

An Approach to Analyzing the Need for Meta-Level Communication *

Keith Decker and Victor Lesser
Department of Computer Science
University of Massachusetts
Amherst, MA 01003
DECKER@CS.UMASS.EDU

Abstract

This paper presents an analysis of static and dynamic organizational structures for naturally distributed, homogeneous, cooperative problem solving environments, exemplified by distributed sensor networks. We first show how the performance of any static organization can be statistically described, and then show under what conditions dynamic organizations do better and worse than static ones. Finally, we show how the variance in the agents' performance leads to uncertainty about whether a dynamic organization will perform better than a static one given only agent *a priori* expectations. In these cases, we show when meta-level communication about the actual state of problem solving will be useful to agents in constructing a dynamic organizational structure that outperforms a static one. Viewed in its entirety, this paper also presents a methodology for answering questions about the design of distributed problem solving systems by analysis and simulation of the characteristics of a complex environment rather than by relying on single-instance examples.

1 Introduction

Organizational theorists have long held that the organization of a set of agents cannot be analyzed separately from the agents' task environment, that there is no single best organization for all environments, and that different organizations are not equally effective in a given environment [Galbraith, 1977]. Most of these theorists view the uncertainties present in the environment as a key characteristic, though they differ in the mechanisms that link environmental uncertainty to effective organization. In particular, the *transaction cost economics* approach [Moe, 1984] focuses on the *relative efficiencies* of various organizations given an uncertain environment, while the modern *contingency theory* approach [Stinchcombe, 1990] focuses on the need for an organization to expand toward the earliest available *information that resolves uncertainties* in the current environment.

*This work was supported by DARPA contract N00014-92-J-1698, Office of Naval Research contract N00014-92-J-1450, and NSF contract CDA 8922572. The content of the information does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

In this paper we use both of these concepts to analyze potential organizational structures for a class of naturally distributed, homogeneous, cooperative problem solving environments where tasks arrive at multiple locations, exemplified by distributed sensor networks [Lesser and Corkill, 1983]. Previous approaches to analyzing organizations in distributed sensor networks have either not focused on the effectiveness of the organization [Davis and Smith, 1983], or have only analyzed organizational effectiveness in particular, single-instance examples [Durfee *et al.*, 1987]. Our approach is to model the task environment mathematically, using a formalism developed specifically to study distributed coordination and scheduling [Decker and Lesser, 1993b]. We then develop expressions for the *expected efficiencies* of static and dynamic organizational structures, in terms of the cost of communication and time to complete a given set of tasks. Finally, we validate these mathematical models by using simulations.

A dynamic organization is one in which the responsibilities of agents can be reassigned based on a developing view of the problem at hand. Due to the uncertainties explicitly represented in the task environment model, there may not be a clear performance tradeoff between static and dynamic organizational structures when agents use just their own local views to make a reorganization decision. Agents that have a dynamic organization have the option of meta-level communication—communicating about the current state of problem solving as opposed to communicating about solving the problem itself. In this way, *information that resolves uncertainties* about the current environment becomes available to the agents, allowing the agents to then create the most efficient organization for the situation.

Section 2 describes the task environment model, the assumptions behind it, and analyzes the uncertainties present. Section 3 describes static and dynamic organizational structures, and develops expressions for the expected performance of each organizational style. In Section 4 we then show how the variance in performance without communication can lead to the efficient use of meta-level communication to customize a dynamic organizational structure. Finally, we discuss how these results can be used by designers of distributed problem solvers, and how our methodology can be used by other researchers. Throughout each section, we will illustrate and confirm the analytical results experimentally, using as an example a simulated distributed sensor network similar to the Distributed Vehicle Monitoring Testbed (DVMT) [Lesser and Corkill, 1983].

2 Task Environment Model

Our task environment model of naturally distributed problems assumes that several independent groups of computational tasks arrive at multiple locations over a period of time called an *episode*. For example, in a distributed sensor network (DSN) episode the movements of several independent vehicles will be detected over a period of time by one or more distinct sensors, where each sensor is associated with an agent. The performance of agents in such an environment will be based on how long it takes them to process all the task groups necessary to interpret their sensed data, which will include the cost of communicating data, task results, and meta-level communication, if any. The organizational structure of the agents will imply which subsets of which task groups are available to which agents and at what cost. For example, if DSN agents have overlapping sensors, either agent can potentially work on data in the overlapping area (from its own sensor) without any extra communication costs. We make several simplifying assumptions: that the agents are homogeneous (have the same capabilities with respect to receiving data, communicating, and processing tasks), that the agents are cooperative (interested in maximizing the system performance over maximizing their individual performance), that the data for each episode is available simultaneously to all agents as specified by their initial organization, and that there are only structural (precedence) constraints within the subtasks of each task group.

