

Evolution of the GPGP/TÆMS Domain-Independent Coordination Framework^{1,2}

CMPSCI 02-03

V. Lesser, K. Decker³, T. Wagner⁴, N. Carver⁵, A. Garvey⁶, B. Horling, D. Neiman⁷, R. Podorozhny, M. Nagendra Prasad⁸, A. Raja, R. Vincent⁹, P. Xuan and XQ. Zhang

Computer Science Department
University of Massachusetts, Amherst

January 11, 2002

ABSTRACT:

The GPGP/TÆMS domain-independent coordination framework for small agent groups was first described almost ten years ago and then more fully detailed in an ICMAS95 paper. In this paper, we discuss the evolution of this framework over the last six years motivated by its use in a number of applications, including: information gathering and management, intelligent home automation, coordination of concurrent engineering activities, distributed situation assessment, and hospital scheduling. First, we review the basic architecture of GPGP and then present extensions to the TÆMS domain-independent representation of agent activities. We next describe extensions to GPGP that permit the representation of situation-specific coordination strategies and social laws as well as making possible the use of GPGP in large agent organizations. Additionally, we discuss a more complex view of commitments that takes into account uncertainty in commitments. We then present new coordination mechanisms for use in resource sharing and contracting, and more complex coordination mechanisms that use a cooperative search among agents to find appropriate commitments. We conclude with a discussion of future research directions.

¹ This material is based upon work that has been sponsored by the National Science Foundation (under Grant No. IIS-9812755 and Grant No. IIS-9988784) and the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF (under agreement number F30602-99-2-0525). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), Air Force Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon.

² This is a significant extension, including additional authors, of a paper with a similar title [38] written in 1998.

³ At Computer and Information Science Dept., University of Delaware.

⁴ At Honeywell Research Lab, Minneapolis, Minnesota

⁵ At Computer Science Dept., Southern Illinois University.

⁶ At Computer Science Dept., Truman State University.

⁷ At River Logic, Inc., Beverly, Massachusetts

⁸ At VerticalNet, Palo Alto, California.

⁹ At SRI International, Menlo Park, California

Introduction

Generalized Partial Global Planning (GPGP) and its associated TÆMS hierarchical task network representation [7, 9, 10, 11, 25] were developed as a domain-independent framework for coordinating the real-time activities of small teams of cooperative agents working to achieve a set of high-level goals. GPGP's development was influenced by two factors: one was to generalize and make domain-independent the coordination techniques developed in the Partial Global Planning (PGP) framework [16]¹⁰; the other was based on viewing agent coordination in terms of coordinating a distributed search of a dynamically evolving goal tree as pictured in Figure 1 [36]. Underlying these two influences was a desire to construct a model that could be used to explain and motivate the reasons for coordination among agents based on a quantitative view of task/subproblem dependency. As discussed in [11], coordination of behaviors among agents requires three things: specification (creating shared goals), planning (subdividing goals into subgoals/tasks, i.e., creating the substructure of the evolving goal tree) and scheduling (assigning tasks to individual agents or groups of agents, creating shared plans and schedules and allocating resources). GPGP is primarily concerned with scheduling activities rather than the dynamic specification and planning of evolving activities (e.g., such as decomposing a high-level goal into a set of subgoals that if successfully achieved will solve the high-level goals)¹¹.

The aim of the GPGP coordination framework is to maximize the overall utility accrued by the group of agents as a result of successful completion of its high-level goals. These goals can be independent or held jointly by two or more agents, can be time and resource sensitive, and can be of different utilities. GPGP deals with high-level coordination issues (scheduling and task selection) involving decisions for each agent about what subtasks¹² of which high-level goals it will execute, the level of resources it will expend in their execution, and when they will be executed. These decisions, if not appropriately made, can lead to the group of agents acting incoherently in achieving the desired set of high-level goals. This incoherence can be characterized by agents wasting resources by: 1) doing redundant, unnecessary or less important subtasks, 2) having nothing to do since the needed prerequisite information or external resources that are necessary for subtask execution are not available in a timely fashion, 3) being overloaded while other agents are underloaded and, finally, 4) missing deadlines on important tasks. The need for these types of scheduling decisions occurs when agents are solving subtasks that are interdependent, either through their contention for scarce resources or through relationships arising among them because they are contributing to the solution of the same higher-level goal. GPGP is thus most appropriate for soft real-time applications operating in dynamic and resource-bounded environments where there are complex interdependencies among activities occurring at different agents.

This line of research also assumes soft real-time multi-agent applications are best structured using groups of sophisticated agents – where the individual agents are responsible for effectively balancing the resources they choose to allocate to multiple, time and resource sensitive activities¹³.

¹⁰ This also involved our understanding that coordination activities could be separated from local agent control if an appropriate bi-directional interface could be established between them.

¹¹ Though as will be discussed, TÆMS provides a rich language for the representation of agent activities and alternative ways of achieving goals.

¹² A subtask can potentially represent a significant body of code or a significant set of activities.

¹³ In recent work on a distributed sensor network application [26], we exploited such a sophisticated agent architecture [58] to create a virtual agent organization of simple, single-threaded agents. These “virtual” agents were created as needed and dynamically assigned as goals to a specific sophisticated (“real”) agent based on gross information about current resource usage of agents and the type of resources available at each agent. In the hardware implementation of this architecture, there was one sophisticated agent for each processor node. The “real” agent then performed detailed planning/scheduling based on local resource availability and priority of these goals, and multiplexed among the different goals that it was concurrently executing in order to meet soft real-time requirements. The use of

An alternative approach to this multiplexing among multiple activities in a single agent (and consequently avoiding the requirement for more sophisticated local agent control) is to use agents that have simpler local control architectures that for instance only allow one goal or task to be performed at any given time. In this latter approach, the underlying operating system handles the multiplexing among the different “simple” agents’ activities when these agents are hosted on the same processor node. Further, the arrival of new tasks/goals requires instantiating the appropriate “simple” agent on a specific processor node to solve the desired task. However, if there are options for not doing goals or doing them in different ways, depending on resource availability, then the former approach will permit an agent to make the decisions based on an appropriate context, while in the latter approach the “simple” agents cannot make the decision unless they get detailed feedback from the underlying operating system about the current set of activities. Even in this case it may be hard because a “simple” agent does not have the appropriate context to understand how to balance its choice with the choice of other agents. We feel that, for many soft real-time multi-agent applications, a certain class of information is required to reason about how to balance resources between multiple tasks or goals that vie for limited resources (e.g. time). Decomposing the problem space completely to “simple” agents does not address the problem or remove the information and decision requirements.

The type of high-level coordination that we are discussing above is very different from the type of low-level, local and fine-grained coordination that commonly occurs when a protocol is created among agents to synchronize their activities. In frameworks such as COOL [2] and AGENTTALK [32] the emphasis is on the flow of messages and how the dialog between agents is structured. Such frameworks generally combine finite state machines with enhancements, e.g., exceptions, to support a flexible and explicit specification of a communication process. In contrast, in GPGP/TÆMS the focus is primarily analytical and domain independent. It is on the quantitative evaluation of the interactions among tasks and the formation of temporal constraints to resolve and to exploit the interactions. GPGP coordination is about exchanging the necessary information for the agents to analyze their activities and perform distributed optimization problem solving to produce better results. While GPGP does contain protocols that define the exchange of messages and the flow of the conversation, the emphasis is not on the machinery for specifying the protocols but on the information that must be exchanged, and when, in order to create a distributed optimization process whose performance approximates that of a centralized process. GPGP and work such as COOL relate in that GPGP could literally make use of COOL or AGENTTALK to structure the flow of information between agents -- GPGP would still be dealing with a higher-level analytical optimization problem. Aspects of the optimization problem considered by GPGP include deadlines and temporal constraints as well as resource interactions. As will be discussed in more detail later, the communication-based joint problem solving performed by GPGP also differs from the coordination that takes places in most BDI agent systems and frameworks such as TEAMCORE/STEAM [47, 56, 57]¹⁴, ARCHON [30]¹⁵ and SHARED-PLAN [24], in that GPGP is focused on optimization – how to concurrently solve multiple, time and resource sensitive goals of varying worth within existing resource constraints so as to maximize the overall utility accrued.

Over the last six years, we have significantly extended the framework based on our experience applying it to a number of applications including information gathering and management, intelligent home automation, distributed situation assessment, coordination of concurrent engineering activities,

sophisticated agents has also helped us to construct a resource negotiation protocol that operates on an abstract model of resources and does not need to resolve all conflicts to be successful [43]. In this case, the agent took on the responsibility of mapping the abstract resource allocation policy generated by the negotiation protocol into a detail resource allocation schedule and where possible, solved resource allocation conflicts at the abstract level through local shifting and modification of tasks.

¹⁴ TEAMCORE/STEAM allows for solving concurrent goals as long as there is no resource contention or interdependencies among the goals, i.e., they are totally independent.

¹⁵ ARCHON does allow for some very simple scheduling of tasks taking into account their time and worth.

and hospital scheduling [14, 26, 39, 40, 41, 45, 46]. Most of these extensions have been already discussed in published papers so the emphasis of this paper is not on providing an in-depth discussion of each extension and empirical results that examine the performance of the extension. Rather, our focus is to present a high-level overview of how this body of work fits together intellectually and how it has been used to create a framework that solves many of the issues necessary for building multi-agent applications that require sophisticated coordination among agents. We start out by first reviewing the basic concepts behind GPGP and TÆMS, and how the GPGP framework relates to other approaches.

Review of Basic Architecture

The GPGP approach to coordination assumes an operating environment consisting of cooperative agents that are pursuing a number of high-level goals with varying worth (importance) and time deadlines, each of whose achievement requires the (partial) satisfaction of a number of subgoals; these subgoals in turn may have their own time constraints imposed as a result of the time constraints on the higher-level goals that they are helping to achieve. This model of agent computation leads to a situation where each agent may be trying to achieve a number of subgoals, possibly associated with more than one high-level goal. We are also assuming that these subgoals may be achieved with varying levels of quality (satisfaction) which in turn will have an impact on the ultimate worth of the higher-level goal. We do not assume that every high-level goal needs to be achieved in order for the collection of agents as a whole to operate effectively¹⁶.

Underlying Conceptual Model

The underlying model of computation implicit in the GPGP approach to agent coordination can be conceptually represented as an extended *AND/OR* goal tree where the leaves of the tree are primitive (non-decomposable) activities. For each goal that the system is solving, there is a tree that represents alternative ways of solving this goal. For example, consider the goal tree for one high-level goal as depicted in Figure 1. In this representation, subgoals are not only related to each other by the *AND/OR* relationship among siblings, but also by constraints and other interdependencies among the goals. Examples of these constraints and interdependencies include: when the result of one subgoal is a precondition for another subgoal, when the result of one subgoal *if available* can improve the performance of another subgoal (for instance, by providing additional constraints), when two subgoal solutions need to be compatible (obey constraints on their solution) since they contribute to the solution of the same higher-level goal, when one subgoal produces a resource (or a state in the world) that is needed by another subgoal, and so forth.

A centralized view of the search control problem associated with solving this goal tree involves at each point choosing which subgoal to work on in order to produce the highest quality result in the available time. Now let us look at Figure 2, which is the distribution of this tree between two agents. Note the circles around certain goals and subgoals indicating that both agents are working on the same parts of the goal tree (jointly-shared goals). The question now arises as to what type of knowledge about the current state of the problem-solving search in *Agent 1* would be helpful to *Agent 2* in deciding which order to solve its subgoals, and how the decisions among the two agents should be coordinated. For example, should they simultaneously pursue alternative ways of achieving some subgoal (e.g., $G_{k,1}^{1,2}$ and $G_{k,2}^2$ solving for $G_k^{1,2}$) or should one agent delay executing its subgoal until it gets the results of another subgoal being solved by the other agent (e.g., $G_{k,2}^2$ waiting for the result from $G_{1,1}^1$). This scheduling problem becomes significantly more

¹⁶In the initial implementation of GPGP using the DTT Scheduler [19], there was a bias in local decision making towards achieving some minimal satisfaction of every high-level goal rather than maximizing the overall worth of the high-level goals achieved. In the more recent DTC Scheduler [60], the bias is not present and the scheduler works to maximize an overall utility function specified by the coordination module or the agent's domain problem solver.

complex if the agents are pursuing more than one high-level goal due to the large number of possible combinations. The computational complexity increases further when there are deadlines for the achievement of goals, there are alternative ways of achieving a goal that trade off the utility of the derived solution and the likelihood of successfully achieving a solution against resource usage, or the goals are only partially overlapping among agents. For example, *Agent 1* is pursuing a subtask associated with high-level goals 1 and 2, while *Agent 2* is pursuing a subtask associated with high-level goal 2 and goal 3. We can further complicate this problem by assuming that the goal tree is not *a priori* decomposed among the agents but is either dynamically assigned as the goal tree is being created (a top-down perspective) or the relationship among subgoals is not known *a priori*. In this latter case, the relationship among goals (e.g., G_0^1 and G_0^2 are the same goal) needs to be discovered through agent communication.

