Quantitative Modeling of Complex Environments [1]

Keith Decker and Victor Lesser
Department of Computer Science
University of Massachusetts

UMass Computer Science Technical Report 93–21
April 18, 1994

Abstract

There are many formal approaches to specifying how the mental state of an agent entails the particular actions it will perform. These approaches put the agent at the center of analysis. For some questions and purposes, it is more realistic and convenient for the center of analysis to be the task environment, domain, or society of which agents will be a part. This paper presents such a task environment-oriented modeling framework that can work hand-in-hand with more agent-centered approaches. Our approach features careful attention to the quantitative computational interrelationships between tasks, to what information is available (and when) to update an agent's mental state, and to the general structure of the task environment rather than single-instance examples. A task environment model can be used for both analysis and simulation, it avoids the methodological problems of relying solely on single-instance examples, and provides concrete, meaningful characterizations with which to state general theories. This paper is organized around an example model of cooperative problem solving in a distributed sensor network.

[1]This paper, which extends our AAAI-93 paper, will appear in the *International Journal of Intelligent Systems in Accounting, Finance, and Management*, in a special issue on mathematical and computational models of organizations, Vol. 2, No. 4, December 1993. 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.

1 Introduction

This paper presents a framework, TÆMS (Task Analysis, Environment Modeling, and Simulation), with which to model complex computational task environments that is compatible with both formal agent-centered approaches and experimental approaches. The framework allows us to both analyze and quantitatively simulate the behavior of single or multi-agent systems with respect to interesting characteristics of the computational task environments of which they are part. We believe that it provides the correct level of abstraction for meaningfully evaluating centralized, parallel, and distributed control algorithms, negotiation strategies, and organizational designs. No previous characterization formally captures the range of features, processes, and especially interrelationships that occur in computationally intensive task environments.

We use the term *computational task environment* to refer to a problem domain in which the primary resources to be scheduled or planned for are the computational processing resources of an agent or agents, as opposed to physical resources such as materials, machines, or men. Examples of such environments are distributed sensor networks, distributed design problems, complex distributed simulations, and the control processes for almost any distributed or parallel AI application. We have recently extended the TÆMS framework to model physical resources as well (see Section 4).

The *reason* we have created the TÆMS framework is rooted in the desire to produce general theories in AI [6]. Consider the difficulties facing an experimenter asking under what environmental conditions a particular local scheduler produces acceptable results, or when the overhead associated with a certain coordination algorithm is acceptable given the frequency of particular subtask interrelationships. At the very least, our framework provides a characterization of environmental features and a concrete, meaningful language with which to state correlations, causal explanations, and other forms of theories. The careful specification of the computational task environment also allows the use of very strong analytic or experimental methodologies, including paired-response studies, ablation experiments, and parameter optimization. TÆMS exists as both a language for stating general hypotheses or theories and as a system for simulation. The simulator supports the graphical display of generated subjective and objective task structures, agent actions, and statistical data collection in CLOS on the TI Explorer and DEC Alpha.

The basic form of the computational task environment framework—the execution of interrelated computational tasks—is taken from several domain environment simulators [4, 5, 16]. If this were the only impetus, the result might have been a simulator like Tileworld [28]. However, formal research into multi-agent problem solving has been productive in specifying formal properties, and sometimes algorithms, for the control process by which the mental state of agents (termed variously: beliefs, desires, goals, intentions, etc.) causes the agents to perform particular actions [7, 30, 35]. This research has helped to circumscribe the behaviors or actions that agents can produce based on their knowledge or beliefs. The final influence on TÆMS was the desire to avoid the individualistic agent-centered approaches that characterize most AI and Distributed AI. The concept of agency in TÆMS is based on simple notions of *execution*, *communication*, and *information gathering*. An agent is a locus of belief (state) and action. By separating as much as possible the notion of agency from the model of task environments, we do not have to subscribe to particular agent architectures (which one would assume will be adapted to the task environment at hand), and we may ask questions about the inherent social nature of the task environment at hand (allowing that the concept of society may arise before the concept of individual agents). The form of the framework is more detailed in structure than many organizational-theoretic models of organizational environments, such as Thompson's notions of pooled, sequential, and reciprocal processes[34], Burton and Obel's linear programs[3], or Malone's queueing models [26], but is influenced by them, and by the importance of environmental uncertainty and dependency that appear in contingency-theoretic and open systems views of organizations [24, 18, 32, 29]

Section 2 will discuss the general nature of the three modeling framework layers. Sections 3 through 6 discuss the details of the three levels, and are organized around a model built with this framework for the study of organizational design and coordination strategies in a multi-agent distributed sensor network environment. Section 6.1 will briefly show the form that mathematical analyses take in the framework using the distributed sensor network example, while Section 6.2 will show the use of the framework in a simulation experiment, based on Burton and Obel, that explores the effect of task decomposability and the availability of non-local information on team performance.

2 General Framework

The principle purpose of a TÆMS model is to analyze, explain, or predict the performance of a system or some component. While TÆMS does not establish a particular performance criteria, it focuses on providing two kinds of performance information: the temporal intervals of task executions, and the *quality* of the execution or its result. *Quality* is an intentionally vaguely-defined term that must be instantiated for a particular environment and performance criteria. Examples of *quality* measures include the precision, belief, or completeness of a task result. We will assume that *quality* is a single numeric term with an absolute scale, although the algebra can be extended to vector terms. In a computationally intensive AI system, several quantities—the quality produced by executing a task, the time taken to perform that task, the time when a task can be started, its deadline, and whether the task is necessary at all—are affected by the execution of other tasks. In real-time problem solving, alternate task execution methods may be available that trade-off time for quality. Agents do not have unlimited access to the environment; what an agent believes and what is really there may be different.

The model of environmental and task characteristics proposed has three levels: *objective*, *subjective*, and *generative*. The *objective* level describes the essential, 'real' task structure of a particular problem-solving situation or instance over time. It is roughly equivalent to a formal description of a single problem-solving situation such as those presented in [17], without the information about particular agents. The *subjective* level describes how agents view and interact with the problem-solving situation over time (e.g., how much does an agent know about what is really going on, and how much does it cost to find out—where the uncertainties are from the agent's point of view). The subjective level is essential for evaluating control algorithms, because while individual behavior and system performance can be measured objectively, agents must make decisions with only subjective information. In organizational theoretic terms, subjective *perception* can be used to predict agent actions or *outputs*, but unperceived, objective environmental characteristics can still affect performance (or *outcomes*) [29]. Finally, the *generative* level describes the statistical characteristics required to generate the objective and subjective situations in a domain. A generative level model consists of a description of the generative processes or distributions from which the range of alternative problem instances can be derived, and is used to study performance over a range of problems in an environment, avoiding single-instance examples.

2.1 Distributed Sensor Network Example

Before we discuss formally the basic components of an objective model, let us turn to an example in which we will build a model using the TÆMS framework. This example grows out of the set of single instance examples of distributed sensor network (DSN) problems presented in [16]. The authors of that paper compared the performance of several different coordination algorithms on these examples, and concluded that no one algorithm was always the best—not a very surprising result and one that can be viewed as the central tenet of contingency theory. This is the classic type of result that the TÆMS framework was created to address—we wish to *explain* this result, and better yet, to *predict* which algorithm will do the best in each situation. Here we do this by extending the analysis to take into account a statistical view of the environmental characteristics. The level of detail to which you build a model in TÆMS will depend on the question you wish to answer—here we wish to identify the characteristics of the DSN environment, or the organization of the agents, that cause one algorithm to outperform another.

In a DSN problem, 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 is based on how long it takes them to create complete vehicle tracks, including the cost of communication. The organizational structure of the agents will imply the portions of each vehicle track that are sensed by each agent.

In our model of DSN problems, each vehicle track is modeled as a task group (problem instance). Several task groups (vehicle tracks) may occur simultaneously in a single problem solving *episode*. The simplest objective model is that each task group $\mathcal{T}_i$ is associated with a track of length l_i and has the following objective structure, based on a simplified version of the DVMT: (l_i) vehicle location methods (VLM) that represent processing raw signal data at a single location resulting in a single vehicle location hypothesis; $(l_i - 1)$ vehicle tracking methods (VTM) that represent short tracks connecting the results of the VLM at time t with the results of the VLM at time $t + 1$; and one vehicle track completion method (VCM) that represents merging all the VTMs together into a complete vehicle track hypothesis. *Non-local effects*, which relate the executions of tasks to one another, exist as shown in Figure 1—two VLMs *enable* each VTM, and all VTMs *enable* the lone VCM. This structure is fairly common in many other environments, where a large amount of detailed work needs to be done, the results of which are collected at a single location or agent and processed again (integrated), and so on up a hierarchy. The question is: to achieve a particular level of performance, how many agents are needed and how should they be organized in their unique and/or overlapping capabilities for accomplishing the necessary tasks. Coordination is used not only to accomplish the necessary transfer of results from one agent to another, but to avoid redundant work on the part of agents with overlapping capabilities, and also to potentially balance the workloads.

