning to create localized regions of interaction. Within these partitions, agents take on particular and different roles which dictate their individual behaviors. The number of sensors in these partitions effects how efficient the system is, as large regions may create unwelcome disparities in load, and small regions cause a more global increase in overhead. In Sections 3 and 4 we will show quantitative evidence of these effects and the tradeoffs that exist between them, and in Section 5 we will show how these effects exist as the environment scales. Before describing our analysis, we first give a more thorough overview of the organization’s structure.

2. Organizational Overview

The environment is first divided by the agents into a series of partitions or sectors, each a non-overlapping, identically sized, rectangular portion of the available area, shown in Figure 1A. The purpose of this division is to limit the interactions needed between sensors, an important element of our attempt to make the solution scalable. In this work the partitions have been defined a priori; we have also explored the dynamic formation of partitions in [5].

Agents may work concurrently on one or more of several high level goals: managing a sector, tracking different targets and producing sensor data. The organizational leader of each sector is a single sector manager, which effectively acts as a hub within a nearly-decomposable hierarchical organization, by serving as an intermediary for much of the local activity. For example, they will generate and distribute plans needed to scan for new targets, store and provide local sensor information as part of a directory service, and assign track managers. They also concentrate nonlocal information, such

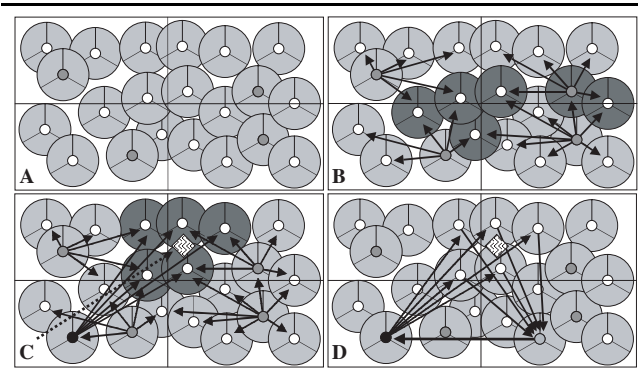


Figure 1. High-level architecture. A: sectorization of the environment, B: distribution of the scan schedule, C: negotiation over tracking measurements, D: tracking data fusion.

as target disambiguation and sensor status data, facilitating the transfer of that knowledge to interested parties. Track managers obtain their local information from their originating sector manager, but can also interact directly with other sector and track managers. They therefore do not follow a fixed chain of command or operate solely within one sector as one might see in a fully-decomposable organization.

To see how the organization works in practice, consider the scenario starting in Figure 1A. Sector managers are represented with shaded inner circles. Agents start by recognizing its manager, and sending it a description of its capabilities (position, range, etc.). The manager then uses this information to generate scanning schedule for detecting new targets in its sector, which it disseminates in Figure 1B. The manager does not strictly assign these tasks - the agents have autonomy to decide locally what action gets performed when. This is important because sensors can potentially scan in multiple sectors; thus there is the possibility that an agent may receive multiple, conflicting requests for commitments from different sector managers. The agent’s autonomy and associated local controller permit the agent to be responsible for detecting and resolving these conflicts as it sees fit.

Once the scan is in progress, individual sensors report any positive detections to the sector manager which assigned them the scanning task. Internally, the sector manager maintains a list of track managers currently in its region, and location estimates for their targets. If the positive detection does not match any of these targets, the manager selects an agent in its sector to be the track manager for that target. Not all potential track managers are equally qualified, and an unformed choice can lead to very poor tracking behavior if the agent is overloaded or shares communication bandwidth with garrulous agents. Therefore, in making this selection, the manager considers each of its agents’ estimated load, communication channel assignment, geographic loca-

tion and history. As we have seen previously, this notion of limited communication is an important motivating factor and recurring theme in this architecture which contributes to the organizational structure, role selection, protocol design and the frequency and verbosity of communication actions.