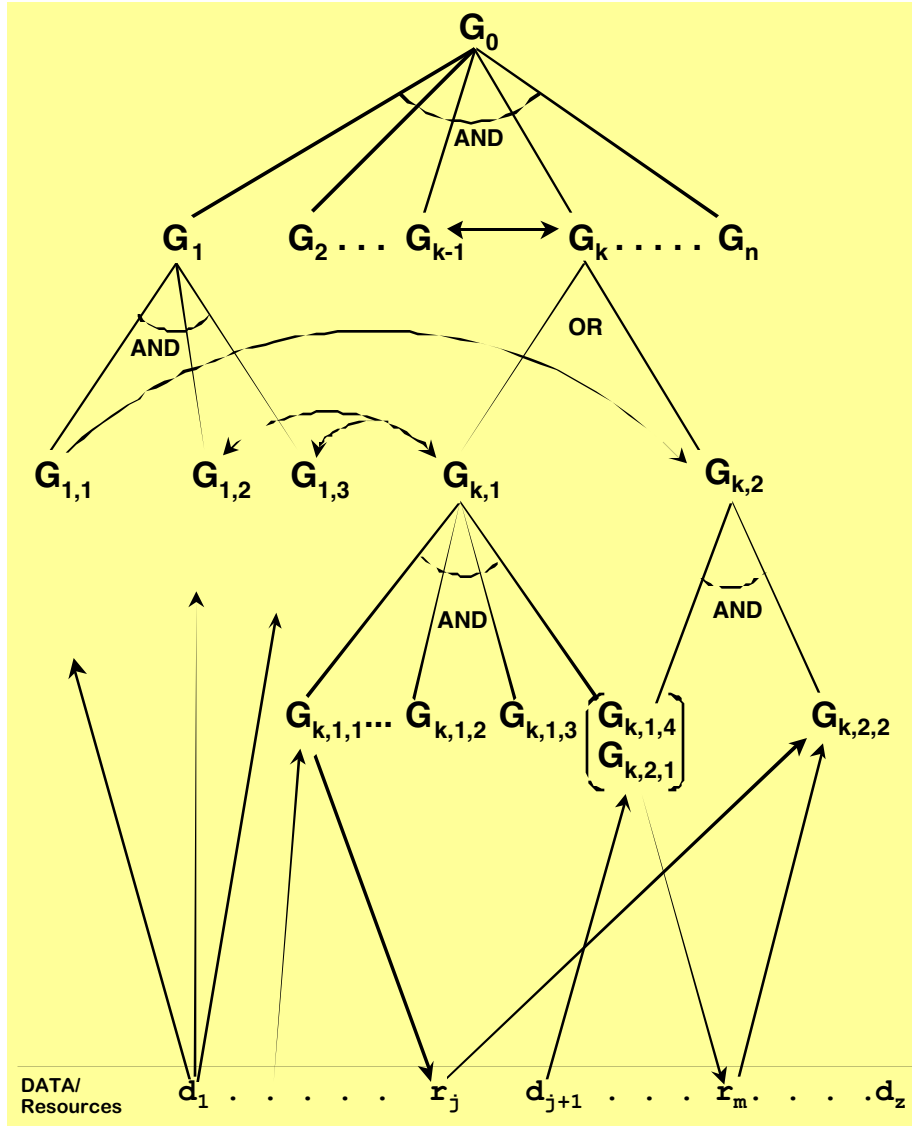


Figure 1: Represents an AND/OR goal tree that specifies a set of subgoals that need to be solved in order to solve the top-level goal. At the bottom of the figure are raw data and resources that are needed to solve specific subgoals. The figure also represents interdependencies among goals and the fact that one subgoal may contribute to solving more than one higher-level subgoal.

Our approach to this optimization problem of necessity needs to be distributed, quantitative and approximate because of limits and delays imposed by the agent communication medium, the uncertainty associated with agent problem-solving behavior, and the cost of decision making due to the combinatorics resultant from even a relatively small number of agents in the system and a small number of higher-level goals that are being pursued concurrently. For these reasons, GPGP is designed so that each agent optimizes its own local computational activities over optimization criteria defined by local considerations but mediated by both short-term and long-term knowledge of the activities of other agents. This mediation occurs as a result of GPGP imposing additional constraints on groups of these local optimization problems. These constraints connect local optimization problems that are not independent (i.e. agents who have dependencies among their activities) and represent a way of locally approximating the more global optimization problem involving the group of agents having interdependent activities.

There is a strong connection between the GPGP coordination module that is resident at each agent and a local scheduler/control module that is part of each agent's architecture. The connection between these two is necessary because the generation of suitable constraints must take into account characteristics of the current solution to the local agent optimization problem, e.g., when resources are available, when tasks will be performed, etc., but also the implications of change such as performing an unplanned task at a particular time. In our current work, the agent's local optimization expert is the Design-to-Criteria Scheduler (DTC) [60, 62, 66] or another module built on top of DTC [58], all of which are descended from the Design-to-Time (DTT) [19] scheduler¹⁷. The scheduler plays a critical role in coordination because the GPGP module may engage in an active search process to create constraints that are suitable to coordinate the activities of the local agent with other agents. During this search process, GPGP may query the DTC scheduler repeatedly to explore the implications of constraints, and their associated temporal/resource bindings, and to obtain feedback from the scheduler about alternative constraints that may be more appropriate [18]. In our work, the local scheduler module must also be more than a conventional scheduler (that makes sequencing decisions) since it is required to carry out planning activities like task selection. In essence, the scheduler must make decisions about which one of a pre-defined (though potentially quite large) set of alternative approaches should be used to satisfy a goal and if there are multiple goals to solve, which ones should be solved in order to generate the maximum local utility.

The GPGP and DTC modules present in each agent thus combine to guide the agent's activities using knowledge of its own local situation and partial knowledge of the activities being carried out by other agents. Each agent's local GPGP module also coordinates or negotiates with the GPGP modules of other agents to generate constraints on local control that will lead to more coherent agent activities. The resulting local control decisions made by GPGP and DTC are only approximately correct because of the limits on communication and processing resources. First, complete and up-to-date knowledge of all agent activities is not used due to the cost of obtaining and processing the information. Second, not all GPGP modules are involved in the decision process; only those having goals that are directly related to the goals of the agent making the local decision are¹⁸. Third, the GPGP modules do not engage in a complete/exhaustive negotiation (search) process that is guaranteed to find the constraints to place on local agent control that lead to

¹⁷ In the work on the DECAF agent architecture [20, 21, 23] that uses a variant of TÆMS, a simpler scheduler than DTC, called DRU, was used [22].

¹⁸ In the current implementation, even if the GPGP module is aware that there are multiple relationships among goals of different agents that need to be taken into account, it will treat them independently rather than as a unified group of relationships. There is current work on trying to take into account these other relationships [74]. Additionally, GPGP currently operates mainly on relationships that only span two agents except for one mechanism (avoiding redundant activity) that can handle the situation where two or more agents are planning to do the same activity. Again, the fact that existing GPGP mechanisms mainly operate in the context of two agents is not inherent in the approach but was initially done to simplify the protocols among GPGP modules.

the most coherent behavior. Finally, because of combinatorics, the local scheduler module is not guaranteed to generate optimal schedules [60, 62, 66].

Additionally, it is important to emphasize that the GPGP approach to coordination is an incremental and evolving one. This also leads to an approximation of an optimal agent coordination policy since the choice GPGP makes about how to coordinate at one instant does not fully reflect possible changes in the coordination scenario as the result of the execution of scheduled activities nor the arrival of new goals/tasks. It is incremental and evolving, since as the execution of activities results in behaviors that do not fall into expected ranges, there are new tasks/goals that an agent encounters, or there is new or revised information that the agent receives about the state of other agents' activities, the agent will need to re-evaluate its current plans and schedules¹⁹. This re-evaluation may in turn cause the agent to revise its coordination strategy that may then cause other agents to revise their plans, schedules and coordination strategies.

This incremental and evolving approach to coordination is in contrast to a distributed MDP optimal policy where all contingencies based on the execution behavior of activities in agents are pre-analyzed and factored into the design of an optimal coordination policy [4, 71, 72]. We do not feel this type of complete pre-planning for every execution contingency is feasible in environments where there are a very large number of possible scenarios over some finite horizon involving how goals are distributed among the agents and when goals become active and need to be completed. In such environments, we do not think it is feasible in soft real-time to do the computation necessary to generate an optimal coordination policy so that when a new or revised coordination scenario presents itself a policy can be generated sufficiently quickly to be useful [3]. Further, off-line generation of policies for all scenarios also seems infeasible because of the large number of scenarios²⁰. However, we are not saying that some level of contingency analysis is not worthwhile [48, 70, 77], and in the future section on *Uncertainty and Dynamics of Commitments* we will discuss how this analysis can be used to construct better coordination strategies. Further, in environments where there is more regularity in distribution and arrival rate of goals, which also implies more structure to the transitions among scenarios, a more pre-planned, contingency-based approach may be appropriate.

In essence, there is a spectrum of approaches to coordination that represent different trade-offs (optimality of coordination, computational costs, communication costs, delays, etc.) based on where they lie in a multi-dimensional space. The dimensions of this space include for example: 1) the amount, detail and accuracy of information used about the current and future set of activities for the local agent and other agents; 2) the time horizon for which a strategy is appropriate; 3) the contingencies that are analyzed and the ones that are pre-planned for; and 4) the optimality of the coordination strategy. From this perspective, GPGP makes the following choices: it only acquires information about other agents based on the existence of potential interdependencies among goals that are deemed important in the current situation; it limits the time horizon and analysis to the

¹⁹Over the years, we have explored a number of techniques for reducing the frequency of complete re-planning, re-scheduling, and re-coordination. The ability to reduce the need for re-coordination is also strongly related to the ability of local agent control to find alternative schedules for its activities, in the event of performance deviation, without breaking existing commitments. These techniques have included adding slack time into schedules, local re-ordering of scheduled activities based on existing constraints and commitments, incremental planning and scheduling as new task arrives, replacement of scheduled activities with semantically equivalent activities that use less resources but have poorer performance [17, 58, 76].

²⁰In this discussion, we have not discussed the possibility of generating non-optimal distributed MDP policies on-line quickly (which, though approximate are very good policies), or trying to dynamically adapt a smaller set of optimal policies for a selected set of scenarios to work non-optimally for a scenario in the larger set of possible scenarios. These are definitely possible approaches that may turn out to be very useful to have in a toolkit of approaches to coordination.

known current goals²¹; it does some level of analysis of contingencies in choosing local schedules and commitments but does not pre-generate alternative local schedules to handle specific contingencies, and as describe above, it makes a variety of decisions that limit computational and communication costs to the detriment of the optimality of coordination. As always, there is no one best approach to coordination; the appropriateness of a specific approach to coordination depends on factors such as the size, regularity and dynamics of the coordination scenario space, the type and degree of inter-dependencies among agent goal and activities, the worth structure of different goals, and the uncertainty in agent behavior, etc.