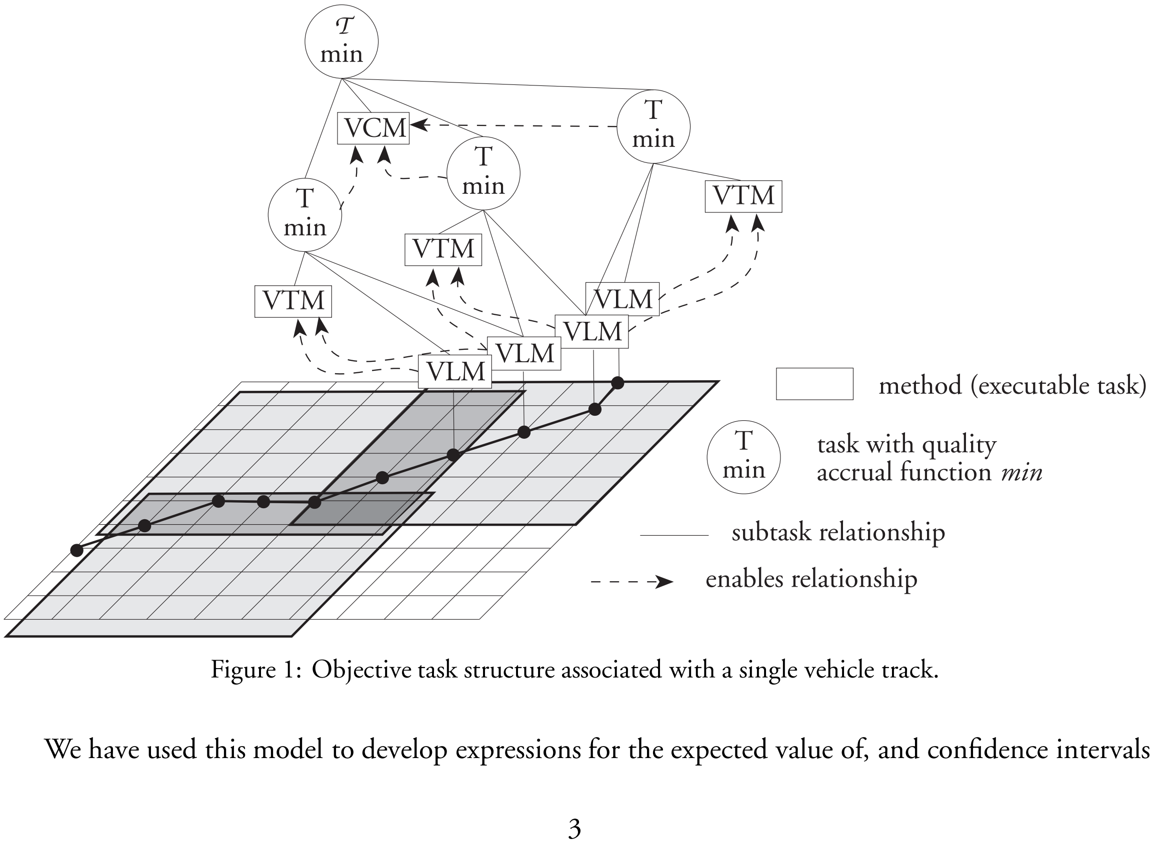

Figure 1: Objective task structure associated with a single vehicle track.

We have used this model to develop expressions for the expected value of, and confidence intervals on, the time of termination of a set of agents in any arbitrary DSN environment that has a static organizational structure and coordination algorithm [12, 13]. We have also used this model to analyze a dynamic, one-shot reorganization algorithm (and have shown when the extra overhead is worthwhile versus the static algorithm). In each case we can predict the effects of adding more agents, changing the relative cost of communication and computation, and changing how the agents are organized (in the range and overlap of their capabilities). These results were achieved by direct mathematical analysis of the model and verified through simulation in TÆMS. We will give a summary of these results later in the paper (Section 6), after discussing the detailed definitions of the objective, subjective, and generative levels.

3 Objective Level

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. The basic model is that *task groups* (problem instances) appear in the environment at some frequency, and induce tasks T to be executed by the agents under study. Task groups are independent of one another, but tasks within a single task group have interrelationships. Task groups or tasks may have deadlines $D(T)$. The *quality* of the execution or result of each task influences the *quality* of the task group result $Q(T)$ in a precise way (Section 3.1). These quantities can be used to evaluate the performance of a system.

An individual task that has no subtasks is called a method M and is the smallest schedulable chunk of work (though some scheduling algorithms will allow some methods to be preempted, and some schedulers will schedule at multiple levels of abstraction). There may be more than one method to accomplish a task, and each method will take some amount of time and produce a result of some *quality*. Quality of an agent's performance on an individual task is a function of the timing and choice of agent actions ('local effects'), and possibly other (previous or future) task executions ('non-local effects'). When local or non-local effects exist between tasks that are known by more than one agent, we call them *coordination relationships*[11]. The basic purpose of the objective model is to formally specify how the execution and timing of tasks affect this measure of quality.

3.1 Local Effects: The *Subtask* Relationship

Task or task group quality ($Q(T)$) is based on the *subtask* relationship. This quality function is constructed recursively—each task group consists of tasks, each of which consists of subtasks, etc.—until individual executable tasks (methods) are reached. Formally, the subtask relationship is defined as $\mathsf{SUBTASK}(T, \mathbf{T}, Q)$, where $\mathbf{T}$ is the set of all direct subtasks of T and Q is a quality function $Q(T, t)$: [tasks $\times$ times] $\mapsto$ [quality] that returns the quality associated with T at time t. In a valid model, the directed graph induced by this relationship is acyclic (no task has itself for a direct or indirect subtask).

The semantics of a particular environment are modeled by the appropriate choice of the quality function Q (e.g., minimum, maximum, summation, or the arithmetic mean). For example, if $\mathsf{SUBTASK}(T_1, \mathbf{T}, Q_{\min})$, then $Q(T_1, t) = Q_{\min}(T_1, t) = \min_{T\in\mathbf{T}} Q(T, t)$. In this case the quality that is associated with task T_1 is the minimum quality associated with any of its subtasks. This is sometimes referred to as an AND because the quality of the parent remains at a minimum until every subtask has been completed. Other functions may be used for modeling particular environments.[1]

[1]The set of possible aggregation operators include three basic classes: *conjunctions*, *disjunctions*, and *trade-offs*. Dubois and Prade have shown that the Triangular norms (including min), averaging operators (including mean), and Triangular conorms (including max and summation) are the most general families of binary functions that

Functions like sum and average indicate the possibility that not all tasks in the environment need to be carried out—this is a major difference between this work and most work in traditional distributed computing. Computational task structures may represent search processes where agents are satisficing, not optimizing. We have now described how quality is modeled at tasks that have subtasks, and now turn our attention to methods.

3.2 Local Effects: Method Quality

Each method M at a time t will potentially produce (if executed by an agent, see Section 5.3) some *maximum quality* $\mathbf{q}(M,t)$ after some amount of elapsed time $\mathbf{d}(M,t)$ (we will defer any further definition of the functions $\mathbf{d}$ and $\mathbf{q}$ until we discuss non-local effects in Section 3.3). The execution of methods is interruptible, and if multiple methods for a single task are available, the agent may switch between them (typically, alternative methods tradeoff time and quality). We model the effect of interruptions, if any, and the reuse of partial results as non-local effects (see Section 3.3).

Let $\mathsf{P}(M,t)$ be the current amount of progress on the execution of M. If M were not interruptible and $\mathsf{S}(M)$ and $\mathsf{F}(M)$ were the execution start time and finish time, respectively, of M, then:

$$\mathsf{P}(M,t) = \begin{cases} 0 & t \leq \mathsf{S}(M) \\ t - \mathsf{S}(M) & \mathsf{S}(M) < t < \mathsf{F}(M) \\ \mathsf{F}(M) - \mathsf{S}(M) & t \geq \mathsf{F}(M) \end{cases}$$

We typically model the quality produced by a method $Q(M,t)$ using a linear growth function Q_{lin}:

$$Q_{\text{lin}}(M,t) = \begin{cases} \frac{\mathsf{P}(M,t)}{\mathbf{d}(M,t)}(\mathbf{q}(M,t)) & \mathsf{P}(M,t) < \mathbf{d}(M,t) \\ \mathbf{q}(M,t) & \mathsf{P}(M,t) \geq \mathbf{d}(M,t) \end{cases}$$

Other models (besides linear quality functions) have been proposed and are used, such as concave quality functions (must execute most of a task before quality begins to accumulate), convex quality functions (most quality is achieved early on in a method, and only small increases occur later), and 'mandatory and optional parts' quality functions [25]. The desired $Q(M,t)$ can be easily defined for any of these.