Any single episode can be specified by listing the task groups, and what part of each task group was available to which agents, given the organizational structure. Our analysis will be based on the statistical properties of episodes in an environment, not any single instance of an episode. The properties of the episodes in a DSN environment are summarized by the tuple $\mathcal{D} = \langle A, \eta, r, o, \mathcal{T} \rangle$ where A specifies the number of agents, n the expected number of task groups, r and o specify the structural portion of the organization by the *range* of each agent and the *overlap* between agents, and \mathcal{T} specifies the homogeneous task group structure (Section 2.5 and Figure 4 describes how task group structures are specified).

Our analysis initially focuses on what *a priori* knowledge agents have about the distribution of task groups in an episode. First we will look at the distribution of the lowest-level sensor subtasks of a single task group among multiple agents (deriving the maximum expected number of subtasks), and then we will look at the distribution of task groups themselves. These results will then be used in subsequent sections to derive the total amount of work, and therefore expected termination performance, under various organizational structures and control schemes.

2.1 Task environment simulation

In the next sections and for the rest of the paper, we will test the model we are developing against simulated DSN problems. Each simulated DSN episode will take place on a grid where the concepts of length and size correspond directly to physical distances. For example, Figure 1 illustrates several simple organizations imposed on such a grid in our simulation.

In the simulation we assume that each vehicle is sensed at

¹in general there are usually more complex interrelationships between subtasks that affect scheduling decisions, such as *facilitation* [Decker and Lesser, 1993b].

discrete integer locations (as in the DVMT), randomly entering on one edge and leaving on any other edge. In between the vehicle travels along a *track* moving either horizontally, vertically, or diagonally each time unit using a simple DDA line-drawing algorithm (see Figure 5). In an 18 x 18 grid, the (empirical) average length of a track is 14 units—the actual length of any one track will range from 2 to 19 units and is not distributed normally. Given the organization (r , o , and A , and the geometry), we can calculate what locations are seen by the sensors of each agent. This information can then be used along with the locations traveled by each vehicle to determine what part of each task group is initially available to each agent. Section 2.5 will detail what the structure of each task group is for the DSN simulation.




Figure 1: Examples of DSN organizations on an 18 x 18 grid

2.2 Expected number of sensor subtasks

In order to analyze the performance of a particular organization, we will want to know what proportion of each task group each agent is likely to process. There will be some upper limit on this proportion (related to the agent's range r), and sometimes the agent will process less than this upper limit. Especially in static organizational structures where tasks are not exchanged, the termination of the system as a whole can be tied to the completion of all tasks at the most heavily loaded agent. Normally, we would use the *average* part of a task group to be seen, but since the focus of our analysis is the termination of problem solving, we need to examine the *expected maximum* portion of a task group to be seen. This section will develop an equation for the expected maximum workload at an agent by counting the expected number of low-level sensor subtasks (each individually associated with a sensed vehicle location) that the maximally loaded agent will have.

The amount of a single task group seen by an agent (which is the same as the number of sensor subtasks in the DSN example) can be viewed as a random variable S with a probability density function and corresponding cumulative distribution function. In the DSN environment, S is discrete, and its probability function (determined empirically) is heavily weighted toward r (the maximum). To simplify the analysis, instead of letting S correspond to the number of subtasks in a single task group seen by an agent, we have it equal 1 if the agent sees the maximum amount, and 0 otherwise. Now S has a Bernoulli (coin-tossing) distribution with parameter p corresponding to the chance of an agent seeing the maximum amount r of a task group. Let's assume we know that $N < n$ is the number of task groups at the maximally loaded agent, and that on average $a < A$ agents see a single task group. The number of times an agent sees the maximum out of N task groups (N coin flips) then has a binomial distribution $(N; p, p^*)$. We need to know, given that a agents each flip N coins, what the distribution is of the

maximum number of 'heads' any agent sees—this is called the binomial max order statistic, $g_{a,N,p}(s)$:

$$\begin{aligned} b_{N,p}(s) &= \binom{N}{s} p^s (1-p)^{N-s} & [\Pr[S_N = s]] \\ B_{N,p}(s) &= \sum_{x=0}^s b_{N,p}(x) & [\Pr[S_N \leq s]] \\ g_{a,N,p}(s) &= B_{N,p}(s)^a - B_{N,p}(s-1)^a & [\Pr[\hat{S}_N = s]] \end{aligned}$$