To summarize, GPGP approximates the solution to the global optimization problem of choosing and sequencing activities of a group of agents that generate the highest combined utility by breaking the problem down into a set of asynchronous and evolving local optimization problems, one per agent, where part of the optimization problem is solved by GPGP and part by the DTC scheduler. Each of these local optimization problems involve choosing and sequencing local agent activities; these local optimization problems are in turn modified to reflect the dynamically evolving interactions of activities in different agents with these local activities. This modification occurs through the addition of constraints²² on local agent activities (commitments to other agents) and expectation of the results of non-local agent activities (commitments by other agents). These modifications occur through a limited sharing of information and search by a small group of agents.

Implementation of Underlying Conceptual Model

GPGP starts by each agent (i.e., the domain problem-solving component of the agent) constructing its own local view of the activities (goals and their associated task structures) that it intends to pursue in the short-to-medium-term time frame, and the relationships among these activities. The TÆMS representation, an enriched representation of a goal tree that includes quantitative information and temporal sequencing constraints plus more dynamic information on the state of subgoal scheduling execution, is the shared representation that is used by the problem-solving, coordination and scheduling components to communicate among themselves. This view of local agent activities may be augmented by information both static and dynamic about the activities of other agents and the relationship between these activities and those of the local agent; thus, the view evolves from a local view to one that is not entirely local (i.e., partially global). Individual coordination mechanisms that are part of GPGP help to construct these partial views, and to recognize and respond to particular inter-agent task structure relationships by making commitments to other agents. These commitments result in more coherent, coordinated behavior by affecting the tasks an agent will execute, when they will be executed, and where their results will be transmitted. Integral to GPGP is a domain-independent scheduler that, based on commitments, agents' goals, the local and non-local values of tasks, and other agent activity constraints, creates an end-to-end schedule of activities for the agent to perform²³ that meets real-time deadlines and addresses cost and resource constraints and preferences. In this way, GPGP coordinates the activity of agents through modulating their local control as a result of placing constraints and commitments on the local scheduler. No one coordination strategy is appropriate for all task environments, but by

²¹ As implemented in [64], a simple model of future task arrival could be easily incorporated into GPGP through an opportunity cost model.

²² These constraints consist of both simple constraints on the temporal ordering of activities, such as the earliest time that an activity can start, the latest time it can finish, and a time interval at which the activity cannot occur and more complex constraints such as guaranteeing that an activity (or related group of activities) is completed with a certain quality by a specified time.

²³ Schedules, as previously discussed, are revised or replaced as new information arrives or the problem solving context changes. This end-to-end view is necessary in order to meet overall real-time and resource constraints. This is in contrast to the myopic view employed in [28].

selecting from a set of possible coordination mechanisms (each of which may be further parameterized), we can create a wide set of different coordination responses that provide a range of trade-offs between the overhead of coordination and the optimality of coordination.

To reiterate, the central representation that drives all of the GPGP coordination mechanisms is that of the local (and non-local) task structures as specified in TÆMS. Several important pieces of static information are captured in a TÆMS representation of agent activities. Figure 11 (see page 25) illustrates pictorially some of these items, which include: (a) the top-level goals/objectives/abstract-tasks that an agent intends to achieve including the deadline for their completion and the earliest time they can begin being executed, (b) one or more of the possible ways that they could be achieved, expressed as an abstraction hierarchy whose leaves are basic action instantiations, called methods, (c) a precise, quantitative definition of the degree of achievement in terms of measurable characteristics, such as solution quality and time, called quality-accumulation-functions (*qaf*'s), (d) task relationships that indicate how basic actions or abstract task achievement affect task characteristics (e.g., quality and time) elsewhere in the task structure, and (e) the resource consumption characteristics of tasks and how a lack of resources affects them²⁴. The existence of interdependencies among subtasks situated in different agents as specified by *qaf*'s, task relationships and resource consumption characteristics are the triggers that are used by GPGP to establish coordination of activities among agents. It is also important to emphasize that the TÆMS representation is not intended to capture the detailed problem-solving state of an activity. The only state information easily representable is the level of resources used and the amount of quality so far achieved²⁵, and it is the responsibility of the domain problem solver to update this information as methods complete execution. This choice of representation is quite unique and powerful. Unlike many systems, it deals with *worth-oriented* [52] environments where a goal is not black and white, but rather has a degree of achievement associated with it. We also allow many task structures to be active at once, representing several objectives all of which must be achieved to some degree. Additionally, the agent's task structure view may change over time due to uncertainty or a dynamically changing environment²⁶.

Qaf's are generalizations of logical constraints among tasks and their subtasks such as *and* and *or* to quantitative functions such as *min* and *max*, respectively. The idea is to provide a repertory of such *qaf*'s that can be used to approximate how the relative success in achieving a low-level problem-solving activity will eventually affect higher-level activities of which they are a part. In a similar way, task relationships represent a quantified view of temporal constraints among activities²⁷ as a result of information sharing relationships. Hard relationships (e.g., *enables*) denote when the result from one problem-solving activity is required to perform another²⁸, or when performing one

²⁴ Though resource consumption characterizations were part of the specification of TÆMS as defined in [11] they were not used as a basis of coordination in the original GPGP.

²⁵ As will be discussed later, TÆMS has been extended to include outcomes on methods. In this way, a limited and local representation of domain state can be represented. Additionally, pseudo resources can be created to add in some capability for representing a more global domain state.

²⁶ An agent's problem-solving component is expected to modify the task structure to reflect its revised view of its expected domain problem solving activities.

²⁷ Task relationships as specified currently in TÆMS only involve relationships between two tasks. There is no reason that more complex relationships irreducibly involving more than two tasks could not be created but we have not yet found a compelling need for such relationships in effectively modeling subtask interdependencies for the applications we have so far studied. It should be noted however that *qaf*'s represent relationships among an arbitrary number of subtasks.

²⁸ The *enables* relationship between tasks at different agents indicates that a result needs to be transferred, after the enabling task is completed, to the enabled task at the other agent before it can execute. There are other types of semantics that could be associated with the type of precedence relationship represented by *enables*. For instance in STEAM [47, 56, 57], there is a role dependency relationship between tasks that does not imply information transfer.

activity precludes the performance of another (e.g., *disables*). Soft relationships, such as *facilitates* (or *hinders*), express the notion that the results of one activity may be beneficial (or harmful) to another activity, but that the results are not required in order to perform the activity. The quantified view of these relationships indicate how the quality of the information generated by an activity will affect the performance characteristics of the activity using this information, such as the length of its execution and the quality of its resulting solution. The quantified aspects of these relationships are parameterized so that the user has some flexibility in specification. Resource consumption characteristics are specified in much the same way as task relationships, except that relationship is between a resource and a task rather than a task and a task. We have used the *produce* and *consume* relationships from task to resource to specify the amount of a resource used and produced, respectively, as a result of task execution; the *limits* relationship from resource to task has been used to represent how the lack of the required resources will affect the performance characteristics of the task such as delaying task execution or increasing the amount of time for the task to execute. The set of task and resource relationships that we have chosen is again an attempt to construct an appropriate repertory of such relationships that will allow the faithful modeling of domain problem-solving activities.

By restricting the set of *qaf*'s, task relationships and resource consumption characterizations to be predefined rather than user defined, it permits more computationally efficient analysis of the TÆMS structure by the DTC scheduler and GPGP but at the cost of less faithful approximation by TÆMS of domain problem-solving activity. This efficiency is important since one of the main analyses done repeatedly on the TÆMS task structure is making predictions on how the choice of particular paths through the task structure will affect the expected quality and duration of the overall task. To re-emphasize, TÆMS takes both a very abstracted view of the state of domain problem solving and a very local view of how intermediate activities affect the state of domain problem solving as reflected in the definition of *qaf*'s and task relationships. The key issue is whether this very abstracted and approximate view of domain activities is sufficient to make effective coordination decisions among agents. We will delay a more detailed discussion on TÆMS until the section on extensions to TÆMS.

In terms of BDI-style architectures [51], an agent selects from among its desires an intended set which are then “planned” for, and whose component actions are then “scheduled” for execution. In this paper, we view “planning” as the process of elucidating a task structure with a manageable number of possibilities, and “scheduling” as the process of selecting specific intended actions and locating them in time. In fact, planning, scheduling, and action execution can be viewed as increasingly specific and less revocable intentions to more and more precise action specifications. As stated earlier, GPGP, despite the name, has so far focused primarily on coordinating the scheduling process²⁹.

The GPGP approach uses this basic TÆMS task structure representation, and adds two important extensions: partial representations of the task structures at other agents, and local and non-local

Rather, it indicates that one task has created a state of the world that is necessary for the execution of the other task. A transfer of a message from one agent to the other can indicate the existence of this state of the world or it can just be observed by the dependent task. We can easily model this type of “role dependency” relationship in our framework, either through the use of a new relationship that is similar to *enables* but without the information transfer, or with the use of the *producer/consumer* relationship on a pseudo resource which represents the state of the world that is desired. In fact in [11, pp. 54–58], a more complete list of relationships (*facilitates*, *hinders*, *precedence*, *causes*, *shares-results*, *cancels*, *favors*) are specified, many of them oriented toward modeling interrelationships among non-computational activities such as “*causes*” where the result of completing one activity results in the completion of another activity.

²⁹The name GPGP derives from Durfee's Partial Global Planning [16], which by today's definition of AI planning also focuses on the scheduling end of the planning-scheduling spectrum.

commitments to task achievement. Each GPGP coordination mechanism is specified with respect to features in these dynamically changing task structures, and so provides flexible coordinated responses across different problem domains. The quantitative, worth-oriented approach has several benefits when compared to symbolic, preplanned, domain-specific approaches [29]. In particular, GPGP is not domain-specific, but rather task-structure-characteristic-specific. The same task relationships may appear across many different domains, and the same coordination mechanisms can be used³⁰. Efficient and effective coordination must account for the benefits and the costs of coordination in the current situation. The current situation includes the goals that the agent is currently pursuing, the goals it will likely pursue in the near term, the characteristics of the abstract tasks and basic actions available to achieve these goals, their relationships to other tasks (possibly at other agents), and the degree of achievement necessary for each goal. Purely symbolic reasoning about costs and benefits can be extremely complex, particularly in large systems and open environments. Thus, the ability to reason quantitatively about the benefits and costs of coordination seems essential for effective system operation where there is a large set of situations that need to be reasoned about [37].³¹

Connections to Joint-Intention Model

The initial set of mechanisms for GPGP was derived from Durfee's work on PGP [16] in the Distributed Vehicle Monitoring Testbed (DVMT). These five mechanisms were: (1) communicate non-local views; (2) communicate appropriate results; (3) avoid redundancy; and (4,5) handle hard and soft coordination relationships that extend between tasks at two agents [10, 11]. These five basic mechanisms formed a domain-independent set for basic agent teamwork. As such, it is interesting to compare them to Tambe's work on flexible teamwork, implemented in the STEAM and TEAMCORE frameworks [47, 56, 57], that is in turn based on Cohen & Levesque's work on joint-intentions³² [42]. In particular, the five initial GPGP mechanisms do not separate the environment-centered coordination actions from actions oriented toward teamwork itself. In other words, in general, GPGP does not communicate solely to maintain some explicit model of teamwork. However, a striking example of the relationship of GPGP to the Cohen and Levesque teamwork model is their prediction that team members should communicate when they believe that the object of a joint persistent goal has been achieved, found to be unachievable, or irrelevant. The GPGP "communicate results" mechanism contains a clause to communicate a "final result" (degrees of achievement) for every task group (implicit joint persistent goal) indicating that the agent believes no further actions can be done for that task group. Interestingly, this is one of the small generalizations that were added when creating the general GPGP mechanisms from the original DVMT PGP heuristics.

GPGP currently only represents joint intentions implicitly. There are no explicit mechanisms to guarantee agents have mutual intentions to achieve their parts of the jointly shared goal. GPGP generates commitments to accomplish subtasks of a jointly shared goal though these detailed commitments are not in response to a more encompassing commitment, explicitly held by each agent, to achieve the jointly shared goal. There is no guarantee that an agent will solve its part of a

³⁰The question of which mechanism to use when more than one is possible is a question about organizational/societal norms.