As an example of the power of this representation, we consider the two main schools of thought on quality accumulation: the anytime algorithm camp [1] and the design-to-time (approximate processing) camp[14, 20]. We can represent their ideas succinctly; in the anytime algorithm model partial results are always available,[2] as in the definition of $Q_{\text{lin}}(M,t)$ above, while in the design-to-time model results are not available (quality does not accrue) until the task is complete, as in the definition of $Q_{\text{DTT}}(M,t)$:

$$Q_{\text{DTT}}(M,t) = \begin{cases} 0 & \mathsf{P}(M,t) < \mathbf{d}(M,t) \\ \mathbf{q}(M,t) & \mathsf{P}(M,t) \geq \mathbf{d}(M,t) \end{cases}$$

Another difference between design-to-time (DTT) and other approaches will show up in our generative and subjective additions to this model—DTT does not assume that $Q(M,t)$ is fixed and known, but rather that it is an estimator for the actual method response. Finally, deadlines can be associated with task groups or individual tasks, and quality accumulation defined such that any work done on a task after its deadline does not increase quality.

3.3 Non-local Effects

Any task T containing a method that starts executing before the execution of another method M finishes may potentially affect M's execution through a *non-local effect* e. We write this relation

respectively satisfy the semantic requirements of the three basic classes of aggregation operators [2, 15].

[2] In Boddy's paper, the assumption is made that $Q(M,t)$ has monotonically decreasing gain.

$\mathsf{NLE}(T, M, e, p_1, p_2, \ldots)$, where the p's are parameters specific to a class of effects. There are precisely two possible outcomes of the application of a non-local effect on M under our model: *duration effects* where $\mathbf{d}(M, t)$ (duration) is changed, and *quality effects* where $\mathbf{q}(M, t)$ (maximum quality) is changed. An effect class e is thus a function $e(T, M, t, d, q, p_1, p_2, \ldots)$: [task $\times$ method $\times$ time $\times$ duration $\times$ quality $\times$ parameter 1 $\times$ parameter 2 $\times \ldots$] $\mapsto$ [duration $\times$ quality].

The amount and direction of an effect is dependent on the relative timing of the method executions, the quality of the effect's antecedent task, and whether information was communicated between the agents executing the methods (in multi-agent models). Some effects are continuous, depending on the current quality of the effect's antecedent $Q(T, t)$. Some effects are triggered by a rising edge of quality past a threshold; for these effects we define the helper function $\Theta(T, \theta)$ that returns the earliest time when the quality surpasses the threshold: $\Theta(T, \theta) = \min(t)$ s.t. $Q(T, t) > \theta$.

Communication. Some effects depend on the availability of information to an agent. We indicate the communication of information at time t about task T_a to an agent A with a delay of δ_t by $\mathsf{COMM}(T_a, A, t, \delta_t)$. There are many models of communication channels that we could take for a communication submodel; since it is not our primary concern we use a simple model with one parameter, the time delay δ_t.[3] For defining effects that depend on the availability of information, we define the helper function $Q_{\text{avail}}(T, t, A)$ that represents the quality of a task T 'available' to agent A at time t. If T was executed at A, $Q_{\text{avail}}(T, t, A) = Q(T, t)$. If T was executed (or is being executed) by another agent, then the 'available' quality is calculated from the last communication about T received at agent A prior to time t.

***Computing* $\mathbf{d}(M, t)$ *and* $\mathbf{q}(M, t)$.** Each method has an initial maximum quality $\mathbf{q}_0(M)$ and duration $\mathbf{d}_0(M)$ so we define $\mathbf{q}(M, 0) = \mathbf{q}_0(M)$ and $\mathbf{d}(M, 0) = \mathbf{d}_0(M)$. If there is only one non-local effect with M as a consequent $\mathsf{NLE}(T, M, e, p_1, p_2, \ldots)$, then $[\mathbf{d}(M, t), \mathbf{q}(M, t)] \leftarrow e(T, M, t, \mathbf{d}(M, t-1), \mathbf{q}(M, t-1), p_1, p_2, \ldots)$. If there is more than one NLE, then the effects are applied one after the other in an order specified in the model (the default is for effects with antecedents closer in the task structure to M to be applied first).

The maximum quality function $\mathbf{q}$ can also be defined for tasks or task groups. The precise definition depends on the set of quality accrual functions in the model. Using the four quality accrual functions we have already discussed (minimum, maximum, summation, mean) the definition of maximum quality for a non-method task $\mathbf{q}(T, t)$ is as follows:

$$\mathbf{q}(T, t) = \begin{cases} \max_{x \in \mathbf{T}} \mathbf{q}(x, t) & \text{if } \mathsf{SUBTASK}(T, \mathbf{T}, Q_{\max}) \\ \min_{x \in \mathbf{T}} \mathbf{q}(x, t) & \text{if } \mathsf{SUBTASK}(T, \mathbf{T}, Q_{\min}) \\ \sum_{x \in \mathbf{T}} \mathbf{q}(x, t) & \text{if } \mathsf{SUBTASK}(T, \mathbf{T}, Q_{\Sigma}) \\ \frac{\sum_{x \in \mathbf{T}} \mathbf{q}(x, t)}{|\mathbf{T}|} & \text{if } \mathsf{SUBTASK}(T, \mathbf{T}, Q_{\text{mean}}) \end{cases}$$

The current duration function $\mathbf{d}$ has no meaningful objective definition when applied to non-method tasks. "Maximum duration" could be defined, but is generally a useless concept. A more useful concept for scheduling—the minimum duration required for achieving maximum quality at a task—is explored in [19]. The clear specification of such concepts is one of the benefits of using our framework.

3.3.1 Non-local Effect Examples

Non-local effects are the most important part of the TÆMS framework, since they supply most of the characteristics that make one task environment unique and different from another. Typically a

[3] Other parameters, such as channel reliability, are being considered. The description of an agent's control and coordination algorithms will describe when and where communication actually occurs (see communication actions in Section 5.3, and the concept of agency in Section 5.1).

model will define different classes of effects, such as *causes*, *facilitates*, *cancels*, *constrains*, *inhibits*, and *enables* [10]. This section contains definitions for four common classes of effects that have been useful in modeling different environments. When non-local effects occur between methods associated with different agents, we call them *coordination relationships* [11, 10].

Enables. If task T_a enables method M, then the maximum quality $\mathbf{q}(M,t) = 0$ until T_a is completed and the result is available, when the maximum quality will change to the initial maximum quality $\mathbf{q}(M,t) = \mathbf{q}_0(M)$. Another way to view this effect is that it changes the "earliest start time" of *enabled* method, because a rational scheduler will not execute the method before it is enabled.

$$\textbf{enables}(T_a, M, t, d, q, \theta) = \begin{cases} [\infty, 0] & t < \Theta(T_a, \theta) \\ [\mathbf{d}_0(M), \mathbf{q}_0(M)] & t \geq \Theta(T_a, \theta) \end{cases} \tag{1}$$

Facilitates. Another effect, used by the PGP algorithm [17] but never formally defined, is the *facilitates* effect. In organization theory, Stinchcombe defines facilitation as an 'external effect' in relation to the spatial properties of environmental activities [31], but also recognizes computational facilitation through the transmission of information. Computationally, facilitation occurs when information from one task, often in the form of constraints, is provided that either reduces or changes the search space to make some other task easier to solve. In our framework, one task may provide results to another task that *facilitate* the second task by decreasing the duration or increasing the quality of its partial result. Therefore the *facilitates* effect has two parameters (called *power* parameters) $0 \leq \phi_{\rm d} \leq 1$ and $0 \leq \phi_{\rm q} \leq 1$, that indicate the effect on duration and quality respectively. The effect varies not only through the power parameters, but also through the quality of the *facilitating* task available when work on the *facilitated* task starts (the ratio R). Note that before work is started on a method, $\mathsf{S}(M) = t$ (i.e. formulae are evaluated as if execution were about to start).

$$\begin{aligned} R(T_a, s) &= \frac{Q_{\rm avail}(T_a, s)}{\mathbf{q}(T_a, s)} \\ \textbf{facilitates}(T_a, M, t, d, q, \phi_{\rm d}, \phi_{\rm q}) &= [d(1 - \phi_{\rm d} R(T_a, \mathsf{S}(M))), \\ &\quad\; q(1 + \phi_{\rm q} R(T_a, \mathsf{S}(M)))] \end{aligned} \tag{2}$$

So if T_a is completed with maximal quality, and the result is received before M is started, then the duration $\mathbf{d}(M,t)$ will be decreased by a percentage equal to the duration power $\phi_{\rm d}$ of the *facilitates* effect. The second clause of the definition indicates that communication after the start of processing has no effect. In other work [11] we explored the effects on coordination of a *facilitates* effect with varying duration power $\phi_{\rm d}$, and with $\phi_{\rm q} = 0$.