The assigned track manager (shown in Figure 1C with a blackened inner circle) is responsible for tracking the given target. To do this, it first discovers sensors capable of detecting the target, and then negotiates with members of that group to gather the necessary data. Discovery is done using the directory service provided by the sector managers. The track manager then determines when the data should be collected, and negotiates with the agents it selects (see Figure 1C). As with scanning, conflicts can arise between the new task and previously existing commitments, which the agent must resolve locally or elevate to the conflicting managers. More details about this process can be found in [4].

The data gathered from individual sensors is collected by an agent responsible for fusing the data and extending the computed track (see Figure 1D). The different measurements are used in a triangulation process, where amplitude and frequency values can place the target's location and heading relative to their source sensor, and several of these relative values can be combined to triangulate an absolute position. Because this is a relatively lightweight process, our organization assigns this fusion task to the track manager itself. If the data values returned are of high enough quality, and the agent determines those measurements were taken from the correct target, then they are used to triangulate what the position of the target was at that time. This data point is then added to the track, which itself is used as to predict where the target is likely to be in the future.

At this point the track manager must again decide which agents are needed and where they should scan. Under most situations, the process above is simply repeated. However, if the target has moved far from where the track manager is, the track managing task may be *migrated* to a new agent in a different sector. This is done to avoid the penalty associated with long-distance communication, which may cause unwanted latency or unreliability transferring information. This technique is covered in more detail in section 4.2.

3. Organizational Types

Below we will describe two of the organizational constructs used in this system, geographic coalitions and functional differentiation. The system also uses structures with characteristics similar to peer-to-peer and hierarchical organizations. An integral part of each of these is the notion of locality. Information propagates and is made available to only the agents which have need of it. In some cases, such as with the environmental sectorization, artificial boundaries are created to encourage locality at the expense of time or




Figure 2. Affect of sector size on messaging.

flexibility. In other cases, as with target tracking, information locality is exhibited naturally through the domain.

3.1. Geographic Coalitions

The partitioning described in Section 2 forms an organization based on the geographic location of sensors. Because much of the information being communicated is contained within sectors, the size and shape of the sector has a tangible effect on some aspects of the system's performance. If the sector is too large, and contains many sensors, then the communication channel used by the sector manager may become saturated, affecting both the manager and any other local nodes which use the same channel. If the sector is too small, then track managers may spend excessive effort sending information to different sector managers as its target moves through the environment. Deciding on the correct sector size is analogous to finding the correct membership of a coalition, although the reward and cost functions are difficult to completely specify and domain uncertainty precludes an allocation that is universally optimal. We hypothesized that a reasonable sector would contain from 6 to 10 sensors, although the physical dimensions of such a sector depend on the density of the sensors, and in different environments one would need to take into account sensor range, communication medium characteristics and maximum target speed. In the following sections, we will show results exhibiting these characteristics, and in section 5 we show how this evidence supports our initial hypothesis.

In these experiments and those that follow in this section, a group of 36 sensors were organized into between 1 and 36 equal-sized sectors with 4 mobile targets. Although the underlying technology can work with arbitrary configurations, the sensors are arranged in a grid pattern and the targets' location and movement spread evenly through the environment to normalize results and simplify analysis. The results were observed over 10 runs per configuration in a simulation environment which closely models the performance of the physical MTI sensors (the same agent code was used for both the simulation and actual hardware tests). Note that these experiments were performed under idealized conditions. Reducing the rate at which simulation time passes pre-




Figure 3. Messaging disparity vs. sector size.

vents bottleneck and overloading conditions from occurring, so behavioral and organizational phenomena associated with the partition changes can be better isolated. Under bounded conditions, excessive message or activity loads would cause performance degradation across many areas of the system. It is our intent to deduce what the issues behind these tradeoffs are, so an informed decision can be made when these limiting effects are present. Conversely, with the same information one could predict the improvements that might be achieved if those limits were relaxed (for instance, if a higher capacity communication mechanism were used).

Figure 2 shows that as the number of agents per sector increases, and there are correspondingly fewer sectors overall, the amount of communication traffic decreases. Because each sector requires a certain amount of control messages, the total number of messages is reduced as the number of sectors decreases. A more detailed view of the effects this change has on messaging will be shown later in Figure 4.