³¹In fact, there are two primary classes of reasoning embodied in the GPGP/scheduler pair. The first class pertains to the reasoning required to form and propose commitments with the other agents, the second class pertains to understanding how the choice of actions is affected by commitments and when/if commitments may realistically be satisfied. In our current work [63], we are further separating these two different types of activities so that commitments may be specified by components other than the GPGP coordination module.

³² STEAM/TEAMCORE framework extends the original joint-intentions model in some important practical ways such as providing mechanism/reasoning so that agents' can 1) switch roles if one of their teammates fails in fulfilling its role, and 2) selectivity choose whether to follow the prescriptions of the joint-intention model based on a decision theoretic analysis of the costs of communication in the current situation.

jointly shared goal though it knows of its existence. In fact from a social utility perspective, an agent may choose to solve other goals/tasks that it knows of rather than its part of a jointly shared goal. However, what is of concern is if the following situation occurs: an agent solves its part of a jointly shared goal while another agent chooses not to solve a critical part of the jointly shared goal, thus making the other agent's effort worthless. Additionally, there are limited mechanisms in the initial GPGP to coordinate about which alternative ways agents should achieve their local subtasks of the jointly shared goal that would be maximally beneficial to them in the context of the other tasks they are doing. This latter problem entails a complex distributed search due to the complex *qaf* relationships possible among distributed subtasks. We have only recently begun to solve this problem in restricted cases as will be described in the section on more complex coordination mechanisms. Of the original five initial mechanisms only the redundancy mechanism is triggered by the existence of a *qaf*'s relationship (i.e., *max*) among sibling subtasks located in different agents, and it uses a very simple heuristic not involving distributed search for making a decision³³. Neither the lack of explicit mechanisms to generate a joint intention nor the lack of mechanism to handle *qaf*'s such as *min* (*and*) turned out to be a serious problem in the experiments [10, 11] that we have conducted. Problems did not often occur because: 1) subtasks in a jointly shared goal were often connected through interrelationships so that failing to schedule a subtask in one agent would lead another agent to not schedule its related subtask—this resulted in implicitly deciding to not solve the shared joint goal³⁴; and 2) the scheduling heuristic used by the local scheduler (DTT [19]) was to first guarantee, if possible, that all local task structures achieved some minimum level of utility before attempting to maximize the overall utility achieved by all local tasks. Based on these reasons, the GPGP mechanisms operated correctly in most cases by either explicitly generating commitments to solve a subtask of the jointly shared goal (or not solve it) due to an interrelationship between a subtask of the shared goal held by another agent or scheduling these subtasks due to local utility concerns. We now feel that with a few additional messages being transmitted among agents³⁵, GPGP could be made to work in all cases within our existing mechanisms; these messages would guarantee that when an agent made a decision not to schedule activities to generate utility for their part of a shared goal other agents having parts of the shared goal would be notified. In essence, we are saying that the joint-intentions model is compatible and integratable with our quantitative view of coordination and the issue of implicit versus explicit representation of jointly shared goals becomes moot as we add additional attributes to the TÆMS task structure representation necessary to support the reasoning associated with these additional messages being exchanged among agents.

³³ We have also explored more complex heuristics involving load balancing [11, pg 163-164].

³⁴ When non-local tasks were connected to local tasks as a result of task relationships such as *enables*, this would trigger a coordination mechanism to be invoked. If successful, it would result in a commitment that would guarantee that the result needed by the local task would be generated in a timely way. If a commitment could not be achieved successfully, then either an alternative way of solving the local task that did not need the non-enabled task was found, or the local task would not be scheduled for execution. This knowledge of the resultant local scheduling decision would be transmitted to the other agent so it could update the expected utility to be produced by doing its local task as part of the larger task. If this utility were zero, the other agent would not schedule tasks associated with its part of the shared goal. However, this local utility as contained in a non-local commitment at a remote agent was not always changed when it changed at the local agent because too much communication was needed.

³⁵ We already send messages alerting agents of failed commitments. What we don't handle correctly in all situations is a case where the lack of successful generation of a commitment may still allow agents to find another way to solve their part of the joint goal, nor the case where there are no task interrelationships among the task structures that make up the shared goal and there is only a *qaf* relationship that defines the connection among activities. However, it should be mentioned that none of these issues arise when one agent contracts with another agent to do part of a higher-level goal, i.e., top-down coordination. All the problematic situations described above occur when agents need to recognize that their local activity is part of a higher-level goal to which other agents are also contributing, i.e., bottom-up coordination. In this case, the execution of their local activity does not make any sense unless other agents are committed to executing their local activity, which is part of the jointly shared goal.

Additionally, Tambe’s STEAM/TEAMCORE work points to a need to more carefully define coordination mechanisms to separate bottom-up, environment-centered coordination activities from top-down, organizationally-prescribed teamwork coordination activities — an exciting area of future work. On the other hand, the success of using GPGP environment-centered mechanisms in multiple environments demonstrates that there are indeed general mechanisms for solving domain-level coordination problems in different environments, because the task relationships in those environments can be characterized in a general way. While STEAM’s teamwork rules are general, it still requires the addition of new, different, domain-specific coordination rules for each domain. Another difference is that TÆMS and GPGP deal in worth-oriented environments, thus goals are achieved to varying degrees, rather than black and white “achieved or not achieved.”³⁶

A Detailed GPGP Example

To illustrate how GPGP operates, consider the task structure shown in Figure 3. The task structure represents a global view of all of the activities being carried out by three cooperative agents, *A*, *B*, and *C*, and the interactions between these activities. In this example, the agents are fully cooperative and working to achieve common objectives so the utility produced by one agent has equal value to another agent.³⁷ In this structure, task *A1* facilitates task *C1* and task *A2* facilitates task *B1* – in other words, agent *A*’s tasks interact with tasks being carried out by each of the other agents. Given this global view, control problem solving, while exponential, is straightforward. However, in GPGP we do not assume that a global view exists or that it is even composable. In large-scale multi-agent systems, particularly those embedded in dynamic environments, construction of the view is generally not possible and the combinatorics of centralized problem solving make the centralized approach intractable.




Figure 3 – The Hypothetical Global View

In GPGP, initially, each agent has a view of only its own local problem-solving activities. Figure 4A shows agent *A*’s initial view, Figure 4B shows *B*’s initial view, and Figure 4C shows *C*’s initial view. In this example, the agents are autonomous and operating on their own local task structures. None of the agents’ activities require input from any other agent and the potential relationships between their activities are not known *a priori*. Each agent’s tasks also have quality, cost, and duration characteristics, as well as potential local constraints such as deadlines and earliest start times. In this example, we will assume that the agents are fully cooperative (attempting to maximize

³⁶ In STEAM, there are some quantitative thresholds that are used to indicate whether a goal was successfully achieved, e.g., 70% of the planes reached their destination. However, there is no reasoning to our knowledge in STEAM about how to dynamically set this threshold given current conditions and how to generate different plans that will likely achieve different thresholds.

³⁷ Because the agents are fully cooperative, in the global view, there is actually a root task common to all the agents that joins their activities. In other words, the cooperative relationship that results in the summing of the utilities produced by the individual agents is represented explicitly in TÆMS by placing a super task over the tasks of the individual agents. In this example we have omitted the root task to clarify the different world views held by the different agents.

overall utility) and that the tasks have the following durations and deadlines: *A1* has a duration of 6, *A2* a duration of 5, *B1* a duration of 5 and a deadline of 10, *C1* a duration of 4 and a deadline of 10. We will assume that the other tasks are unconstrained and ignore their characteristics, to clarify the example.




Figure 4A - Agent A's Initial View




Figure 4B - Agent B's Initial View




Figure 4C - Agent C's Initial View

Each agent constructs a local schedule based on its local view. This is generally done by the Design-to-Criteria (DTC) agent scheduler [60]. DTC analyzes the agent's task structure, its constraints, and evaluates the different trade-offs of different possible courses of action and custom tailors a schedule for the agent. Given that the agents' initial views are only of their local problem-solving activities and their local constraints, the initial schedules of the agents are shown in Figure 5. Note that time 10, which is the deadline for tasks *B1* and *C1*, is identified in all of the schedules.




Figure 5 - The Agents' Initial Schedules

If agent *A* encounters agent *B* during the course of problem solving, the agents may then engage in a dialogue during which they compare their activities, or portions thereof, and look for interactions between the activities. In GPGP, the mechanism used to exchange and compare task structures is called *exchange-non-local-views*. It is during this dialogue that the facilitation relationship between agent *A*'s task *A2* and agent *B*'s task *B1* is discovered. To ground the example, if the agents were information-gathering agents, the facilitation might model agent *A* providing a URL to agent *B* so that agent *B* does not need to search for said URL. After the exchange, agent *A* and agent *B* have views of each other's activities as shown in Figure 6. Note that if the agents have larger task structures, it is unlikely that the agents will exchange entire task structures³⁸.




Figure 6 - Agent A's and B's View After the Information Exchange

³⁸ In the experimental work cited in [10, 11], the norm for most experiments was to exchange only partial views.

After *A* and *B* detect the facilitation relationship they can then compute that performing *A2* before *B1*, and sending agent *B* the result from agent *A*, will improve the overall utility of their problem solving. Thus, to improve problem solving the agents deploy the GPGP *handle-soft-interactions* mechanism and agent *A* reschedules, using DTC, and offers agent *B* a *commitment* to produce a result for *A2* by time 5 so that agent *B* can execute task *B1* between time 5 and 10 and still meet its deadline. Agent *B* then reschedules, again using DTC, to take advantage of this result. Their respective schedules after the commitment is formed are shown in Figure 7.




Figure 7 - Agents *A* and *B* Revise Schedules to Leverage Positive Task Interaction

The example thus far illustrates a single exchange, discovery, and commitment formation between two cooperative agents. However, GPGP was designed to operate in an open environment where this type of activity occurs frequently as agents discover one another and as the environment changes or the tasks being performed by the agent's change. Control in this context is a difficult problem because little remains static – the goal in GPGP is to perform control from a satisficing perspective and to respond to these dynamics of the environment.

Returning to the example, let us now assume that agent *C* encounters agent *A*. As with the exchange between *A* and *B*, agents *A* and *C* will now exchange their local views. Agent *A*'s revised view is shown in Figure 8 and agent *C*'s in Figure 9. Note that agent *C* does not have a view of agent *B*'s activities – only agent *A* now has a full view of the activities happening at the other agents.




Figure 8 - Agent *A*'s Revised View




Figure 9 - Agent *C*'s Revised View

Let us assume that when task *C1* is facilitated by *A1* it will produce greater utility than *B1* when it is facilitated by *A2*. Note that because of the time 10 deadline held by both *C1* and *B1* agent *A* is unable to facilitate both tasks. In this case, agent *A* must break its commitment with agent *B*, and then reschedule in cooperation with agent *C* to form a new commitment to provide the result of task *A1* to agent *C* by time five so that agent *C* can perform *C1*. The act of decommitting from agent *B* may involve a penalty depending on how the decommitment is structured. In the original GPGP

work, decommitment penalties were not used³⁹. The revised schedules for the agents are shown in Figure 10.




Figure 10 - The Final Schedules of Agents A, B, and C

This example shows the basic style of coordination in the original formulation of GPGP, which involves an incremental, dynamically evolving process. It also indicates the close interplay between the local scheduler and coordination mechanisms.

Extensions to TÆMS

In its original conceptualization, TÆMS was a hierarchical task representation language that featured the ability to express alternative ways of performing tasks, characterization of methods according to expected quality and expected duration, and the representation of interactions between tasks. As a result of using it in a number of applications, TÆMS has evolved considerably in its representational power and in the scope of applications that can be appropriately represented.