Hinders. The *hinders* effect is the opposite of *facilitates*, because it increases the duration and decreases the maximum quality of the consequent. It can be used expressively in a model to represent situations that reduce some baseline performance measures (as opposed to *facilitation*, which represents an increase in the baseline). Such situations occur computationally when there are multiple methods for a task and one fails to provide as many constraints as the other. *Hinders* can also be used as a high-level model of distraction [16], where inappropriate constraints communicated from one agent to another lead the second agent on a wild goose chase.

Precedence. We define the *precedence* effect as a combination of *enables* and *hinders*. That is, one task must be done before another *and* the first task must be done well or later tasks will suffer for it. If T_a *precedes* M, then M has infinite duration and 0 maximum quality until some quality is accrued at T_a. Afterwards, the duration drops toward the initial value and the maximum quality increases to the initial value according to the ratio of available and maximum quality and the precedence effect's power parameters ($\phi_{\rm d} \geq 0$ and $\phi_{\rm q} \geq 0$). The following formula is more easily understood if one keeps in mind that, in general, the ratio of available quality to maximum quality will go from 0 to 1 as methods are executed.

$$\textbf{precedes}(T_a, M, t, d, q, \phi_{\rm d}, \phi_{\rm q}) = \begin{cases} [\infty, 0] & Q_{\rm avail}(T_a, \mathsf{S}(M)) \leq 0 \\ [\mathbf{d}_0(M)/R(T_a, \mathsf{S}(M))^{\phi_{\rm d}}, & \\ \mathbf{q}_0(M) R(T_a, \mathsf{S}(M))^{\phi_{\rm q}}] & Q_{\rm avail}(T_a, \mathsf{S}(M)) > 0 \end{cases} \tag{3}$$

3.4 Expanding the DSN Model

We will now add some complexity to the distributed sensor network model. The length of a track l_i mentioned earlier is a *generative* level parameter. Given a set of these generative parameters, we can construct the objective model for a specific problem-solving instance, or *episode*. Figure 1 shows the general structure of episodes in our DSN environment model, rather than a particular episode. To display an actual objective model, let us assume a simple situation: there are two agents, A and B, and that there is one vehicle track of length 3 sensed once by A alone (T^1), once by both A and B (T^2), and once by B alone (T^3). We now proceed to model the standard features that have appeared in our DVMT work for the past several years. We will add the characteristic that each agent has two methods with which to deal with sensed data: a normal VLM and a 'level-hopping' (LH) VLM (the level-hopping VLM produces less quality than the full method but requires less time; see [14, 9] for this and other approximate methods; a similar technique can be used to model agents who have different capabilities such as processor speed). Furthermore, only the agent that senses the data can execute the associated VLM; but any agent can execute VTMs and VCMs if the appropriate enablement conditions are met.

Figure 2 displays this particular problem-solving episode. To the description above, we have added the fact that agent B has a faulty sensor (the durations of the grayed methods will be longer than normal); we will explore the implications of this after we have discussed the subjective level of the framework in the next section. An assumption made in [16] is that redundant work is not generally useful; this is indicated by using *max* as the combination function for each agent's redundant methods. This would *not* be a good assumption for an environment where the independent derivation of a result is a *positive* thing—we can alter the assumption by simply changing this function (to *mean* or *sum*). Another characteristic that appeared often in the DVMT literature is the sharing of results between methods (at a single agent); we would indicate this by the presence of a sharing relationship (similar to *facilitates*) between each pair of normal and level-hopping VLMs (note: this is not indicated in the figure). Sharing of results could be only one-way between methods.

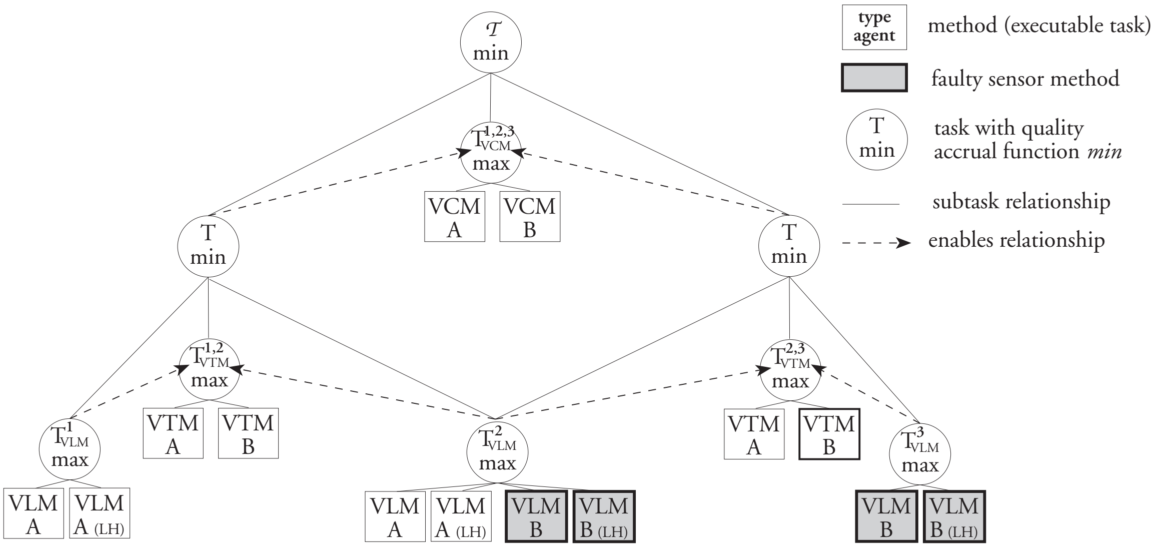

Figure 2: Objective task structure associated with two agents

Now we will add two final features that make this model more like the DVMT. First, low quality results tend to make things harder to process at higher levels. For example, the impact of using the level-hopping VLM is not just that its quality is lower, but also that it affects the quality and duration of the VTM it enables (because not enough possible solutions are eliminated). To model this, we will use the *precedence* relationship instead of *enables*: not only do the VLM methods enable the VTM, but they can also hinder its execution if the enabling results are of low quality. Secondly, the first VLM execution provides information that slightly shortens the executions of other VLMs in the same vehicle track (because the sensors have been properly configured with the correct signal processing algorithm parameters with which to sense that particular vehicle). A similar *facilitation* effect occurs at the tracking level. These effects occur both locally and when results are shared between agents—in fact, this effect is very important in motivating the agent behavior where one agent sends preliminary results to another agent with bad sensor data to help the receiving agent in disambiguating that data. Figure 3 repeats the objective task structure from the previous figure, but omits the methods for clarity. Two new tasks have been added to model facilitation at the vehicle location and vehicle track level.[4] T_{VL} indicates the highest quality initial work that has been done at the vehicle level, and thus uses the quality accrual function *maximum*. T_{VT} indicates the progress on the full track; it uses *summation* as its quality accrual function. The more tracking methods are executed, the easier the remaining ones become. The implications of this model are that in a multi-agent episode, then, the question becomes when to communicate partial results to another agent: the later an agent delays communication, the more the potential impact on the other agent, but the more the other agent must delay. We examined this question somewhat in [11].

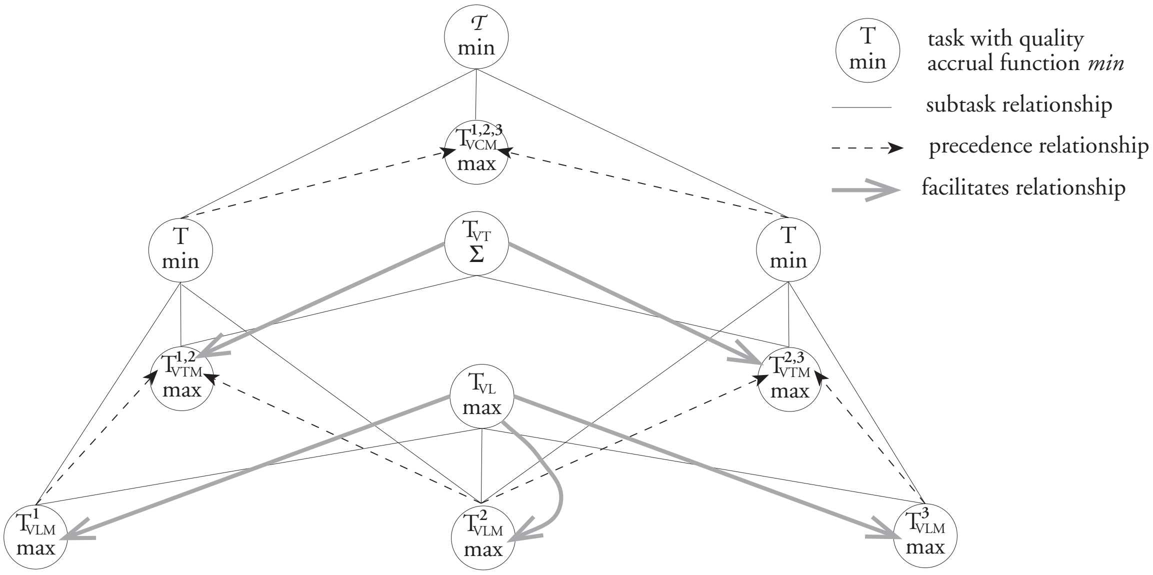

Figure 3: Non-local effects in the objective task structure

At the end of the next section, we will return to this example and add to it subjective features: what information is available to agents, when, and at what cost.