In the DSN example we have $p = 0.5$ (empirical) and the amount of a task group seen when an agent does not see the maximum amount averages $r/2$ (empirical). The expected heaviest load seen by any agent when a agents see N task groups with a probability p of seeing r and probability $1 - p$ of seeing $r/2$ is:

$$E[\hat{S}_N] = \sum_{s=0}^N g_{a,N,p}(s) \left(rs + \frac{r}{2}(N-s) \right) \quad (1)$$

Figure 2 demonstrates the correctness of Eq. 1 by showing the heaviest load actually observed in the simulation plotted against the expected value from Eq. 1, for 1000 runs with n tracks and all square DSN organizations [$2 \leq r \leq 10, 0 \leq o \leq r, 1 \leq \sqrt{A} \leq 10, 1 \leq N \leq 10$] ($R^2 > .99$).




Figure 2: Actual versus predicted heaviest load for various values of A , r , o , and n

2.3 Expected Number of Task Groups

Given the maximum number of task groups seen by an agent (N), we can calculate the expected heaviest agent load using Eqn. 1. But this begs the question of what is the maximum number of task groups an individual agent will see, given the actual (n) or expected number (\hat{n}) that the entire system will see. The solution is similar—each agent either sees or does not see each of the n task groups, another binomial process. Let N_i be the number of task groups sensed by agent i , with a binomial distribution of parameters n and q . If a is again the number of agents that see a single task group and A the total number of agents, then $q = a/A$, the probability that each agent will see a particular track. By the same derivation as in the last section, the max order statistic \hat{N}_n has the probability function $g_{A,n,a/A}(s)$, and expected value:

$$E[\hat{N}_n] = \sum_{s=0}^n s g_{A,n,a/A}(s) \quad (2)$$

Figure 3 shows the actual mean value of the maximum number of tracks seen by an agent over 1000 runs of the DSN simulation, versus the predicted value, for n tracks and square DSN organizations [$2 \leq r \leq 10, 1 \leq \sqrt{A} \leq 10, 1 \leq n \leq 10$] without any overlap ($R^2 > .99$)




Figure 3: Actual versus predicted maximum number of task groups (tracks) seen by any one agent for various r , A , and n (left); Actual versus predicted average number of agents seeing a single task group (track) for various r , o , and A (right).

2.4 Expected Number of Agents

The only remaining term we need to analyze before deriving an expression for system performance is a , the expected number of agents that will see a single task group. In general, a will depend on the total number of agents A and the organization (r and o). When there is only one agent, it will see every task group ($a = 1$). When the agents overlap completely, every agent sees every task group ($[o = r] \rightarrow [a = A]$). When the agents in a square environment do not overlap, $a = \sqrt{A}$. The relationship follows the ratio of the area solely covered by an agent plus the area of the overlapping section, to the total area covered alone:

$$a = A \left(\frac{r^2 + o^2}{2ro} \right) \quad (3)$$

Figure 3 shows a regression of the actual average value of a over 1000 runs versus the predicted value for all 630 DSN organizations [$2 \leq r \leq 10, 0 \leq o \leq r, 1 \leq \sqrt{A} \leq 10$] ($R^2 > .99$).

2.5 Work Involved in a Task Structure

Finally we turn to modeling the performance of the system as a whole, which is based on the structure of the tasks involved. We have developed a characterization of task environments that formally captures the range of features, processes, and especially interrelationships that occur during computationally intensive coordination and scheduling [Decker and Lesser, 1993b]. The model of environmental and task characteristics we propose has three levels: *objective*, *subjective*, and *generative*—the subjective level is not discussed here.