Partitioning can also reduce reactivity, because an extra step may be required to fetch information. A track manager must perform queries to obtain sector information as its target moves to new sectors, so smaller, more numerous sectors will result in delays caused by the additional queries. This delay will be revisited in Section 6.

3.2. Functional Differentiation

The varied assignment of roles forms a different, functional organization [1] in the system. Agents specialize their functionality in order to restrict the type of interactions which must take place between agents. For example, to obtain information about available sensors, a track manager must only contact the relevant sector managers [7]. Concentrating the track management functionality into individual agents serves a similar role, by limiting the number of interactions necessary to resolve conflicts in sensor usage.

Interestingly, although this type of functional decomposition does reduce the total number of interactions an agent might need to make, it can also increase that number for particular individuals in the environment. For example, we have seen how the sector manager is responsible for disbursing information about the sensors in its sector, thus providing a

single point of contact for such data. However, by serving in this capacity, it makes itself a center of attention, which can adversely affect its overall performance.

Consider Figure 3, which shows how sector size affects the standard deviation in communication activity exhibited by individual agents. This metric captures how much agents in the population differ in their communication habits. If all agents are roughly the same they will have a low deviation, while a population that has a handful of outlier agents with significantly higher message traffic will have a high deviation. As the number of agents in each sector increases, this graph shows a marked increase in disparity, because a few agents are communicating much more than their peers. As the sector sizes scale, specialized agents become “hotspots” of activity. In a bounded environment this could lead to significant data loss as the communication channel becomes overloaded. In conjunction with Figure 2 which shows the average *total* communication, we see a tension between sector sizes: smaller sectors lead to increased overall message traffic, and larger sectors can imbalance load in the population. Since the environment and target spacing are uniform, the differences can be attributed to the roles those agents take on. The slight rise at the left end of the graph, where there is a single agent per sector, represents the coexistence of the sector and track manager roles at a single node.

4. Maintaining Organizations

Although we have seen how organizations can be effective, there are costs associated with creating and maintaining these structures which have the potential to significantly degrade their benefits. These costs differ from those described in the previous section in that they are more dependent on the dynamics of the environment and organization.

4.1. Measurement Collection

The most frequently updated organization in the environment is the manager-worker hierarchy formed between track managers and sensors, because the tracking sensors change as the target moves. This results in a class of organizational control messages dependent on sector size. For example, as the target moves into part of the environment the manager is not familiar with, the track manager must send a directory service query to the sector manager of that area to discover which sensors are available. Once those sensors are found, additional messages are needed to create and maintain data collection commitments with them. Finally, as the target is tracked, the relevant, nearby sector managers must be notified of the target’s estimated position.

Figure 4 provides a quantitative view of this overhead. As the number of agents in each sector increases, fewer directory service and tracking control messages are necessary,




Figure 4. Message types vs. sector size.




Figure 5. Effect of sector size on RMS error.

because there are a fewer sectors which must be interacted with as the target moves. In addition, the number of measurements increases as the sector size increases, which produces a lower root-mean-squared (RMS) error between the measured and actual track, as seen in Figure 5. This is also due to communication control overhead; the reduced time spent by the manager interacting with the additional sector managers allows more time to be spent collecting data. This is caused primarily by front-loading sensor discovery, and the reduced probability of track manager migration.

4.2. Track Manager Migration

The technique of migrating the tracking responsibility through the agent population as the target moves is another aspect of local information exploitation. It should be clear that, lacking the capacity for movement, the initial manager selected to track a target will gradually become less effective as the target moves away from it. Simple signal latency and attenuation conspire to make communication over distance less reliable. By migrating this task to follow the target, the organization is able to retain locality despite the fact that the sensors themselves are immobile. This results in a reduction in the average distance that messages must travel.

Figure 6 shows the effects track manager migration has on the average distance of communication. Because migration is triggered by sector boundaries, the tracking task will migrate less frequently when sectors are large simply because they covers more area. Thus, a lower average communication distance is observed when sectors are smaller.




Figure 6. Average communication distance.




Figure 7. Communication disparity with varied sector sizes and target densities.