4 Resources

We have recently extended TÆMS so that resources can be represented directly in task structures. From the point of view of task environment specification, the only effect of resources, with respect to computational tasks, is to change the duration or quality of other tasks, and so we represent this as a set of non-local effects (NLE's). Because the NLEs are linked (i.e., not just between two nodes) we add a

[4] Note that these tasks were added to make the model more expressive; they are not associated with new methods.

resource node to the task structure and associate any state information needed with it. From the point of view of the state-based mathematics behind TÆMS, a *resource* looks like a task—but instead of a quality vector, a nominal state vector is used, and there is no associated duration or subtask relationships. This idea is very similar to the way we added tasks to represent the amount of work done on the first vehicle location and the vehicle track in Figures 2 and 3 in Section 2.1.

Instead, methods are related to the resource via one kind of NLE (e.g. *uses, replenishes, consumes, reads, writes*), and another NLE runs from the resource to the methods (e.g. *exclusive access, limited bandwidth, consumable*). The method-to-resource NLE's change the state of the resource, and the resource-to-method NLE's affect duration and max quality as usual.

For example, take the situation of a low-bandwidth communication link as described in [33]. Two agents have a diagnosis method that uses the low-bandwidth link. When more than one of these diagnosis methods are executed at the same time, the link is saturated and the durations of the methods are lengthened. We represent this situation as shown in Figure 4; the two new NLE's are defined as follows:

$$\textbf{uses}(T_a, R, t, d, q, \alpha) = \begin{cases} [d, q+\alpha] & \mathsf{S}(T_a) < t < \mathsf{F}(T_a) \\ [d, q] & \text{otherwise} \end{cases} \tag{4}$$

$$\textbf{limits}(R, M, t, d, q, \sigma, \phi_{\mathrm{d}}) = \begin{cases} [\phi_{\mathrm{d}} d, q] & Q(R,t) > \sigma \\ [d, q] & \text{otherwise} \end{cases} \tag{5}$$

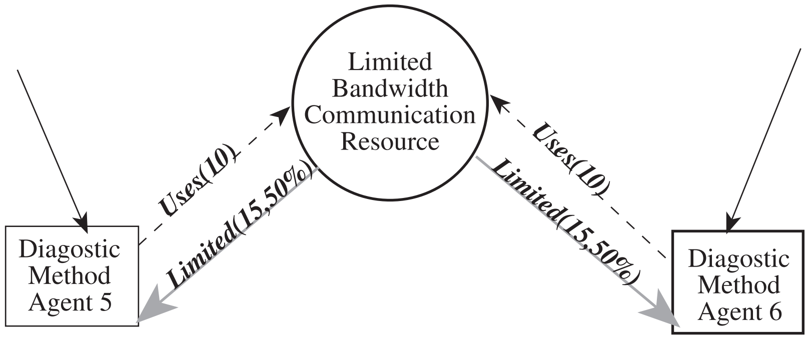

Figure 4: Example of two methods sharing a limited resource of capacity 15. In this example, if both methods execute in temporally overlapping time periods, the durations of each method will be lengthened by 50%.

The NLE *uses* indicates the amount of limited resource capacity used by the method with the parameter α. The NLE *limits* indicates that the resource has maximum capacity (saturation) σ and that beyond that point duration is affected by percentage ϕ_{d} (alternately, ϕ_{d} could be defined as a function of the level of oversaturation). Similar NLE's can be defined for representing consumable resources and their associated operations. Note that when methods can be interrupted, the lengthening of method durations due to a blocked resource is not associated with lengthening the amount of continuous computation—the blocked method can be interrupted and other computations performed until the method is no longer blocked.

5 Subjective Level

The purpose of a subjective level model of an environment is to describe what portions of the objective model of the situation are available to 'agents'. It answers questions such as "when is a piece of information available," "to whom is it available," and "what is the cost to the agent of that piece of information". This is a description of how agents might interact with their environment—what options are available to them.

To build such a description we must introduce the concept of *agency* into the model. Ours is one of the few comprehensive descriptions of computational task environments, but there are many formal and informal descriptions of the concept of agency (see [21, 22, 27]). Rather than add our own description, we notice that these formulations define the notion of *computation* at one or more agents, not the environment that the agents are part of. Most formulations contain a notion of *belief* that can be applied to our concept of "what an agent believes about its environment". Our view is that an "agent" is a locus of belief and action (such as computation).

The form of the rest of this section is as follows: how does the environment affect the beliefs of the agents; how do the beliefs of agents affect their actions, and how do the actions affect the environment.

5.1 Agent Beliefs

We use the symbol Γ_A^t to denote the set of beliefs of agent A at time t. A subjective mapping of an objective problem solving situation $\mathcal{O}$ is a function $\varphi : [A \times \mathcal{O}] \mapsto \Gamma_A$ from an agent and objective assertions to the beliefs of an agent. For example, we could define a mapping φ where each agent has a probability p of believing that the maximum quality of a method is the objective value, and a probability $1 - p$ of believing the maximum quality is twice the objective value. Any objective assertion has some subjective mapping, including **q** (maximum quality of a method), **d** (duration of a method), deadlines, and the relations SUBTASK, NLE, and COMM.

5.2 Control

The beliefs of an agent affect its actions through some control mechanism. Since this is the focus of most of our and others' research on local scheduling, coordination, and other control issues, we will not discuss this further (but see Section 6.2). The agent's control mechanism uses the agent's current set of beliefs Γ_A to update three special subsets of these beliefs (alternatively, *commitments*[30]) identified as the sets of information gathering, communication, and method execution actions to be computed.

The models we build typically further divide control into *local scheduling* and *coordination* (see [11]), but this is not required. Besides describing how an agent's beliefs entail commitments to particular information gathering, communication, and method execution actions, a control component model must also describe the *duration* of its deliberations. This feature allows us to analyze questions concerning the cost of control without becoming mired in implementation details.[5]

5.3 Computation

Our model can support parallel computation, but for brevity we will just describe single processor computation as a sequence of agent states. Agent A's current state is uniquely specified by Γ_A. We provide a meta-structure for the agent's state-transition function that is divided into the following 4 parts: control, information gathering, communication, and method execution. First the control mechanisms assert (commit to) information-gathering, communication, and method execution actions and then these actions are computed one at a time, after which the cycle of meta-states repeats.

A simple model of parallel computation, similar to the implementation in [8], is to allow control, information gathering, and communication to run on one (abstract) processor, and multiple method executions on the other processors. Any important interactions between methods executing in parallel would be represented by non-local effects.

[5]Understanding the details of the control costs of particular algorithm implementations is important, but usually not at early stages of research. Detailed information about control costs can be used by TÆMS if it is available.

Method Execution. How do the actions of an agent affect the environment? Both the objective environment (e.g. quality of the executing method) and the subjective mapping (e.g. information available via φ) can be affected. We use two execution models: simple method execution, and execution with monitoring, suspension, and preemption. These follow from the discussion of Q_{DTT} and Q_{lin} in Section 3.2, and are simple state-based models. Basically, for non-interruptible, single processor method executions, the agent enters a method execution state for method M at time $\mathsf{S}(M)$ and remains in that state until the time t when $t - \mathsf{S}(M) = \mathbf{d}(M, t)$. Method execution actions are similar to what Shoham terms 'private actions' like DO[30].

We have also considered pre-emptable method execution, where a method execution action is given a set upper time limit, after which computation will proceed to the next meta-state. The agent can then monitor the execution of long methods, and interleave their execution with other actions or pre-empt them entirely [20].

Communication. How do the actions of an agent affect other agents? Communication actions allow agents to affect each others' beliefs to a limited extent. Many people have worked on formalizing aspects of communication; the semantics of communication actions can be freely defined for each environment. The simplest communication act is to send another agent the 'current result/value' of a method—the effect is to change the available quality $Q_{\text{avail}}(T, t, A)$ at the remote agent after the message has been received. What happens when a communication is 'received'? The reception of information, by changing the available quality of a task, may trigger a non-local effect as we described earlier, and may influence the behavior of an agent as specified by its control algorithm.