The *objective* level describes the essential structure of a particular problem-solving situation or instance over time. It focuses on how task interrelationships dynamically affect the *quality* and *duration* of each task. In this paper we will concentrate only on duration as a performance metric. The basic model is that *task groups* T occur in the environment at some frequency, and induce *tasks* to be executed by the agents under study. Task groups are independent of one another, but tasks within a single task group have interrelationships. An individual task that has no sub tasks is called a *method* M and is the smallest schedulable chunk of work. The quality and duration of an agents performance on an individual task is a function of the timing and choice of agent actions ('local effects*'), and possibly previous task executions ('non-local effects*'). The basic purpose of the objective model is to formally specify how the execution and timing of tasks affect quality and duration.

At the lowest level, each method (leaf task) M at time t can produce, if executed, some maximum quality $q(A_f, t)$ in some amount of time $d(M, t)$ (each method has an initial maximum quality $q_0(A_f)$ and duration $d_0(M)$). Any task execution that starts before the execution of M completes may potentially affect M 's execution through *non-local effects*. The effect is dependent on the relative timing of the two task executions, the quality of the task causing the effect, and whether information was transmitted between the two tasks.

This work considers a single non-local effect, *precedence*. If task A precedes task B , then the maximum quality $q(B, t) = 0$ until A is completed and the result is available, when the maximum quality will change to the initial maximum quality $q(B, 0 = q_0(B)$.

2.5.1 Execution Model

For this paper we use an extremely simple model of execution. Agents can perform three actions: method execution, communication, and information gathering. The control component of an agent determines the next action an agent will perform based on the agent's current set of beliefs [Cohen and Levesque, 1990; Shoham, 1991]. A method execution action, of method M , that is begun at time t will conclude at time $t + d(M, t)$. An information gathering action has duration $d_0(\tau)$ and updates the agents set of beliefs with any new information in the environment, for example, the arrival of data at the start of an episode, or communications from other agents. A communication action has duration $d_0(C)$ and, after a communication delay, makes information (such method execution results) available to other agents. The agent on the receiving side must perform an information gathering action before the communication can affect its local beliefs.

2.5.2 Simple Objective DSN Model

Recall that the summary of a DSN environment was the tuple $D = \langle A, \eta, r, o, T \rangle$; this will become our generative model, especially the parameter n (expected number of task groups). A particular episode in this environment can be described by the tuple $D = \langle A, r, o, T_1, \dots, T_n \rangle$ where n is a random variable drawn from an unknown distribution with location parameter (central tendency) of η . Note that we make almost no assumptions about this distribution; its characteristics will differ for different environments. For example, in the description of our DSN simulation early in Section 2 we noted the physical process by which vehicle tracks were generated and that the length of the tracks was not normally distributed.

Each task group T_i is associated with a track of length l_i and has the same basic objective structure, based on the DVMT:

- l_i Vehicle Location Methods (VLM's) that represent processing raw signal data at a single location to a single vehicle location hypothesis.
- $l_i - 1$ Vehicle Tracking Methods (VTM's) that represent short tracks connecting the results of the VLM at time t with the results of the VLM at time $t + 1$.
- 1 Vehicle Track Completion Method (VCM) that represents merging all the VTM's together into a complete vehicle track hypothesis.

Non-local precedence effects exist between each method at one level and the appropriate method at the next level as shown in Figure 4—two VLMs precede each VTM, and all VTM's precede the lone VCM.

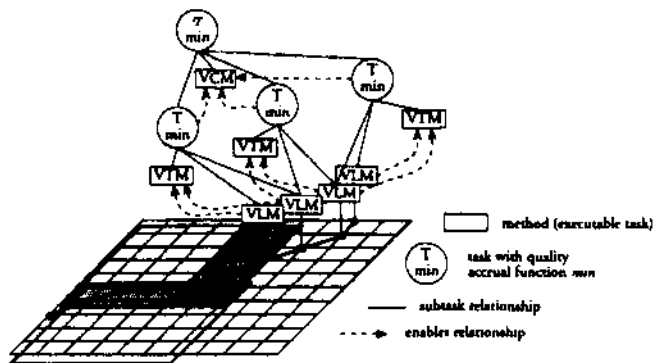


Figure 4: Objective task structure associated with a single vehicle track.