5. Scalability Results

To explore the generality of these conclusions, we performed similar tests with varied target numbers and different sensor population sizes. The first set of tests kept the sensor population static, but introduced between 1 and 24 targets. The scenario was otherwise identical to previous tests. Figure 7 shows that our original communication disparity profile is maintained if the target density is varied, although the level of disparity is reduced as the number of targets increases. Intuitively, this is because more agents are doing more work, and thus the effect of distinguished overworked nodes is minimized. As in Figure 3, when there is only a single node per sector, overlapping roles exacerbates the problem. Similarly consistent results are shown in Figure 8. As one would expect, the baseline RMS error increases with the number of targets, since the bounded sensing capabilities result in fewer average measurements per target.

In the second set of additional tests, we varied the size of the sensor population from 9 to 81, while maintaining a similar sensor-to-target ratio. The communication deviations shown in Figure 9 are consistent with the earlier findings. The flattening and lowering of the profile as the number of sensors increase is expected, because the larger number of targets spreads work out among more agents. This decrease is somewhat misleading, however, as it hides the true burden that is imposed under these more extreme circum-



Figure 8. RMS error differences with varied sector sizes and target densities.




Figure 9. Communication disparity with varied sector and sensor population sizes.

stances. The deviation, although slight, can represent a significant load when the population is large. Figure 10 shows the actual communication burden incurred by different roles for a single-sector environment (i.e. all sensors are in the same partition). The sector manager’s burden increases at an undesirable rate, while the track manager and median (non-manager) roles remain relatively constant.

Although space precludes a more thorough examination, similarly consistent trends were observed in the RMS error, message totals, and average communication distance for other target and sensor population sizes.

These experiments suggest a tradeoff exists between the overall volume of message traffic and its distribution over the agent population. Message volume decreases when there are more agents per sector because fewer interactions are needed to obtain information, as shown in Figure 2. However, this shift can cause individual agents to incur a disproportionate communication burden, as shown in Figures 3 and 9. Figure 10 in particular shows the large-sector solution does not scale well. Figures 4, 5, and 6 show that maintenance of these organizations has a similar tradeoff, since larger sectors require a lower control overhead and better RMS error, while smaller sectors allow track migration to take advantage of information locality.




Figure 10. Average communication totals by role for a single-sector environment.




Figure 11. Finding the appropriate configuration from normalized results.

6. Analysis

Normalizing and overlapping the earlier results produces the graph in Figure 11. By searching for a common inflection point in these results, we can conclude that a sector size between 4 and 9 is most appropriate in general for this environment. This supports our hypothesis that sector sizes between 6 and 10 were “reasonable”. However, if more robust managers were available to handle the increased load, this graph also shows that better RMS performance can be obtained by using larger sector sizes. In general, the requirements imposed by goals and capabilities of the system and environment guide an appropriate selection.

Although these profiles give empirical evidence of the system’s performance, it is usually preferable to work with a more formal model at design time. We therefore wish to capture the system’s behaviors in an abstract, quantitative model that provides a good approximation of the real system. We will concentrate our analysis on the communication load by role or task over the lifetime t of the role or task. As before, we will assume that the sensors and targets are uniformly distributed throughout the environment, and that targets move with constant velocity. One could relax these assumption by estimating explicit interaction probabilities; although the calculations would be more complex, the spirit of

the analysis we present would remain the same. Similarly, one could determine worst-case peak performance (which could be localized) by assuming worst-case densities. The formulas presented below do not represent actual message totals, but are meant to reflect relative growth rates. As we will show in Figure 12, quantitative total estimates can be obtained through the addition of appropriate constants.

We start by looking at the measurement gathering (\mathcal{M}) role. Sensors producing measurements do so in response to track manager requests, which are in turn prompted by the belief that the target is within range of the sensor. Therefore the role's communication load is dependent on the likelihood that a target is within its range r .

$$\mathcal{M} = \sum_t \min\left(\frac{\pi r^2}{A} T m, m\right) \quad (1)$$

T is the number of targets in the environment, A is the environment's area, and m is the number of measurements requested per time unit. So, as the number of targets increase, or the environments area decreases, the number of measurements will approach tm . This model is an upper bound, however, as it does not take into account the track managers' specific behaviors. To better understand this, we will look at how many measurements are generated for an entire track.