Information Gathering. An information gathering action trades-off computational resources (time that could be spent executing methods) for information about the environment. For example, one useful information gathering action is one that queries the environment about the arrival of new tasks or task groups. Another information gathering action causes any communications that have arrived at an agent to be 'received' (added to the agent's set of beliefs). A third kind of information gathering may identify coordination relationships—non local effects that span multiple agents. Both communication and information gathering actions take some period of time (not necessarily constant) to execute, as specified in the model.

5.4 Subjective Modeling Example

Let's return to the example we began in Section 2.1 to demonstrate how adding a subjective level to the model allows us to represent the effects of faulty sensors in the DVMT. We will define the default subjective mapping to simply return the objective value, i.e., agents will believe the true objective quality and duration of methods and their local and non-local effects. We then alter this default for the case of faulty (i.e., noisy) sensors—an agent with a faulty sensor will not initially realize it ($\mathbf{d}_0(\text{faulty-VLM}) = 2\mathbf{d}_0(\text{VLM})$, but $\varphi(A, \mathbf{d}_0(\text{faulty-VLM})) = \mathbf{d}_0(\text{VLM})$).[6] Other subjective level artifacts that are seen in [16] and other DVMT work can also be modeled easily in our framework. For example, 'noise' can be viewed as VLM methods that are subjectively believed to have a non-zero maximum quality ($\varphi(A, \mathbf{q}_0(\text{noise-VLM})) > 0$) but in fact have 0 objective maximum quality, which the agent does not discover until after the method is executed. The strength with which initial data is sensed can be modeled by lowering the subjectively perceived value of the maximum quality $\mathbf{q}$ for weakly sensed data. The infamous 'ghost track' is a subjectively complete task group appearing to an agent as an actual vehicle track, which *subjectively* accrues quality until the hapless agent executes the VCM method, at which point the true (zero) quality becomes known. If the track (subjectively) spans multiple agents' sensor regions, the agent can potentially identify the chimeric track through

[6] At this point, one should be imagining an agent controller for this environment that notices when a VLM method takes unusually long, and realizes that the sensor is faulty and replans accordingly.

communication with the other agents, which may have no belief in such a track (but sometimes more than one agent suffers the same delusion). In general, having different subjective mappings for different agents or classes of agents allows us to model situations where some agents are more, less, or simply differently 'informed' than others.