If we assume that each VLM has initial duration $d_0(\text{VLM})$ and each VTM has the initial duration $d_0(\text{VTM})$, then we can see from the task structure that for each task group the total execution time taken by a single processor agent will be:

$$l_i d_0(\text{VLM}) + (l_i - 1) d_0(\text{VTM}) + d_0(\text{VCM}) \quad (4)$$

This task structure is a simplification of the real DVMT task structure. For example, there is no sensor noise (which will cause *facilitation* relationships between tasks, and there is no confusion caused by 'ghost tracks'. Adding these features to the task structure will cause some interesting phenomena that we will discuss briefly in the conclusions (see also [Decker and Lesser, 1993b] for a representation of these extensions).

3 Static vs. Dynamic Organizational Structures

Now we have the necessary background to analyze static and dynamic organizational structures. The key to static structures is to divide up the overlap area *a priori* (rather than to penalize agents for doing redundant work in the overlap area [Durfee *et al.*, 1987]). The key to dynamic organizational structures is to transfer tasks so that all the agents' resources are used efficiently. We will repeat the assumptions we discussed at the start of Section 2 on page 2: the agents are homogeneous, cooperative, the data for each episode arrives in a single burst, and the only non-local effect is precedence.

3.1 Analyzing Static Organizations

In a static organization, agents divide the overlapping areas of their ranges as evenly as possible. The result is a new area of responsibility $r' = r - \frac{o}{2}$ for each agent with no overlap (see Figure 5)² Given the task structure as described in Section 2.5 and shown in Figure 4, and any raw data or communicated task results provided by information gathering actions, the agent can at any time build a list of currently executable methods (under the set of precedence constraints). Also, at any time an agent can build a list of methods that need to be executed, but cannot be executed because their precedence constraints have not yet been met. The communication action in this algorithm is a broadcast of the highest level results of all the task groups an

²The reason for overlap will be apparent in dynamic structures—multiple agents can work in an overlapping area without paying any cost for communicating raw data between them. Overlap can also provide redundancy in case of agent failure.

agent has worked on. Each agent follows the same control algorithm (remember, all the raw data is available at the start) and terminates when all task groups are completed (either locally or by reception of the result from another agent):

```
(Repeat
  Do Information-Gathering-Action
  (Repeat
    Let E = [get set of currently executable methods]
    (For method In E
      Do Method-Execution-Action(mcthod))
    Until (null E))
  Do Communication-Action(broadcast highest-level results)
  Let W = [get set of methods waiting on precedence constraints]
  Until (null W))
```




Figure 5: Example of a 3x3 organization, $r = 11, o = 5$, with 5 tracks. The thick dark grey boxes outline the default static organization, where there is no overlap among agents.

In the environment $D = \langle A, \eta, r, o, T \rangle$, if we let S' represent the largest amount of initial low-level data in one task group seen by any agent, and a the total number of agents that see the task group (from Eqn. 3), then the amount of time it will take that agent to construct a complete solution is equal to the amount of time it will take for the initial information gathering action ($d_0(I)$) plus the amount of time to do all the local work ($S'd_0(VLM) + (S' - 1)d_0(VTM)$), communicate that work ($d_0(C)$), get the other agents' results ($d_0(I)$), plus the amount of time to combine results from the other $a - 1$ agents ($(a - 1)d_0(VTM)$), plus time to produce the final complete task group result ($d_0(VCM)$), plus communicate that result to everyone ($d_0(C)$). For simplicity we will assume that $d_0(I)$ and $d_0(C)$ are constant and do not depend on the amount of data. Note that the maximally loaded agent will be the last to finish any local work.

If the system sees n total task groups, then the expected size of the initial low-level data set at the maximally loaded agent can be derived from the marginal expected value for S given the joint distribution of \hat{S}_N (Eqn. 1) and \hat{N}_n (Eqn. 2):

$$E[\hat{S}] = \sum_{N=0}^n \sum_{s=0}^N g_{A,n,s}(N) g_{a,N,p}(s) \left(\tau s + \frac{r}{2}(N - s) \right) \quad (5)$$

Similar to the single task group case, the total time until termination for an agent receiving an initial data set of size S is the time to do local work, combine results from other agents, and build the completed results, plus two communication and information gathering actions:

$$\hat{S}d_0(VLM) + (\hat{S} - \hat{N})d_0(VTM) + (a - 1)\hat{N}d_0(VTM) +$$

$$\hat{N}d_0(VCM) + 2d_0(I) + 2d_0(C) \quad (6)$$

We can use Eq. 6 as a predictor by combining it with the probabilities for the values of \hat{S} and \hat{N} given in Eq. 5 (i.e., substitute Eq. 6 for $(\tau s + \frac{r}{2}(N - s))$ in Eq. 5).