Addition of Discrete Distributions

As TÆMS has evolved, we have moved beyond simple quality/time trade-off reasoning to a multi-dimensional satisficing methodology [60]. Methods are no longer simply characterized in terms of expected quality and duration values but are characterized statistically via discrete probability distributions in three dimensions: quality, cost, and duration; and the work is extensible to multiple other dimensions as well. The switch from expected values to discrete distributions enables reasoning about the certainty [61] of particular courses of action, perhaps performing contingency analysis [50, 77], as well as the quality/cost/time trade-offs. This more complete view is crucial in highly constrained situations (e.g., where deadlines are present). For example, in meeting an important commitment, it may be preferable to choose a highly certain course of action that has only moderate overall quality to ensure that results are available by the deadline, rather than a less certain solution path, even though it may have the possibility of a very high-quality payoff. The enhancement of multiple attribute dimensions, including uncertainty, is also important because the representation enables the multi-agent system designer to better characterize the desired system performance [60, 62]. This new performance characterization takes into account both hard and soft constraints on quality, cost and duration of the schedule of tasks intended to meet a high-level goal or commitment. Additionally, the amount of uncertainty that is tolerable in meeting these constraints can be specified [61].

Additional Quality Accumulation Functions

In TÆMS, progress toward the problem-solving objective is expressed and tabulated in terms of quality. Tasks accumulate quality from their subtasks according to *quality accumulation functions* (*qafs*). Originally, *qafs* defined only which combinations of subtasks could be performed to achieve the parent task and how the subtask qualities are tabulated at the parent task. For example, the *sum()* *qaf* denotes that the power set of the child tasks, minus the empty set, can be performed to achieve the parent task and that the parent task's quality is the sum of the children's qualities. In the current

³⁹ Although decommitment penalties were not used, an alternative, more qualitative mechanism was used to make decisions about what commitments could be decommitted from and in what circumstances. A commitment when initially constructed was marked with different "levels of negotiability" which reflected its relative importance.

TÆMS, *qafs* may also define orderings⁴⁰. In all, we have developed a suite of nine different *qafs* to represent the needs of different application domains.

<i>q_min</i>	- Minimum single quality of all subtasks.
<i>q_max</i>	- Maximum single quality of all subtasks.
<i>q_sum</i>	- Total aggregate quality of all subtasks.
<i>q_all</i>	- As with <i>q_sum</i> , except all subtasks must be completed.
<i>q_seq_min</i>	- As with <i>q_min</i> , except all subtasks must be completed in order.
<i>q_seq_max</i>	- As with <i>q_max</i> , except all subtasks must be completed in order.
<i>q_seq_sum</i>	- As with <i>q_sum</i> , except all subtasks must be completed in order.
<i>q_exactly_one</i>	- Quality of single subtask. Only one subtask may be performed.
<i>q_sigmoid</i>	- Similar to <i>q_sum</i> except there is a parameterized sigmoid function that defines how quality accumulates

Table 1: Definition of Quality Accumulation Functions

Another important enhancement is the representation of different possible outcomes for primitive activities. For example, an action may produce a result of type *A* or type *B*, or it may fail entirely. With the outcome enhancement, clients can represent cases where a result of type *A* is beneficial to some other method, say method *A*; type *B* is beneficial to yet another method, method *B*; and the failure outcome enables a recovery action. While outcomes have always been included conceptually in TÆMS, either via entries in a method's quality, cost, and duration distributions or via a contribution to the expected values in the early versions of TÆMS, this new explicit outcome representation enables the expression of task interactions that are based on particular outcomes⁴¹. The new outcomes feature also allows the representation of different outcomes that have different statistical characteristics, i.e., different discrete probability distributions for quality, cost, and duration.

⁴⁰ The addition to *qafs* of task ordering information such as in *q_seq_sum* has provoked much discussion. One argument against inclusion of this type of information is that with the appropriate set of temporal relationships among tasks there is no need for the addition of ordering information to *qafs*. However, we eventually came down on the side of ease of representation that occurs when in certain situations the *qafs* contain such information.

⁴¹ However, we have not developed the equivalent of *qafs* that are able to specify how outcomes at the method level should be propagated to outcomes of higher-level tasks. The reason is that these functions seem very domain-dependent and not amenable to systemization. This can be solved in part by having code in the domain problem solver that will compute the higher-level outcome to be associated with a specific higher-level task. This can be represented by slightly modifying the task structure such that each higher-level task that requires an outcome will have a brother task, which is a method with the set of outcomes that are possible for its brother higher-level task. The method, when executed after its brother higher-level task completes, is in reality a call to the appropriate code in the domain problem solver that computes the desired outcome to be associated with the method. Unfortunately, this does not solve the problem completely since there is no way for the scheduler to reason about the likelihood of different higher-level tasks' outcomes given a specific plan to achieve the higher-level task. An alternative approach to solving this problem was taken in the DECAF architecture [20, 21, 23]. In DECAF, task outcomes are signaled by the child tasks (and thus the outcome probability is the same as that of the child outcome). The outcome of a parent task does not need to be signaled by the same, or even a unique, child.

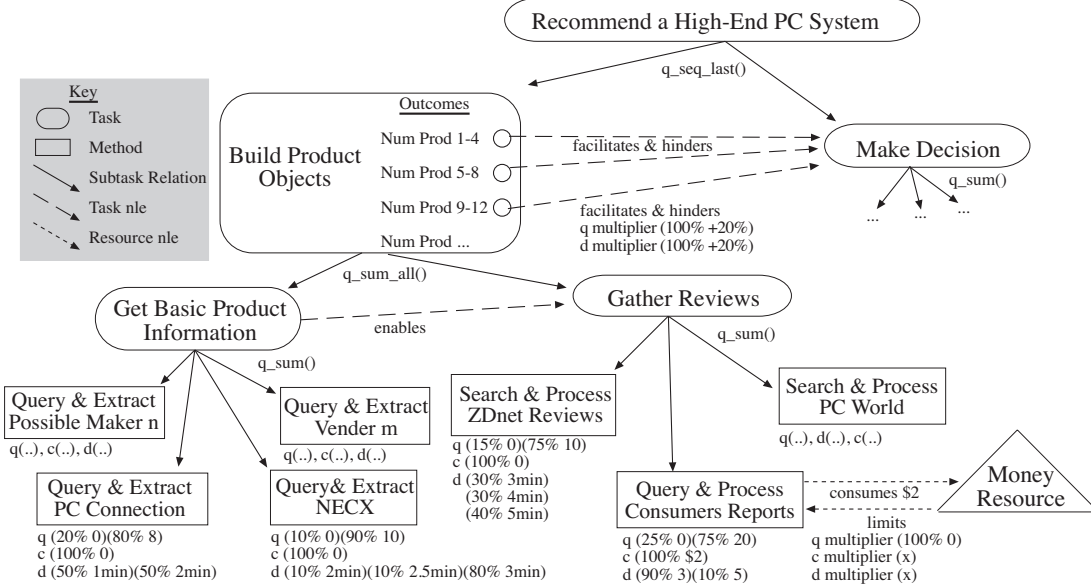


Figure 11: An Example of a TÆMS Task Structure for an Information-Gathering Agent

Figure 11 shows an example information-gathering TÆMS task structure, from the BIG information-gathering agent [41] that contains some of these new extensions to TÆMS. The top-level task is to recommend a high-end PC system. It has two subtasks: one that pertains to finding information about various products and constructing models of the products, and one for making the decision about which product to purchase. The two tasks are governed by a *seq_last()* *qaf*, which indicates that the tasks must be performed in order and that the quality of the top-level task is determined by the quality returned by the last subtask executed, i.e., the decision process task. This models the fact that the decision process described here takes into consideration the quality, coverage, and certainty of the information used to make the decision and reflects these attributes in the quality of its final decision. The task that handles finding, extracting, and building information objects *Build-Product-Objects* is decomposable into tasks for finding basic product information and for gathering reviews. The *sum_all()* *qaf* denotes that both of these tasks must be performed and that the quality of *Build-Product-Objects* is a sum of the qualities produced by the subtasks. *Build-Objects* has a set of different outcomes⁴² that indicate the number of objects produced during the information-gathering phase. Note the dotted edges leading from *Build-Objects* to *Make-Decision*. These denote task interactions also called *non-local-effects* (or *nles*). The *nles* pictured denote facilitation and hindering effects to model the notion that the decision maker's quality is improved by more objects but that more objects also increases the time required to make the decision. *Get-Basic-Product-Information* has four children governed by a *sum()* *qaf*, which means that the power-set of the children (minus the empty set) may be performed and that the quality of *Get-Basic* is the sum of the qualities of its children. Note that the child tasks have different quality, cost, and duration characteristics—essentially, there are 2^4-1 different alternative ways to perform *Get-Basic* and the agent will determine which approach to employ based on the current problem-solving context. For example, if time is limited, it may only get information from one fast site as opposed to a set of sites.

⁴² The actual information-gathering task structure does not incorporate outcomes at the task level. This example is a conceptual abstraction of the actual class of task structures produced by the information-gathering agent's planner and is simplified for example purposes.

Representation of Non-Computational Resources

The explicit representation of non-computational resources, such as network bandwidth, databases, and physical resources like printers, is another important TÆMS enhancement. To do this, we represent each resource explicitly in the existing task structure. As with task–task relations, a task–resource relation indicates how the execution of a task affects the state of the resource, and how the state of the resource affects the characteristics (quality, cost, duration) of the tasks to which it is related. A monetary resource is shown in Figure 11 with appropriate consumption and limiting resource relationships. Iteration and repetitive tasks are still more enhancements recently added to TÆMS but not yet fully implemented.

Virtual Tasks

Figure 12 shows the same information-gathering task structure as it might be decomposed and assigned to several agents. Obviously, many such decompositions are possible; the decomposition in Figure 12 is according to the WARREN [13] approach to agent specialization in multi-agent information-gathering systems. The *Task Agent* is responsible for producing the result but *Information Agents* are the retrieval and processing experts. The interrelationships between the different task structures denote GPGP coordination points. The task structure in Figure 12 actually contains several different types of coordination-centric interactions. The inability of the *Task Agent* to carry out the information-gathering actions means it must contract those duties out to the specialists — this is handled by new GPGP mechanisms for virtual tasking or contracting (discussed later). The *enables* relation between the two information agents requires the use of the original GPGP coordinate-hard-task-interaction mechanism. The monetary resource interaction requires resource centric coordination like that found in the recent hospital patient scheduling [14] and the intelligent home automation [39, 40] applications.

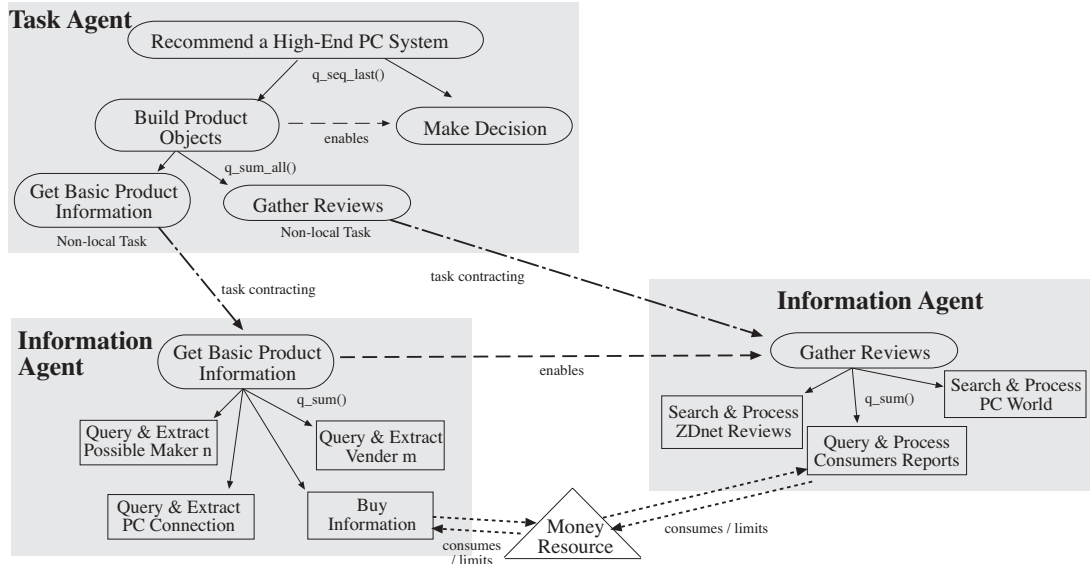


Figure 12: WARREN Style Model of Multi-Agent Information Gathering

The task structures shown in Figure 12 provide a small-scale example of TÆMS and interaction-motivated GPGP. One can easily imagine a pool of information-gathering and decision-making agents cooperating to equip an office. Interactions would exist between information-gathering activities at multiple levels of abstraction. For example, two agents shopping for computer systems might share free-format text reviews of the structured data produced by abstraction. Interactions

would also exist over shared resources and to support data-assimilation and data-driven decision making.