Next we turn to the control of the agents in an environment. As an example, we will now present a very simple static control algorithm that uses no meta-level communication (more information about this and other algorithms can be found in [12]). In a static organization, agents divide their overlapping areas of responsibility as evenly as possible, resulting in new areas of responsibility for each agent with no overlap (to avoid redundant processing). Given a subjectively believed task structure as described in Section 2.1 and any communicated task results received by information gathering actions, the agent can at any time build a list of currently executable methods (under the set of precedence constraints). At any time an agent can also build a list of methods that need to be executed, but cannot be 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 agent has worked on. Each agent follows the same control algorithm, 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(method))
        Until (null E))
        Do Communication-Action(broadcast highest-level results)
        Let W = [get set of methods still waiting on precedence constraints]
Until (null W))
```

It would be useful to know the performance of a system using this coordination algorithm. We could use the simulation capability of TÆMS, but in this case a more powerful analytic result can be derived. If we let S' represent the largest amount of low-level data in one task group seen by any agent, and a the total number of agents that see the task group, 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 ($\mathbf{d}_0(I)$) plus the amount of time to do all the local work ($S'\mathbf{d}_0(\mathrm{VLM}) + (S'-1)\mathbf{d}_0(\mathrm{VTM})$), communicate that work ($\mathbf{d}_0(C)$), get the other agents' results ($\mathbf{d}_0(I)$), plus the amount of time to combine results from the other $a-1$ agents ($(a-1)\mathbf{d}_0(\mathrm{VTM})$), plus time to produce the final complete task group result ($\mathbf{d}_0(\mathrm{VCM})$), plus communicate that result to everyone ($\mathbf{d}_0(C)$).

We have explained the objective and subjective levels of our modeling framework, and presented an example of a moderately complex task structure and a simple control algorithm for an agent that can accomplish this task. Next we turn to the generative level, where we specify the statistical properties of an environment across many episodes.

6 Generative Level

By using the objective and subjective levels of TÆMS we can model any *individual* situation; adding a *generative* level to the model allows us to go beyond that and determine what the expected performance of an algorithm is over a long period of time and many individual problem solving episodes. Our previous work has created generative models of task interarrival times (exponential distribution), amount of work in a task group (Poisson), task durations (exponential), and the likelihood of a particular non-local effect between two tasks [12, 11, 20]. Generative level statistical parameters can also be used by agents in their subjective reasoning, for example, an agent may make control decisions based on the knowledge of the expected duration of methods.

A generative level model can be constructed by careful analysis of the real environment being modeled, or by observing the statistical properties of real episodes (if that is possible). Even when certain parameters of the real world are unknown, they can be made variables in the model and then you can ask questions about how much they affect the things you care about. Our approach so far has been to verify our assumptions about the environment with simple statistical approaches [23]. Detailed model verification will be more important when using our framework to optimize parameters in a real application, as opposed to learning the general effects of parameters on a coordination or negotiation algorithm (see Section 7).

In our DSN model example, 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 analyses are based on the statistical properties of episodes in this environment, not any single instance of an episode. The properties of the episodes in a DSN environment are summarized by the tuple $\mathcal{D} =< A, \eta, r, o, \mathcal{T} >$ where A specifies the number of agents, η 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[7], and $\mathcal{T}$ specifies the homogeneous task group structure (as discussed in Sections 2.1 and 5.4). A particular episode in this environment can be described by the tuple $D =< A, r, o, \mathcal{T}_1, \ldots, \mathcal{T}_n >$ where n is a random variable drawn from a Poisson distribution with an expected value of η. If we were to extend this generative model to cover every feature we added to the objective and subjective models in Sections 2.1 and 5.4, we would need to add the likelihood of a sensor being faulty (noisy), the likelihood of a ghost track, etc.

6.1 Analysis Summary

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 [18]. 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 [27] focuses on the *relative efficiencies* of various organizations given an uncertain environment, while the updated versions of the *contingency theory* approach [32] focus on the need for an organization to expand toward the earliest available *information that resolves uncertainties* in the current environment.

We have used these general concepts, in conjunction with our modeling framework, to analyze potential organizational structures for the environment that has been our example in this paper, i.e., the class of naturally distributed, homogeneous, cooperative problem solving environments where tasks arrive at multiple locations, exemplified by distributed sensor networks. Uncertainty can occur both at the level of episodic variation in a single environment, and in the differences between the objective task structure of a single episode and the subjective view of each agent. Our approach was to first construct the task environment model we have been using as an example in this paper, and then to develop expressions for the *expected efficiencies* of static and dynamic organizational structures, in terms of performance—the cost of communication and time to complete a given set of tasks. Finally, we validated these mathematical models by using the simulation capabilities of TÆMS.

We briefly showed at the end of Section 5.4 how we can predict the performance of a system given the objective and subjective models of a particular episode. This is very useful for explaining or predicting

[7] We also assume the agents start in a square geometry, i.e, 4 agents in a 2×2 square, 25 agents arranged 5×5.

agent behavior in a particular episode or scenario, but not over many episodes in a real environment. To do this, we build probabilistic models of the relevant objective and subjective parameters (now viewed as random variables) that are based on generative level parameters. Another paper, [12], details this process, and shows how the distributions of objective parameters such as "the number of VLM methods seen by the maximally loaded agent" ($\hat{S}$) and "the max number of task groups seen by the same agent" ($\hat{N}$) can be defined from just the generative parameters $\mathcal{D} =< A, \eta, r, o, \mathcal{T} >$.

For example, the total time until termination for an agent receiving an initial data set of size $\hat{S}$ is the time to do local work, combine results from other agents, and build the completed results, plus two communication and information gathering actions, as was discussed in Section 5.4:

$$\hat{S}\mathbf{d}_0(\text{VLM}) + (\hat{S} - \hat{N})\mathbf{d}_0(\text{VTM}) + (a-1)\hat{N}\mathbf{d}_0(\text{VTM}) + \hat{N}\mathbf{d}_0(\text{VCM}) + 2\mathbf{d}_0(I) + 2\mathbf{d}_0(C) \tag{6}$$

We can use Eq. 6 as a predictor by combining it with the probabilities for the values of $\hat{S}$ and $\hat{N}$. Again, we refer the interested reader to [12] for derivations, verification, and applications of these results. Note that if the assumptions behind our generative model change (for instance, if we assume all agents initially line up side-by-side, instead of in a square, or if vehicles made a loop before exiting the sensed area) the probability distributions for $\hat{S}$ and $\hat{N}$ might change, but that the form of Eqn. 6 does not. If the agent's coordination algorithm changes, then Eqn. 6 will change (see [12]).

The coordination algorithm discussed in Section 5.4 results in a static organization; 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. 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. In [12] we present equations similar to Eq. 6 that show the potential benefits of dynamic reorganization in an arbitrary environment $\mathcal{D}$, and discuss when the overhead of meta-level communication is worthwhile.

6.2 Simulation

Simulation is a useful tool for learning parameters to control algorithms, for quickly exploring the behavior space of a new control algorithm, and for conducting controlled, repeatable experiments when direct mathematical analysis is unwarranted or too complex. The simulation system we have built for the direct execution of models in the TÆMS framework supports, for example, the collection of paired response data, where different or ablated coordination or local scheduling algorithms can be compared on identical instances of a wide variety of situations (generated using the generative level of the model). We have used simulation to explore the effect of exploiting the presence of *facilitation* between tasks in a multi-agent real-time environment where no quality is accrued after a task's deadline [11]. The environmental generative characteristics here included the mean interarrival time for tasks, the likelihood of one task facilitating another, and the strength of the facilitation (ϕ_{d}).

The TÆMS framework is not limited to experimentation in distributed problem solving. In [20], Garvey and Lesser used the framework to describe the effects of various task environment and agent design features on the performance of their real-time 'design-to-time' algorithm. They manipulate the objective-to-subjective mapping to examine questions about monitoring the execution of tasks when true method durations are not known. They show that monitoring does provide a reduction in missed deadlines but that this reduction may be significant only during 'medium' loads. Garvey is now using a more complex model of *enabling* and *hindering* task structures to design an optimal design-to-time algorithm for certain task environments.

In contrast to the previous analytical example, in this section we will consider the use of TÆMS as a simulator to explore hypotheses about the interactions between environmental and agent-structural characteristics. We return to a question explored by Burton and Obel: is there a significant difference in performance due to either the choice of organizational structure or the decomposability of technology? *Technology* is used here in the management science sense of "the physical method by which resources are converted into products or services" or a "means for doing work" [3, 29].

We equate a technology with a task structure, instead of a linear program. Task structures allow us to use a clear interval measure for decomposability, namely the probability of a task interrelationship (in this example enables, facilitates, and overlaps). We define a nearly decomposable task structure to have a base probability of 0.2 for these three coordination relationships and a less decomposable task structure to have a base probability of 0.8 (see Figure 5). We will continue in this example to look at purely computational task structures, although the use of physical resources can also be represented (see Section 4).

Burton and Obel were exploring the difference in M-form and U-form hierarchical structures; we have recently been analyzing a team-oriented, non-hierarchical coordination algorithm for computational agents called the Generalized Partial Global Planning (GPGP) algorithm. Without going into too much detail, the GPGP algorithm specifies three basic areas of the agent's coordination behavior: how and when to communicate and construct non-local views of the current problem solving situation; how and when to exchange the partial results of problem solving; how and when to make and break *commitments* to other agents about what results will be available and when. Each agent also has a heuristic local scheduler that decides what the agent should do next, based on its current view of the problem solving situation (including the commitments it has made), and a utility function.