We tested these predictions of Eqn. 6 versus the mean termination time of our DSN simulation over 10 repetitions in each of 43 randomly chosen environments from the design space [$2 \leq r \leq 10, 0 \leq o \leq r, 1 \leq \sqrt{A} \leq 5, 1 \leq N \leq 10$]. The durations of all tasks were set at 1 time unit, as were the duration of information gathering and communication actions. We used the simulation validation statistic suggested by Kleijnen [Kleijnen, 1987] (where \hat{y} = the predicted output by the analytical model and y = the output of the simulation):

$$z = \frac{y - \hat{y}}{(\text{Var}(y) + \text{Var}(\hat{y}))^{1/2}} \quad (7)$$

where $\text{Var}(\hat{y})$ is the predicted variance.³ The result z can then be tested for significance against the standard normal tables. In each case, we were unable to reject the null hypothesis that the actual mean termination equals the predicted mean termination at the $\alpha = 0.05$ level.⁴ Figure 6 shows the mean of 10 repetitions in each environment versus the expected value and its confidence intervals. Regression of the measured and predicted values results in $R^2 = 0.96$.

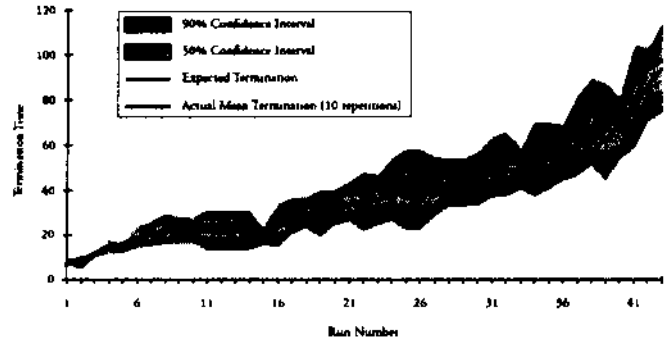


Figure 6: Actual system termination versus analytic expected value and analytically determined 50% and 90% confidence intervals. Runs arbitrarily ordered by expected termination time.

3.2 Control Costs

The control algorithm presented above is simple and not necessarily optimal. By communicating only when there is no local work to be done, the heaviest-loaded agent gives up the chance for other agents to do the high-level composition in a task group by incrementally transmitting partial results, a maximum potential savings of $(\hat{N} - 1)(a - 1)d_0(VTM) + (\hat{N} - 1)d_0(VCM)$. However, this needs to be balanced with the cost of multiple communication actions, which is $(\hat{N} - 1)d_0(C)$. Thus 'when to communicate' rests directly on the cost of communication relative to $(a - 1)d_0(VTM) + d_0(VCM)$ (which depends on

³The predicted variance of Eqn. 6 can be easily derived from the statistical identity $\text{Var}(x) = E[x^2] - (E[x])^2$.

⁴For non-statisticians: this is a good thing. The null hypothesis is that our prediction is the same as the actual value, we did not wish to reject it, and we did not.

both the basic method durations and the agents' organizational structure). This simple control algorithm can be analyzed easily, unlike many other systems where control costs are ignored. If we view the cost of control as the time spent by an agent when *not* performing an action (executing a method, information gathering, communication), then our algorithm runs in constant time between actions except for the two tests [*get set of currently executable methods*] and [*get set of methods still waiting*]. Each of these in the worst case requires a constant-time test of each element of the full task structure, which is of size $O(\eta r)$. Thus we see how the control costs, too, are related to organizational structure.

3.3 Analyzing Dynamic Organizations

In the dynamic organizational case, agents are not limited to the original organization and initial distribution of data. Agents can reorganize by changing the initial static boundaries (changing responsibilities in the overlapping areas), or by shipping raw data to other agents for processing (load balancing).