TÆMS as an Abstract Model of Problem Solving

There are two views of how TÆMS is related to agent problem solving. One views TÆMS as a way to represent agent problem-solving activities, that is, agents would literally use TÆMS internally. The other view TÆMS as an abstract model of agent problem-solving activities⁴³. It is this latter view that is typical of our current uses in real applications: a domain problem solver uses its own internal representation for problem solving and this representation is abstracted into a TÆMS task structure for use by GPGP. In this way of using TÆMS/GPGP, we are arguing that for most coordination needs an appropriately abstracted view of activities is sufficient for effective agent coordination. This perspective motivates the continuous evolution of TÆMS. The addition of new features and new representations to TÆMS is caused by the constant tension among representational power of TÆMS to accurately model an agent's problem solving activities, the ease of being able to map these activities into TÆMS, and the ability to reason without significant computational cost about these activities. There is a trade-off between representing large amounts of problem-solving possibilities, and contingencies, and maintaining a somewhat unwieldy modeling framework. For example, if the problem solver enumerates all possible problem-solving actions for every step in the problem-solving process, and all possible actions caused by carrying out the first set of actions, and all possible contingencies in the case of failure or a change in the environment, the TÆMS model becomes computationally intractable⁴⁴. However, through new constructs like iteration, outcomes, and *qafs* that impose orderings and other semantics, the representational power of TÆMS is enhanced without increasing the computational complexity of reasoning with the TÆMS models. It also allows the domain problem solver to more accurately model key contingencies and key repetitive actions, and thus provide GPGP with more information while not overloading it with details. Without these features, an agent might have to constantly communicate with GPGP about how its view of its future activities needs to be modified as the result of its incremental execution of its tasks; this modified view may entail significant re-coordination activities since other agents may also need to be notified of these changes. Instead, GPGP can now, when it makes its original coordination decisions based on the agent's description of its activities, factor into these decisions that an agent's activities may change in certain ways as a result of the partial execution of its activities. For example, instead of establishing a commitment that is highly uncertain because there is a significant possibility that a task established to satisfy the commitment may not be completed successfully, it is now possible to analyze whether there would be another way that the agent could successfully complete the commitment with additional time. Thus, the coordination mechanism can, in making the original commitment, factor in some of the implications of the original task failing and an alternative method being required, which can result in making a different commitment that has a much higher certainty of being successfully honored.

Other Extensions to TÆMS

We have also explored ideas for a number of other task relationships. We mention them here to emphasize that the TÆMS/GPGP framework can be easily extended to represent other types of task interdependencies. One relationship involves a new type of temporal synchronization among tasks that denotes the constraint that both tasks need to be started and completed within the same temporal window; the more quantitative aspects of this relationship would be the tolerable overlap

⁴³ In this more abstracted view, a schedule of activities generated is more like a high-level policy indicating to the domain problem solver which approaches to problem solving to use when, and what amount of resources to expend on each approach [41].

⁴⁴ Examples of modeling frameworks that have more complex control constructs are [53, 69].

among the temporal windows of tasks and how the task's solution quality is affected by the amount of overlap. Another relationship indicates that there are consistency constraints among solutions to a set of tasks that will require sharing intermediate information among them to converge on a solution; the more quantitative aspects of this relationship would be the expected frequency of information exchange among tasks. Knowledge of both relationships could lead to more coordinated scheduling of task activities in the different agents. Another interesting relationship represents a meta-level exchange of information between tasks. In this case, the relationship indicates that the completion of one task would provide information about the performance characteristics of another task; this information could then be used by the agent's problem-solving component to revise its TÆMS description of the task. This relationship is like *facilitates* in that the task can be performed without an updated view of its performance characteristics, but waiting for this information in certain situations would lead to more optimized behavior⁴⁵.

We have also implemented a more complex view of a method's execution that permits us to model activities that do not use the processor resource for their entire duration [59]⁴⁶. This is accomplished by specifying an additional attribute associated with a method that indicates what percentage of the processor the method will use. Using this information, the scheduler can then overlap method execution. In recent work [58], we have developed a more comprehensive view of method execution that allows us to reason about concurrent execution of methods where there are multiple processors and multiple copies of resources at a single agent and where methods may not be using processor resources. This allows us to build agents that can not only schedule computational activities but also act as resource controllers for shared resources [39, 40, 58].

Extensions to GPGP

There are three main areas of extensions to GPGP. The first concerns how the partial global view used in making coordination decisions is constructed. It also has implications for how GPGP can interact with other components in the agent. The second involves a more complex representation of commitments. The third focuses on new and more complex mechanisms for establishing commitments among agents. These extensions have been made for a number of reasons. One reason has been to reduce the communication and computational overhead of using GPGP — the full use of all the mechanisms is not warranted in all environments. This is also especially important if GPGP is to operate in an environment with a large number of agents where agents take on organizational roles. Another reason is to avoid forcing users of GPGP into only one style of building multi-agent systems. The final reason is the desire to exploit the additional information representable in the extended TÆMS on uncertainty and on usage of non-computational resources for effective coordination.

It is important to mention that a current version of GPGP that includes all these extensions does not exist. Some have been done in different versions of GPGP while others have been done outside of GPGP though they have used TÆMS as their basic representation of task activities and interrelationships. To our knowledge, there is no impediment to creating a version of GPGP/TÆMS that incorporates all of these extensions in a unified way; in fact, this is a current goal of ours.

Constructing and Using the Partial Global View

⁴⁵ The need for this meta-level relationship has been recognized in our own work [55] and by others [1].

⁴⁶ We found this very useful in an information-gathering application [41] where we need to model methods that access a WWW server site for information. In this case, there was not that much processor time spent in executing the setting up of the call to the server and processing the received information but there was a significant delay before the results were returned by the remote activity.

We have significantly enriched how the partial global view used in making coordination decisions is constructed. Two issues motivated this first extension. One was a desire for increased efficiency in coordination by being able to exploit default knowledge whenever it was available. The second was being able to handle not only an agent-centric and bottom-up view on how goals and their subgoals were created and distributed among the agents (which was the original orientation of GPGP) but also a hierarchical and top-down perspective on their creation and distribution. The ability to handle both perspectives in an integrated way significantly increases the scope of multi-agent applications that can be represented in GPGP.

Conditioned View

As mentioned earlier, TÆMS task structures are typically used to encode the different plans for achieving a goal, and the constraints and trade-offs associated with each potential plan. They are also used to describe those instances during execution where it has been determined that coordination is to take place. For instance, a TÆMS structure might indicate that one method enables another. If these two methods are to be performed by different agents, this enablement indicates a point where coordination between the two parties would be beneficial.

Interrelationships in TÆMS therefore have additional meanings. A (*consumes*) relationship between a method and a resource would cause the usage of that resource to be appropriately coordinated over with other entities in the system prior to use. A hard or soft (*enables, facilitates*) interrelationship arising from a method or task represented at a remote agent with a local target would cause the local agent to contact the remote one to coordinate their activities in a satisfactory manner. A relationship indicating potential negative effects on a remote agent (*disables, hinders*) would require the agent to first determine if performing the activity is acceptable before execution. Scheduling a method flagged as being non-local would cause the agent to coordinate over the execution of that method by a remote agent.

In this way, an agent possessing a task structure containing methods or constraints that may be satisfied by many different remote agents may actually represent quite a large range of coordination possibilities. Under some conditions the availability of so many alternatives can be beneficial, but when operating under a deadline or tight resource constraints, too many alternatives can be distracting and wasteful. What is needed then, is a way to represent an agent's true operating potential, while limiting the scope of this representation to only those parts that are reasonable and necessary.

To achieve this, we have added a second view of an agent's task structure; each agent will possess two different versions of its local task structure, called the subjective and conditioned views. The subjective view contains what the agent believes to be a complete view of its local execution alternatives. The conditioned view is a copy of the subjective that has gone through a process of conditioning — it may contain task, method or interrelationship deletions and modifications to better adapt that structure to current operating conditions. When the conditioned view is used for plan construction, these modifications indirectly allow the problem solver performing the conditioning process to focus the attention of the scheduling and coordination mechanisms. These changes serve two purposes: to prevent certain subtrees from being considered during scheduling, and to remove interrelationships that are not to be coordinated over. In the first case, if the problem solver has determined that a particular action should not be performed, it can simply remove it from the conditioned view to prevent it from being considered. The second use is more germane to the issue at hand — by convention any interrelationships present in the conditioned view indicates that the interrelationship should be coordinated over when the related tasks or methods are scheduled for execution. Removing the interrelationship will either prevent the coordination from taking place, or prevent the method from being selected for execution. By making these changes, the overhead required to reason about the task structure is reduced, and the agent's attention is more focused.

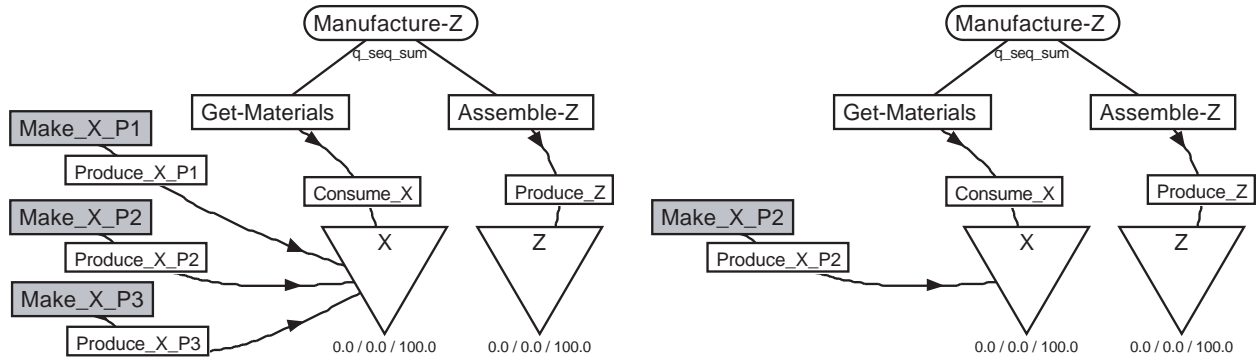


Figure 13: A Simple Example of a Subjective and Conditioned View of a TÆMS Task Structure

Consider the simple conditioning process that has taken place in the task structure in Figure 13. On the left, the subjective view shows that resource X , which is needed as part of Z 's manufacturing process, can be produced by three different agents: $P1$, $P2$ and $P3$. While the local agent could elect to coordinate with all three suppliers, this is time and resource intensive, and can be avoided with the simple use of historical data. In this case, the agent knows that $P2$ did a good job of producing X the last few times it has been used, so as part of its conditioning process the other two candidates are removed, simplifying the selection process and eliminating potentially unnecessary communication. If $P2$'s performance fails to meet expectations this time, the next conditioning process might elect to leave one or two of the other producers in, to increase the probability of an acceptable result at the expense of higher scheduling and reasoning costs.

We have also recognized that many of the additional attributes that are dynamically attached to the TÆMS task structure as a result of coordination mechanisms (e.g., learning about the existence of non-local task relationships and the related task structure of other agents, and dynamic generation of commitments) can be statically predefined. As will be described next, this has led us to be able to not only develop highly efficient coordination protocols for specific situations but also to allow GPGP to operate in a wide range of application environments.