Ignoring the effects of uncertain measurements or faulty data fusion (which are outside of the organization's scope), the RMS error of the tracking process is dependent on the number of measurements produced for the track over its lifetime (\mathcal{R}). In the absence of hindering factors, the track will ideally receive measurements at a uniform rate m from each of c sensors used (we assume c is sufficient for triangulation purposes). The actual rate of measurement is affected by the number of sensors that are used and any delays incurred by overhead tasks. In particular, the collection of sector directory information, and task migration when the target has grown too distant can reduce the total number of measurements are obtained. Competition for sensors by other targets can also reduce the measurement rate.

$$\hat{c} = \min\left(c, \frac{\pi(b+r)^2}{A} N\right) \quad (2)$$

$$l = \min\left(0, v\left(\frac{N}{\hat{c}T} - c\right)\right) \quad (3)$$

$$\mathcal{R} = \sum_t \hat{c} m \left(\min\left(1, \frac{N}{\hat{c}T}\right) \lambda^l \left(1 - \frac{v}{\sqrt{S}} \left(d - \frac{g}{2}\right)\right) \right) \quad (4)$$

Equation 2 defines \hat{c} , the number of sensors that will actually be used to track the target. It is bounded above by the desired quantity c , and below by the expected proportion of the total number of sensors N that are in range of the target. The radius b of the target bound will depend on a number of factors, including measurement uncertainty and the target's velocity. The first term of Equation 4 models the proportion of a potentially contended sensor's time usable




Figure 12. Comparison of predicted and actual results of \mathcal{R} for 2, 4 and 12 targets.

by the target. If we assume that each target using the common sensor will share it equally, then target density will be inversely proportional to the measurement rate obtained by an individual target. However, as sensors come under contention, an allocation strategy must be employed to resolve the conflict [4]. An additional reducing factor models the effect of this optimization process; l estimates the amount of conflict, while λ controls how much the conflict degrades performance. When the target moves into a new area, there will be a delay d before the appropriate directory information is received. The rate at which this happens depends on the velocity v of the target and the average distance across the sector. This of course depends on the probability of target turns and the shape of the sector itself; we model this with a very coarse estimate of the average chord length in the sector \sqrt{S} . A delay g is incurred by tracking task migrations occur when the target has moved two sectors away from that of the track manager. The net effect of these delays and the corresponding increase in measurements when sector sizes grow is supported by Figures 4 and 5. A comparison of the predicted \mathcal{R} obtained from Equation 4 and the observed load is shown in Figure 12, which was produced after finding appropriate constants for our system.

Returning to the estimated measurement gathering load, we can see that \mathcal{M} is more accurately represented by $\frac{\mathcal{R}T}{N}$, as \mathcal{R} models the managing behaviors absent in Equation 1.

The sector manager's load (\mathcal{S}) is dependent on both the size of the sector and the number of targets. As we have observed earlier larger sectors mean more sensors must be registered, as well as an increased probability that a target will be in the area. \mathcal{S} can be broken down into the one-time costs associated with sector creation, and the continuing costs derived from targets moving through the sector:

$$\mathcal{S} = \frac{S}{A} N + \sum_t \left(\frac{\hat{S}}{A} T u + \frac{v}{\sqrt{S}} \frac{S}{A} N \right) \quad (5)$$

u is the frequency at which target updates are supplied to the sector manager by the track manager, S is the size

of the sector's area while \hat{S} is the *effective* size of the sector's area. S and \hat{S} are differentiated by what they represent. S is a strict bounding area; membership in the sector is defined by containment within that area. \hat{S} is the area over which measurements can be taken by those sensors; if for example each sensor has a range of $r = 20$, then \hat{S} will be the area bounded by S unioned with a perimeter of width 20 surrounding S . The second term in the summation represents the directory queries it must respond to as targets enter its sector, which depends on the velocity of the target. \mathcal{S} therefore grows in proportion with \hat{S} , which is supported by the results observed in Figure 10.