In the case of reorganized overlapping areas, agents may shift the initial static boundaries by sending a (very short) message to overlapping agents, telling the other agents to do all the work in the overlapping areas. The effect at the local agent is to change its effective range parameter from its static value of $r' = r - o/2$ to some value r'' where $r - o/2 \geq r'' \geq r - o$, changing the first two terms of Eqn. 6, and adding a communication action to indicate the shift and an extra information gathering action to receive the results. Another paper [Decker and Lesser, 1993a] discusses a particular implementation of this idea that chooses the partition of the overlapping area that best reduces expected differences between agent's loads and averages competing desired partitions from multiple agents.

In the second case, an agent communicates some proportion p of its initial data to a second agent, who does the associated work and communicates the results back. Instead of altering the effective range and overlap, this method directly reduces the first two terms of Eqn. 6 by the proportion p . The proportion p can be chosen dynamically in a way similar to that of choosing where to partition the overlap between agents (see [Decker and Lesser, 1993a]).

Whether or not a dynamic reorganization is useful is a function of both the agents local utility and also the load at the other agent. We will again be concentrating on the agent with the heaviest load. Looking first at the local utility, to do local work under the initial static organization with n task groups, the heaviest loaded agent will take time:

$$S' d_0(\text{VLM}) + (S' - n) d_0(\text{VTM}) \quad (8)$$

When the static boundary is shifted before any processing is done, the agent will take time:

$$d_0(C_{\text{short}}) + S'' d_0(\text{VLM}) + (S'' - n) d_0(\text{VTM}) + d_0(I) \quad (9)$$

to do the same work, where C_{short} is a very short communication action which is potentially much cheaper than the result communications mentioned previously, and S'' is calculated using r'' . When balancing the load directly, local actions will take time:

$$d_0(C_{\text{long}}) + \rho S' d_0(\text{VLM}) + \rho(S' - n) d_0(\text{VTM}) + d_0(I) \quad (10)$$

where $d_0(C_{\text{long}})$ is potentially much more expensive than the communication actions mentioned earlier (since it involves sending a large amount of raw data). If the other agent had no work

to do, a simple comparison between these three equations would be a sufficient design rule for deciding between static and either dynamic organization.

4 Using Meta-Level Communication

For some environments $\mathcal{D} = \langle A, \eta, r, o, T \rangle$ one of the three organizational choices may be clearly better in the long run, but for most environments the choice is not so clear given the variance in system performance. The choice that optimizes performance over the long run is often not optimal in any particular episode. Taking the equations for local work in Section 3.3 along with Eqn. 5, we can compute confidence intervals on the predicted performance of an organization under each of the three coordination regimes by combining the local confidence interval on the expected load of the heaviest loaded agent, and the confidence interval on the average agent load. These results, for the 50% confidence interval, are shown in Figures 7 and 8. Again we have assumed that all execution, communication, and information gathering action durations have the same value (making communication relatively expensive). The first figure, Figure 7, highlights how the relationship between performance under a static organization and a dynamically load balanced organization changes as the number of agents increases. As expected, load balancing becomes more desirable as the number of agents increases (in relation to the average number of tracks): when there are many agents, the average agent load becomes very low, which offsets the cost of transferring tasks. In this figure the performance difference between static and overlap reorganization remains nearly constant relative to the number of agents.




Figure 7: Predicted 50% confidence intervals on the expected termination of a system under three coordination regimes, different numbers of agents, and three values of n (number of tracks). This figure is based on Eqns. 8, 9, and 10.

The second, Figure 8, points out how dynamically reorganizing the overlap area increases the performance over static organization as the amount of overlap increases. For this graph we assumed that the agents would shrink their entire area of responsibility (as opposed to minimizing the difference in maximum versus average work as described in [Decker and Lesser, 1993a]). This graph shows the need for dynamically calculating the shrinkage parameter (p) especially at high levels of overlap (note how the dynamically reorganized organization is predicted to do worse at high levels of overlap in the $n = 20$ portion). The expected performance difference between the static organization and load balancing remains relatively constant across changing values of o . In both figures we have let $d_0(C) = d_0(C_{\text{short}}) = d_0(C_{\text{long}})$; changing these values will move the corresponding curves directly up or down.




Figure 8: Predicted 50% confidence intervals on the expected termination of a system under three control regimes, different overlaps, and three values of n (number of tracks).