Exploiting Default Knowledge in Coordination

We have also introduced the idea of coordination mechanisms that are based on default knowledge that allows us to represent coordination based on social laws. As previously discussed, coordination in GPGP is achieved through the use of *commitments*, that is, inter-agent contracts to perform certain tasks by certain times. The commitments which are dynamically constructed generally fall into three categories: 1) *deadline* commitments are contracts to perform work by a certain deadline; 2) *earliest start time* commitments are agreements to hold off on performing certain tasks until a certain time has passed; and 3) *do* commitments are agreements to perform certain tasks without a particular time constraint. By allowing these commitments to be specified *a priori* as part of the TÆMS task structure, social law-type coordination, which has low overhead, can be achieved among agents. For example, *a priori* commitments could indicate that an agent could expect another agent to generate a specific result and transmit it to this agent by a certain time [45]. The transmitting agent, in turn, has an *a priori* commitment to generate the results by a specific time. In this way, agent activities can be coordinated without the agents exchanging information about their current activities, and then negotiating over a suitable commitment⁴⁷. More generally, knowledge of potential task relationships can be created based on 1) organizational knowledge of the activities of other agents, 2) another agent decomposing a task structure and then assigning parts of the task

⁴⁷ An easy extension of this idea that we have not yet implemented is for either agent to notify the other that it either cannot honor the *a priori* commitment, or no longer has any need for the commitment to be honored.

structure to different agents, 3) the past history of agent interactions, and 4) a process of extensive broadcast and discovery. In summary, our goal is to have a suite of coordination mechanisms that can reason about in a controlled manner any existing information useful for making coordination decisions, and to acquire in a controlled manner additional information that is important to making the coordination decision in the current situation.

Bottom-Up versus Top-Down Agent Coordination

The TÆMS task structure can be viewed as not only representing an agent’s problem-solving activities but also as a repository of information about how these activities relate to activities in other agents (including knowledge about in what situations will coordination take place, and what commitments have been made among agents and their current status). From this perspective, it is not important how this repository is constructed (e.g., generated locally through dynamic information sharing and negotiations with other agents, statically predefined by the system designer, generated by another agent and then transferred to this agent, etc.); this is true as long as the mechanisms that use this repository make no assumptions about how the repository was constructed and how it will evolve. On the face of it, this does not seem a radical shift in thinking. However, it has important implications for GPGP since it extends the range of applications for which GPGP is appropriate⁴⁸ and it allows GPGP to inter-operate with other coordination schemes.

GPGP traces its origins to PGP, which was designed to handle coordination problems occurring in the DVMT (Distributed Vehicle Monitoring Testbed [35]). In order for agents to coordinate and communicate effectively in this distributed situation-assessment application domain, they need to understand that the vehicle they are about to track is being tracked by another agent who has overlapping sensors, that the previous track positions of a vehicle they are tracking had been or is in the process of being constructed by another agent, or that the vehicle they are tracking is part of a larger pattern of vehicle movements that are being tracked by other agents. Each of these situations can be thought of as recognizing that the tracking goals of specific individual agents were, in reality, part of a higher-level goal to track a specific vehicle or formation of vehicles. This recognition needs to be dynamic because neither the number of vehicles nor the pattern of their movements was predefined. From this perspective, the agents were dynamically recognizing and constructing, where appropriate, shared high-level goal trees from a bottom-up perspective. The GPGP mechanism-1 (“communicate non-local views”) was responsible for this bottom-up construction process⁴⁹. Let us change this problem slightly so that once an agent recognized the track of a vehicle locally, it could make very good predictions about what and when agents will sense that vehicle, and possibly other associated vehicles. In this case, an approach to coordinating the agents would be to have the initial recognizing agent generate a goal structure for tracking the vehicle in the different agents and then appropriately distribute it to relevant agents. This goal structure would consist of individual tracking subgoals for specific agents, and would also indicate how the subgoals were related to each other. For example, two tracking subgoals (one for *agent1* and the other for *agent2*) involved tracking the same fragment of track, since both agents were capable of sensing the vehicle in a specific region. Thus, in this case the generation of the goal tree is top-down and there is no reason for agents to discover interrelationships among their goals and other agents since all this

⁴⁸ In the original STEAM framework [47, 56, 57], they used an approach to coordination where it was not necessary for agents to exchange knowledge about what their activity tasks were. This occurred for two reasons: one reason was that there was only one higher-level goal active at any one time, and the other was that once an agent decided to pursue its part of a higher-level goal the coordination framework made the assumption that the other agents who had roles that mandate work on this goal would also reach a similar conclusion that this specific higher-level goal was the goal that they should now work on. This type of social law can now be easily implemented in our system through appropriate generation of commitments by the domain problem solver when it generates new goals/task structures that it wants to pursue.

⁴⁹ Though the PGP mainly involved the bottom-up creation of subgoals and the recognition of how these subgoals were part of a larger goal, PGP did allow for hierarchical control [16] and task contracting [15].

information is part of the goal structure that is distributed to specific agents⁵⁰. In this way, GPGP can provide effective coordination independent of whether the application employs a bottom-up creation and synthesis of agent task structures, or a top-down elaboration and assignment of agent task structures⁵¹. In fact, we foresee applications in which both approaches to generation of agent task structures co-exist and operate concurrently. Figure 14 embodies the complex goal creation process [33] that we see possible in a complex multi-agent system that we feel can be mirrored in GPGP.

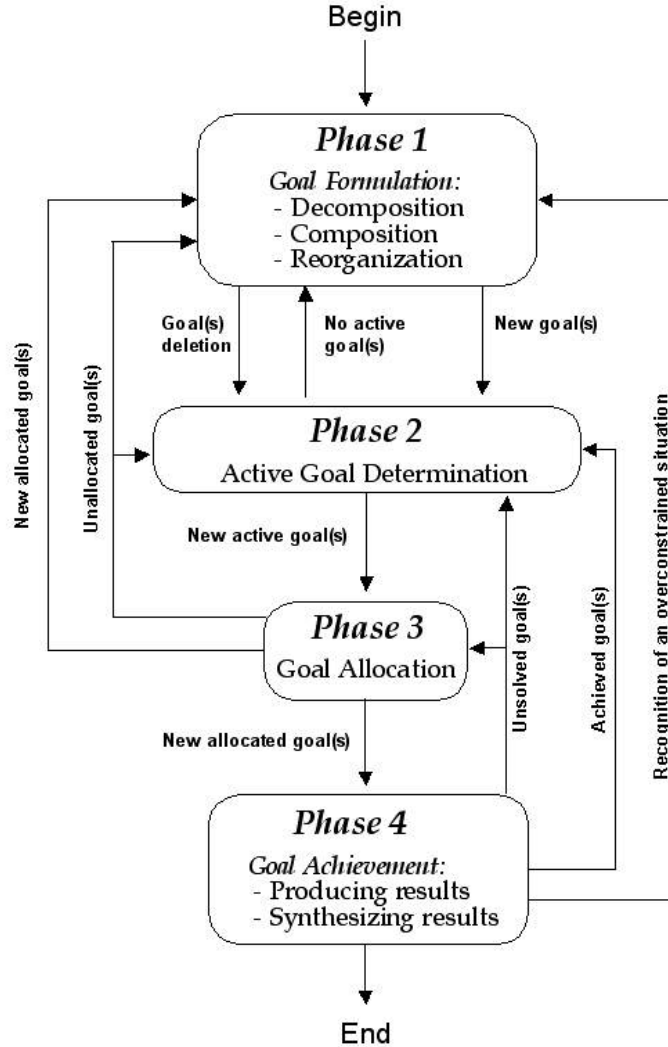


Figure 14: Interaction Patterns among Phases of Goal Creation, Selection, Allocation and Achievement

⁵⁰ In fact, if you knew that there would be only one vehicle in the environment the information contained in the generated goal structure could contain commitments on exactly when certain activities would be completed. In [17], we explored such an approach to hierarchical control where significant deviations in meeting commitments would then cause the higher agent to be notified and for it to generate a new more appropriate global coordination plan. However, in a multiple vehicle environment you want the agents to coordinate with appropriate agents since one agent may be working on more than one vehicle-tracking task.

⁵¹ In early research [54], bottom-up creation/synthesis of goals was associated with a style of problem solving called *result sharing*, while the top-down elaboration/assignment of goals was associated with the style of problem solving called *task sharing*.

GPGP interactions with other components

In using GPGP as part of a real multi-agent system, we have developed a more complex view of the interaction among GPGP, the scheduler and the agent's problem-solving component. We have introduced an additional component: a (domain-specific) task assessor (see Figure 15). The problem solver and its associated task assessor use TÆMS to represent the tasks that will likely need to be carried out in the future. The methods that the task assessor provides to the scheduler represent higher-level entities, which may still require that numerous decisions be made before the actual selection of executable actions is accomplished. It is the problem-solving component that makes these decisions—possibly in a reactive manner. This is one way that we have accommodated the need for some reactivity (without necessitating repeated reschedulings)—by having the task assessor only abstractly plan how to accomplish tasks⁵². The agents' local coordination modules interact, possibly resulting in modifications to the agents' local task structures to represent inter-agent task relationships. This information is used in implementing domain-independent coordination protocols. The coordination-adapted task structures are then used by each agent's real-time scheduler to find the best sequences of activities to meet both the local and network-wide objectives. The resulting schedule of tasks (methods) directs the problem solver, which ultimately determines what actions to take to best achieve a task. The TÆMS description of the expected tasks created by the task assessor is used to update the current task model that is accessible to the scheduler and coordination module. This triggers the scheduler to generate a new schedule, which is passed back to the problem solver to direct the selection of actions. The schedule consists of an ordered list of methods to be carried out. It implicitly represents the order in which the top-level tasks should be pursued as well as the approaches that should be used to pursue them. The scheduler also associates monitoring information with each method that indicates in what situations the problem solver should alert the scheduler that the partial execution of the method is not performing as expected.

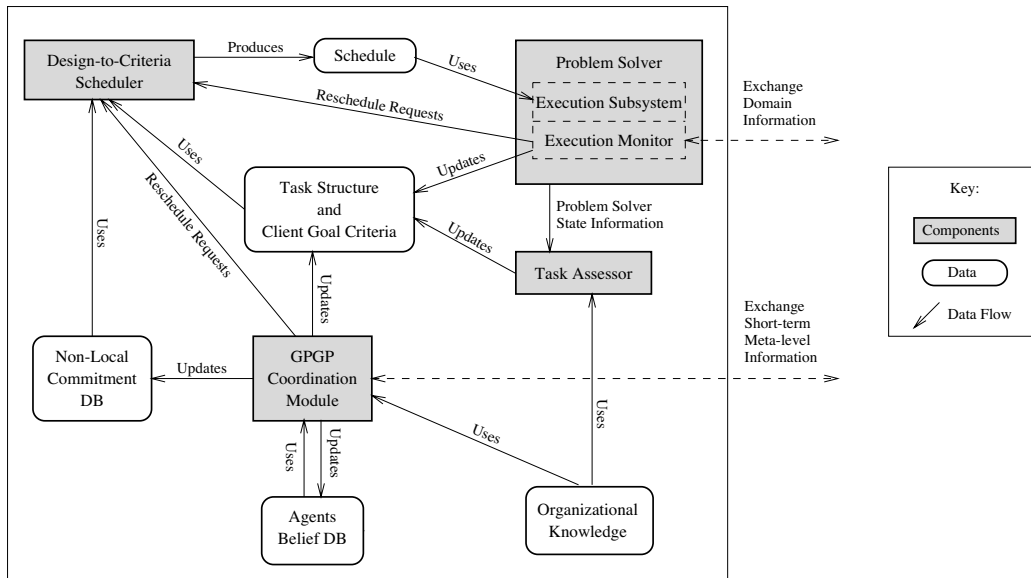


Figure 15: A Typical Agent Architecture

⁵² In recent work on SRTA a soft real-time agent architecture [58], we have developed a variety of mechanisms to avoid complete rescheduling when there are ways to do local shifting of activities—replacing activities with functionally similar but less resource-intensive activities, and inserting a precomputed sequence of activities to recover from execution failures (such precomputed recoveries may be learned over time).

We can also think of GPGP interacting with other coordination mechanisms (e.g., domain-specific, user-defined coordination rules built into a BDI system integrated into the domain problem-solving component) that, for instance, would generate and assign parts of goal trees to agents and make tentative commitments to the achievement of certain goals by specific agents. GPGP would then attempt to verify the feasibility of achieving these commitments in the context of the given goals structures using more detailed reasoning about resource usage and analysis of the comparative worth of achieving specific goals, and to further elaborate the commitments with, for instance, detailed temporal constraints and the level of quality desired. This leads to a layered view of the use of GPGP [75], as pictured in Figure 16, where other non-GPGP coordination mechanisms would be embedded in the domain problem-solving component of the agent.