The track manager's load (\mathcal{T}), depends in part on the measurement rate described above. It will receive $\mathcal{R}_{\mathcal{T}}$ measurements over the lifetime of the role, where $\mathcal{R}_{\mathcal{T}}$ is the portion of \mathcal{R} generated for that particular track manager. This will be the same as Equation 4, but without the migration term g and using a role-specific t . It will also need to receive updates of other targets' locations and periodically notify relevant sector managers of the target's location. Because we are looking at the communication load of a single track manager role over the time it exists at a particular node, the load is not affected when the tracking task is migrated to a new node.

$$B = \frac{\hat{S} + 2b(\hat{S}_x + \hat{S}_y) + \pi b^2}{\hat{S}} \quad (6)$$

$$\mathcal{T} = \mathcal{R}_{\mathcal{T}} + \sum_t \left(\frac{\hat{S}}{A} T u + B u \right) \quad (7)$$

The number of location updates u received by the target manager, represented by the second term of \mathcal{T} , is proportional to (but less than) the number received by the sector manager above. The number of target updates that must be sent to sector managers depends on how many sectors currently intersect with the target bound's area. The equation for B is the average such number for a uniform environment, derived from a geometric analysis of bound-sector interactions. \hat{S}_x and \hat{S}_y are the dimensions of the effective sector area, and $\hat{S} = \hat{S}_x \hat{S}_y$ because our sectors are rectangular, while b is the radius of the target bounds. Interestingly, because of the relationship of S and b in B , increasing the target bounds (from an increased target speed, for instance) has roughly the same effect as shrinking the sector bounds. Note that as the dimensions of sector increase, \hat{S} will increase quadratically, causing B to approach 1, while the number of target updates will increase. Because the rate of target location receipts is less than that sent out, the net effect is a reduction in communication load. This is supported by the reduction in tracking control seen in Figure 4.

7. Conclusions

The quantitative results we have presented are quite domain specific. They depend on the communication characteristics of the environment, the actions needed to achieve the scenario goals, and the behaviors exhibited by the agents. However, we feel that the types of issues raised by these experiments, such as information locality, specialization bottlenecks and organizational control overhead, are applicable to many different domains, particularly those which are communication intensive. For instance, our sector size results can be directly related to the estimated load incurred by a distributed collection of middle agents [7].

More generally, we feel that multi-agent organizations can have significant positive effects on performance. By specifying roles, authority relationships and working groups, the system can both reduce runtime combinatorics by restricting search as well as improve global coherence without requiring a global view. However, we have seen that these benefits come with costs and side effects, which must be well understood for the organization to be used successfully. In this paper, we varied just one organizational parameter, and observed the ramifications of this change across several distinct dimensions. With continued research in this area, the complete space of organizational types and their corresponding characteristics can be more fully understood and exploited.

References

- [1] Mark S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(1):70–80, January 1981.
- [2] V. Lesser, C. Ortiz, and M. Tambe, editors. *Distributed Sensor Networks: A Multiagent Perspective (Edited book)*, volume 9. Kluwer Academic Publishers, May 2003.
- [3] V.R. Lesser and L.D Erman. Distributed Interpretation: A Model and an Experiment. *IEEE Transactions on Computers Special Issue on Distributed Processing*, C-29(12):1144–1163, December 1980.
- [4] Roger Mailler, Victor Lesser, and Bryan Horling. Cooperative Negotiation for Soft Real-Time Distributed Resource Allocation. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, pages 576–583, Melbourne, July 2003. ACM Press.
- [5] Mark Sims, Claudia Goldman, and Victor Lesser. Self-Organization through Bottom-up Coalition Formation. In *Proceedings of Second International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2003)*, pages 867–874, Melbourne, AUS, July 2003. ACM Press.
- [6] Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104–1113, 1980.
- [7] Katia Sycara, K. Decker, and M. Williamson. Middle-agents for the internet. In *Proceedings of IJCAI-97*, January 1997.