These figures bring us to the final point of this paper: often system performance can be improved significantly by dynamic reorganization, but it will rarely *always* be improved. Therefore, meta-level communication between agents about their local loads can, with a small communication cost, pinpoint the true costs and benefits of the various organizational structures, allowing an informed organizational decision to be made. Instead of an agent making a decision about restructuring or load balancing by assuming the *average* load, the agent will have the *actual* load for the neighboring agents. As we said in the introduction, the proper organization is often one that exploits *information that resolves uncertainties* about the current environment as it becomes available to the agents, allowing the agents to then create the most efficient organization for the situation.

5 Conclusions

The results of this paper can be looked at from three points of view. From the practitioners viewpoint, the analysis presented here resulted in a set of design equations that can be used directly to optimize the performance of a simple DSN, or explore the design space given some model of how expensive agents are and what bounds (mean, median, 90% quantile) on their performance are required. Several of the simplifying assumptions we used, such as constant communication and information gathering costs, can be easily replaced with submodels chosen by the designer. From the viewpoint of the distributed AI community, we have returned to look at some of the problems first studied by Durfee, Lesser, and Corkill [Durfee *et al.*, 1987]. They concluded that "Our intent is to show that overly specialized organizational structures allow effective network performance in particular problem-solving situations, but that no such organization is appropriate in all situations." In this paper we reach the same abstract conclusion, but also show precisely what the effect is of a particular organizational structure (characterized by both its structural components and its coordination algorithm) in an environment (characterized by the structure and frequency of its tasks) in a clear way that not only allows us to predict performance but to explain it. The technique of using binomial approximations should also prove useful in different domains. From the viewpoint of the general research community this paper presents a methodology for answering questions about the design of a system by analysis and simulation. In such a methodology, the observation of particular phenomena in a complex system (the DVMT) leads to the building and verification of general models that predict and explain such phenomena.

In the short term, this work leads to the explanation of other interesting distributed problem solving phenomena displayed in [Durfee *et al.*, 1987]. The addition of noise at DSN sensors leads to the necessity of more complex coordination with the introduction of more complex task interrelationships (such as *facilitation* [Decker and Lesser, 1993b]). The addition of correlated noise in the environment can then cause these new, more complex coordination mechanisms to break down, producing the phenomenon recognized as *distraction*. In the long term, we are working towards a complete characterization of generalized partial global planning [Decker and Lesser, 1992] as a first step towards a theory of coordination in distributed problem solving.

A longer technical report version of this paper is available from the authors.

References

- [Cohen and Levesque, 1990] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3), 1990.
- [Davis and Smith, 1983] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63-109, January 1983.
- [Decker and Lesser, 1992] K. S. Decker and V. R. Lesser. Generalizing the partial global planning algorithm. *International Journal of Intelligent and Cooperative Information Systems*, 1(2), June 1992.
- [Decker and Lesser, 1993a] K. S. Decker and V. R. Lesser. A one-shot dynamic coordination algorithm for distributed sensor networks. In *Proc. Eleventh National Conference on Artificial Intelligence*, Washington, July 1993.
- [Decker and Lesser, 1993b] K. S. Decker and V. R. Lesser. Quantitative modeling of complex computational task environments. In *Proc. Eleventh National Conference on Artificial Intelligence*, Washington, July 1993.
- [Durfee and Montgomery, 1991] E. H. Durfee and T. A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6): 1363-1378, November 1991.
- [Durfee *et al.*, 1987] E. H. Durfee, V. R. Lesser, and D. D. Corkill. Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, 36(11): 1275-1291, November 1987.
- [Galbraith, 1977] J. Galbraith. *Organizational Design*. Addison-Wesley, Reading, MA, 1977.
- [Kleijnen, 1987] J. P. C. Kleijnen. *Statistical Tools for Simulation Practitioners*. Marcel Dekker, New York, 1987.
- [Lesser and Corkill, 1983] V. R. Lesser and D. D. Corkill. The distributed vehicle monitoring testbed. *AI Magazine*, 4(3):63-109, Fall 1983.
- [Moe, 1984] T. M. Moe. The new economics of organization. *American Journal of Political Science*, 28(4):739-777, November 1984.
- [Shoham, 1991] Y. Shoham. AGENT0: A simple agent language and its interpreter. In *Proc. Ninth National Conference on Artificial Intelligence*, pages 704-709, Anaheim, July 1991.
- [Stinchcombe, 1990] A. L. Stinchcombe. *Information and Organizations*. Univ. of California Press, Berkeley, CA, 1990.