For our structural variable we will vary the communication of non-local views. Informally, we will be contrasting the situation where each agent makes commitments and communicates results based only on local information (*no* non-local view) with one where the agents freely share task structure information with one another across coordination relationships (*partial* non-local view). Figure 6 shows an example—note that in neither case does the agent have the *global* view of Figure 5.

Burton and Obel used a profitability index as their performance measure, derived from the percentage of optimal profit achieved. In general, the scheduling an arbitrary TÆMS task structure is an NP-hard problem and so we do not have access to optimal solutions. Instead we compare performance directly on four scales: the number of communication actions, the amount of time spent executing methods, the final quality achieved, and the termination time. Simulation runs for each of the four combinations of non-local view policy and level of task decomposability were done in matched sets—the randomly generated episode was the same for each combination with the exception of more coordination relationships being added in the less decomposable task structures. Following Burton and Obel, we used the non-parametric Friedman two-way analysis of variance by ranks test for our hypotheses. The assumptions of this test are that each block (in our case, randomly generated episode) is independent of the others and that each block contains matched observations that may be ranked with respect to one another. The null hypothesis is that the populations within each block are identical.

We generated 40 random episodes of a single task group, each episode was replicated for the four combinations in each block. We used teams consisting of 5 agents; the other parameters used in generating the task structures are summarized in Table 1 and a typical randomly generated structure is shown in Figure 5. Figure 6 shows the difference in the local view of one agent with and without creating partial non-local views. We first tested two major hypotheses:

Hypothesis 1: *There is no difference in performance between agents with a partial non-local view and those without.* For the communication and method execution performance measures, we reject the null hypothesis at the 0.001 level. We cannot reject the null hypothesis that there is no difference

Parameter	Value
Mean Branching factor (Poisson)	1
Mean Depth (Poisson)	3
Mean Duration (exponential)	10
Redundant Method QAF	Max
Number of task groups	1
Task QAF distribution	(50% min) (50% max)
Decomposition parameter	p = 0.2 or 0.8
Hard CR distribution	(p enables) ((1-p) none)
Soft CR distribution	(p facilitates) (10% hinders) ((.9-p) none)
Chance of overlaps (binomial)	p

Table 1: Parameters used to generate the 40 random episodes

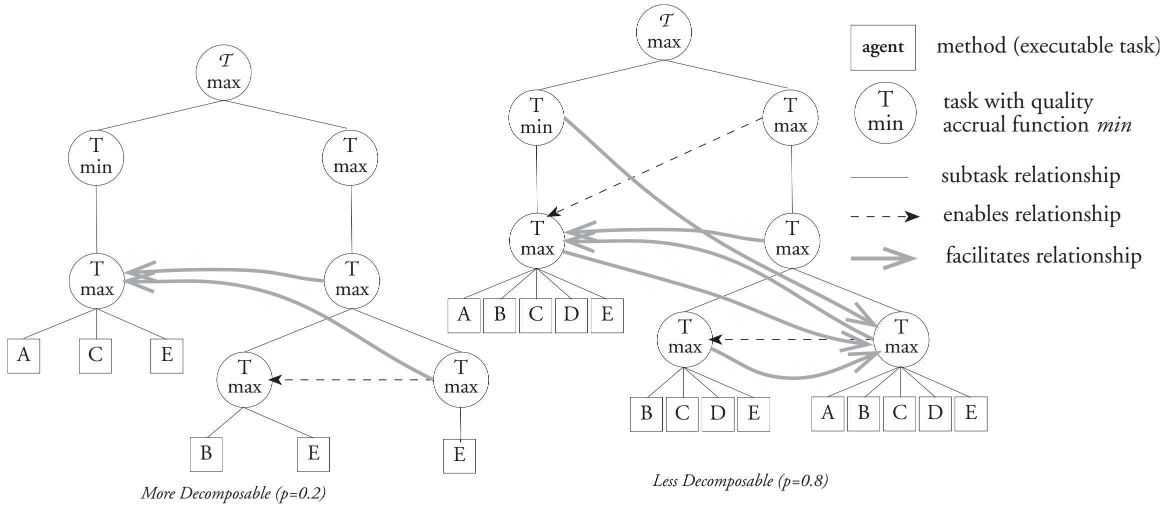

Figure 5: Example of a randomly generated objective task structure, generated with the parameters in the previous table.

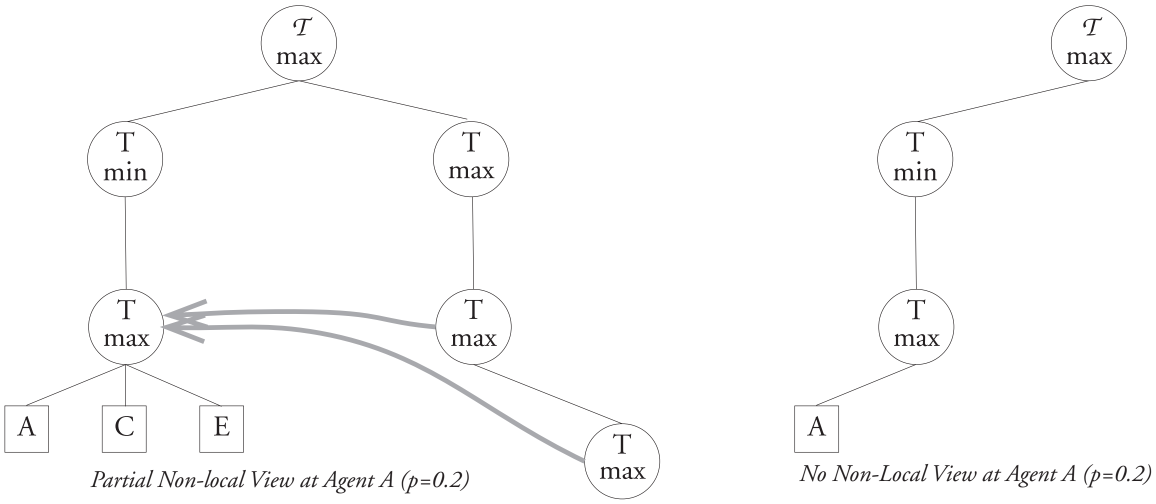

Figure 6: Example of the local view at Agent A when the team shares private information to create a partial non-local view and when it does not.

in final quality and termination time. Teams of computational agents that exchange information about their private, local views consistently exchange more messages (in this experiment, a mean increase of 7 messages) but do less work (here, a mean decrease of 20 time units).

Hypothesis 2: *There is no difference in performance due to the level of decomposability of technology.* For the communication and method execution performance measures, we reject the null hypothesis at the 0.001 level. We cannot reject the null hypothesis that there is no difference in final quality and termination time. Teams of computational agents, regardless of their policy on the exchange of private, local information communicate more messages (in this experiment, a mean increase of 47 messages) and do more work (here, a mean increase of 24 time units) when faced with less decomposable computational task structures (technology).

Again following Burton and Obel, we next test for interaction effects between non-local view policy and level of technology decomposability by calculating the differences in performance at each level of decomposability, and then testing across non-local view policy. This test was significant at the 0.05 level for communication, meaning that the difference in the amount of communication under the two policies is itself different depending on whether task decomposability is high or low. This difference is small (two communication actions in this experiment) and was not verified in a second independent test.

7 Conclusions

This paper has presented TÆMS, a framework for modeling computationally intensive task environments. TÆMS exists as both a language for stating general hypotheses or theories and as a system for simulation. The important features of TÆMS include its layered description of environments (*objective* reality, *subjective* mapping to agent beliefs, *generative* description of the other levels across single instances); its acceptance of any performance criteria (based on temporal location and *quality* of task executions); and its non-agent-centered point of view that can be used by researchers working in either formal systems of mental-state-induced behavior or experimental methodologies. TÆMS provides environmental and behavioral structures and features with which to state and test theories about the control of agents in complex computational domains, such as how decisions made in scheduling one task will affect the utility and performance characteristics of other tasks.

TÆMS is not only a mathematical framework, but also a simulation language for executing and experimenting with models directly. The TÆMS simulator supports the graphical display of generated subjective and objective task structures, agent actions, and statistical data collection in CLOS on the TI Explorer and DEC Alpha. These features help in both the model-building stage and the verification stage. The TÆMS simulator is being used not only for research into the coordination of distributed problem solvers[12, 11, 10], but also for research into real-time scheduling of a single agent[20], scheduling at an agent with parallel processing resources available, and soon, learning coordination algorithm parameters.

TÆMS does not at this time automatically learn models or automatically verify them. While we have taken initial steps at designing a methodology for verification (see [12]), this is still an open area of research [6]. Our future work will include building new models of different environments that include physical resource constraints, such as airport resource scheduling. Other extensions we envision include specifying dynamic objective models that change structure as the result of agent actions—such models add yet another important source of uncertainty that can be a factor in influencing agent organization. We also wish to expand our analyses beyond the questions of scheduling and coordination to questions about negotiation strategies, emergent agent/society behavior, and organizational self-design.

Acknowledgments

The authors would like to thank Marty Humphrey, Alan Garvey, and the reviewers for their helpful comments and suggestions.

References

[1] Mark Boddy and Thomas Dean. Solving time-dependent planning problems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, August 1989.

[2] Piero P. Bonissone and Keith S. Decker. Selecting uncertainty calculi and granularity: An experiment in trading-off precision and complexity. In L. N. Karnak and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence*. North Holland, 1986.

[3] Richard M. Burton and Børge Obel. *Designing Efficient Organizations: Modelling and Experimentation*. North Holland, Amsterdam, 1984.

[4] Norman Carver and Victor Lesser. A new framework for sensor interpretation: Planning to resolve sources of uncertainty. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 724–731, August 1991.

[5] Paul Cohen, Michael Greenberg, David Hart, and Adele Howe. Trial by fire: Understanding the design requirements for agents in complex environments. *AI Magazine*, 10(3):33–48, Fall 1989. Also COINS-TR-89-61.

[6] Paul R. Cohen. A survey of the Eighth National Conference on Artificial Intelligence: Pulling together or pulling apart? *AI Magazine*, 12(1):16–41, Spring 1991.

[7] Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(3), 1990.

[8] Keith S. Decker, Alan J. Garvey, Marty A. Humphrey, and Victor R. Lesser. Control heuristics for scheduling in a parallel blackboard system. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(2):243–264, 1993.

[9] Keith S. Decker, Alan J. Garvey, Marty A. Humphrey, and Victor R. Lesser. A real-time control architecture for an approximate processing blackboard system. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(2):265–284, 1993.

[10] Keith S. Decker and Victor R. Lesser. Generalizing the partial global planning algorithm. *International Journal of Intelligent and Cooperative Information Systems*, 1(2):319–346, June 1992.

[11] Keith S. Decker and Victor R. Lesser. Analyzing a quantitative coordination relationship. *Group Decision and Negotiation*, 2(3):195–217, 1993.

[12] Keith S. Decker and Victor R. Lesser. An approach to analyzing the need for meta-level communication. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, Chambéry, France, August 1993.

[13] Keith S. Decker and Victor R. Lesser. A one-shot dynamic coordination algorithm for distributed sensor networks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 210–216, Washington, July 1993.

[14] Keith S. Decker, Victor R. Lesser, and Robert C. Whitehair. Extending a blackboard architecture for approximate processing. *The Journal of Real-Time Systems*, 2(1/2):47–79, 1990.

[15] D. Dubois and H. Prade. Criteria aggregation and ranking of alternatives in the framework of fuzzy set theory. In H.J. Zimmermen, L. A. Zadeh, and B.R. Gains, editors, *TIMS/Studies in Management Science*, volume 20, pages 209–240. Elsevier Science Publishers, 1984.

[16] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill. Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, 36(11):1275–1291, November 1987.

[17] E.H. Durfee and V.R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5):1167–1183, September 1991.

[18] J. Galbraith. *Organizational Design*. Addison-Wesley, Reading, MA, 1977.

[19] Alan Garvey, Marty Humphrey, and Victor Lesser. Task interdependencies in design-to-time real-time scheduling. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 580–585, Washington, D.C., July 1993.

[20] Alan Garvey and Victor Lesser. Design-to-time real-time scheduling. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(6), 1993. Special Issue on Scheduling, Planning, and Control.

[21] Les Gasser. Social conceptions of knowledge and action. *Artificial Intelligence*, 47(1):107–138, 1991.

[22] Carl Hewitt. Open information systems semantics for distributed artificial intelligence. *Artificial Intelligence*, 47(1):79–106, 1991.

[23] Jack P. C. Kleijnen. *Statistical Tools for Simulation Practitioners*. Marcel Dekker, New York, 1987.

[24] Paul Lawrence and Jay Lorsch. *Organization and Environment*. Harvard University Press, Cambridge, MA, 1967.

[25] J. W. S. Liu, K. J. Lin, W. K. Shih, A. C. Yu, J. Y. Chung, and W. Zhao. Algorithms for scheduling imprecise computations. *IEEE Computer*, 24(5):58–68, May 1991.

[26] Thomas W. Malone. Modeling coordination in organizations and markets. *Management Science*, 33:1317–1332, 1987.

[27] Terry M. Moe. The new economics of organization. *American Journal of Political Science*, 28(4):739–777, November 1984.

[28] Martha E. Pollack and Marc Ringuette. Introducing Tileworld: Experimentally evaluating agent architectures. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 183–189, July 1990.

[29] W. Richard Scott. *Organizations: Rational, Natural, and Open Systems*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1987.

[30] Yoav Shoham. AGENT0: A simple agent language and its interpreter. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 704–709, Anaheim, July 1991.

[31] Arthur L. Stinchcombe. *Constructing Social Theories*. University of Chicago Press, Chicago, 1987.

[32] Arthur L. Stinchcombe. *Information and Organizations*. University of California Press, Berkeley, CA, 1990.

[33] Toshiharu Sugawara and Victor R. Lesser. On-line learning of coordination plans. Computer Science Technical Report 93–27, University of Massachusetts, 1993.

[34] James D. Thompson. *Organizations in Action*. McGraw-Hill, New York, 1967.

[35] G. Zlotkin and J. S. Rosenschein. Incomplete information and deception in multi-agent negotiation. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 225–231, Sydney, Australia, August 1991.