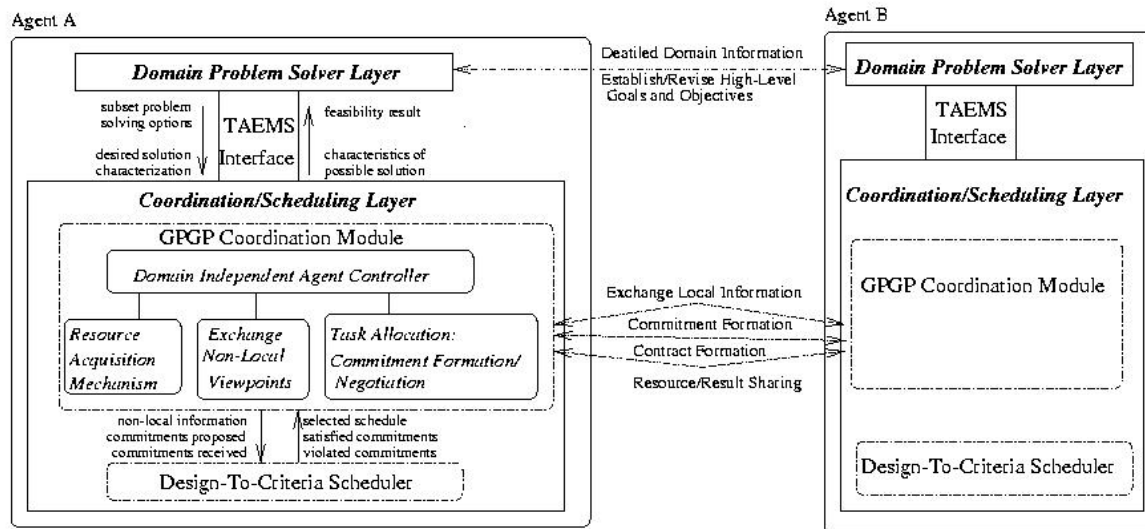


Figure 16: The Layered Use of GPGP

More Complex View of Commitments

The use of a more complex representation of commitments is motivated by two disparate issues. One was the development of coordination protocols that involved resource sharing among activities. (In the original GPGP, protocols only involved information-sharing relationships among agents.) The other was the desire to exploit more fully the knowledge about contingent/uncertain behavior that could now be represented in the extended version of TÆMS. This new knowledge involved the specification of distributions on duration, quality and cost behaviors of activities and the specification of alternative outcomes of activities.

Commitments for Resource Sharing

In the process of developing new mechanisms in GPGP to coordinate over shared resources, a new class of commitments called *don't* commitments have been introduced. They indicate to the local scheduler that a specific task cannot be executed during a certain time interval. Through the use of this type of commitment, a coordination mechanism can serialize access to a resource by a group of agents, or restrict the number of agents that will access the resource during a specific time interval. For example, we have defined a new coordination mechanism that uses a simple negotiation protocol to schedule access to any such mutually exclusive resource [14]. The idea behind this resource-constraint coordination mechanism is that when an agent intends to execute a resource-constrained task, it sends a bid of the time interval it needs and the local priority (projected utility) of its task. After a communication delay, it knows all the bids given out by the other agents at the same time as its own bid. Since all the agents who bid will have the same information about others,

if they all use the same commonly accepted rule to decide who will get the time interval, they can get the same result on this round of bidding. The agent who wins will keep its schedule and execute that task within its time interval, and everyone else will mark this time interval with a *don't* commitment, and never try to execute the resource-constrained task in it unless the owner gives it up. All the agents who didn't get their time intervals at the first round will recompute their schedules and bid again⁵³.

We have also explored another type of commitment for use in a distributed resource allocation problem [26, 43]. This commitment, called a *periodic* commitment, allows us to represent commitments that reoccur at some fixed time interval. The use of this type of commitment reduces coordination costs by incurring the overhead of coordination only for the first time the commitment is established.

Uncertainty and the Dynamics of Commitments

Although commitments are often regarded as “contracts” that have a binding effect, it would be a misconception if commitments were treated as static objects. In fact, commitments are highly dynamic objects. Commitments are built on top of various sources of uncertainties in the agent problem-solving activities, therefore they are inherently uncertain. Let us look at the issue of de-commit, for example. De-commit often occurs for one of these reasons:

1. “statistically unlucky.” This refers to the fact that the tasks often have uncertain, sometimes undesirable outcomes (such as failure). In this case, if the task being promised in the commitment simply fails, the commitment fails simply because the agent does not have luck on its side for its current problem-solving episode.
2. “infeasible.” Sometimes, the task being promised in the commitment may depend on a number of conditions in order to be ready to run. These conditions include precondition tasks, and sufficient computational and non-computational resources needed for its execution. If one or more of these conditions are not satisfied, it is infeasible for the agent to execute the task successfully.
3. “undesirable.” It is also possible that although the agent can manage to satisfy the conditions for the commitment to be kept, it decides not to do so because keeping the commitment is not in the agent's current best interest; for example, the utility gain associated with the commitment has become less than the gain the agent may receive if choosing a different action due to for instance a new task of high utility arriving at the agent after it made the original commitment.

The above discussion is focused on the agent who offers the commitment. Similarly, the agent who requested the commitment may find the commitment unnecessary if the conditions for its local use of the commitment are no longer valid, or the commitment offers no gain since the agent has decided on a more advantageous use of its resources. In these cases, the receiving agent may also

⁵³ We have also developed a resource coordination protocol [39, 40] that uses a different approach that involves the agent owning the resource to have a statistical model of the current reservations for this resource [58]. In this protocol, each resource request contains a priority that denotes the importance of the task for which the resource is being requested. Priority ranges are associated with the different tasks and these reflect the client's preferences. Conflicts occur when resources for a particular time period are insufficient to meet the demand. When this happens, agents holding resource confirmations for the time slot may be requested to release their reservation or may simply be told that their reservation is canceled. The affected agents can then decide to raise their priority and negate the release request or they can simply release the resource.

prefer to retract its request for the commitment, and hence its need to de-commit. If we can predict the likelihood of these changes in the desirability of commitments, which is now possible with representation of outcomes and distributions in TÆMS, the question then arises of how to exploit this information so that agents can coordinate their activities more effectively.

As an example, let us study a commitment offered at time 0, indicating that agent *X* would complete task *A* for agent *Y* before time 160. The following are some scenarios that may occur due to the uncertainty and dynamics of problem solving:

Scenario 1. In typical commitment semantics, this commitment means that *Y* can expect to receive information about the completion of *A* by time 160. However, if task *A* has a 20% chance of failure, *Y* can only receive the utility gain 80% of the time. Furthermore, it means that *Y* needs to make contingency plans after time 160 when *A* fails. Thus, if *Y* knows about the uncertainty about this commitment, it may evaluate the commitment differently at time 0, and may have a better picture of its own overall utility.

Scenario 2. Continuing with the change added to the commitment, let us assume that task *A* actually depends on a preceding task *B*, which has a success rate of 70%, and *B* finishes at time 80 according to the schedule of *X*. Immediately we now know that 30% of the time *A* becomes infeasible. For *Y*, it would mean that the commitment has only a 56% chance of being delivered. This is another source of uncertainty for *Y*. New information would be added to the semantics of the commitment to indicate that there is a 30% chance of de-commit from *X*. More importantly, the new semantics should point out that the decision to de-commit or not will arrive at time 80. Thus, for *Y*, it means that another contingency schedule needs to be made at time 80. Given this information, agent *Y* will now have a more complete picture of the commitment at time 0.

Scenario 3. Again, following the change made in the previous scenario, we will address the issue of a commitment becoming undesirable or unnecessary. Here we need to introduce the concept of marginal gain (marginal loss). Namely, marginal gain (loss) is the difference between the agent's expected utility with and without receiving (offering) the commitment. Clearly, we need to take into account the uncertainty associated with the commitment in order to evaluate the expected utility with the commitment. For cooperative agents, the agents need to make sure that the sum of marginal gain and loss is positive, so that the total expected utility increases with the commitment. As such, if the value of the commitment decreases (in terms of the sum of marginal gain and loss), it is desirable for the agents to agree to de-commit, although the commitment is otherwise satisfiable. For example, let us assume that the task *B* has two outcomes when it does not fail: 40% of the time the outcome quality is 5, and 30% of the time the quality is 2. Also, task *A*'s outcome quality depends on *B*'s outcome: *A* will have quality 6 if *B* has quality 5, and 3 if *B* has quality 2. As a result, the potential marginal gain in *Y* when *A*'s quality is 3 may be much less when *A*'s quality is 6, if *Y*'s reward depends on the quality of *A*'s outcome. In this case, agents might want to de-commit at time 80 when *B*'s outcome is 2 — if the total expected reward would be better for the agents to forgo the commitment and choose some other alternative. As such, it would be beneficial if the agent provides marginal gain/loss information in the commitment, so that the agent would be able to know the value of the commitment to the other agent. Then it could decide if the commitment is desirable/necessary, thereby allowing agents to make contingency plans earlier in the planning process.

In summary, we need to incorporate these three types of uncertainty information in the commitments, and treat commitments as dynamic objects in order to fully understand and evaluate the role of commitments, which serve as the key to agent coordination. In [66] the need for more complex coordination protocols based on uncertainty of commitments is more fully addressed together with implementation and experimental details.

New and More Complex Coordination Mechanisms

The development of new and more complex mechanisms for establishing commitments among agents evolved out of a desire to provide a more sophisticated and parameterized set of coordination protocols that offer a wider range of trade-offs between the overhead of coordination protocols and the optimality of the resulting coordination. In the original GPGP, the coordination mechanisms were basically single-shot in that there was no negotiation process among agents. A commitment was established, or not, based on the requesting agent's local view of reasonableness, and the recipient agent's local view of whether the requested commitment was reasonable for it to commit to⁵⁴. There was no idea of a cooperative search among agents to find an acceptable (or for that matter, best) commitment among agents based on some more encompassing view of the effect of the commitment on both agents. We begin this discussion by first examining how task allocation/contracting can be done in GPGP.

Contracting in GPGP

In the original conception of GPGP, all coordination decisions were based on agents already intending to perform certain activities as a result of the local goals that they were currently pursuing. There was no way for one agent to get another agent to generate results for it if that agent was not already intending to generate those results. Obviously, for many applications, agents may be structured so that they can pursue a large number of different types of local activities and only when there is a specific need for the result of an activity will the agent want to execute that activity. In order to handle the coordination that arises out of this contracting model of agent interaction, we have introduced a new coordination mechanism with associated additional information in the TÆMS task structure. The development and character of this mechanism is consistent with the work described above on the use of default knowledge and situation-specific control. It is also interesting to note that this new mechanism is valuable even in applications which use bottom-up, data-directed processing strategies such as distributed situation assessment.

The new contracting mechanism models the *potential* for external interaction using “virtual tasks.” A task structure is created at an agent with a “possible-task” or “virtual task” or “partially instantiated” task that abstractly describes results that could be generated by another agent. These results are described in terms of a goal that another agent could satisfy. Virtual tasks are related to an agent's task structures via *non-local-effect* (*nle*) (e.g., *enables*, *facilitates*, and so on). If an agent wishes to take advantage of this *nle*, it generates a goal, describing its desired outcome, and transmits it to the appropriate agent. The receiving agent then instantiates a task structure, based on the characteristics of the goal. The actual question of whether the receiving agent is capable of or interested in performing the task is handled by that agent's domain problem-solving component.

The purpose of the virtual task is to serve as a placeholder in scheduling. An agent can attempt to develop schedules including the contributions of the virtual task to determine whether it can indeed develop feasible schedules were that task to be executed. This can take place before the actual communication with the remote agent. It is the modeling of the utility of the virtual task using the discrete probability distribution features of the TÆMS framework that distinguishes the GPGP contracting mechanism from a conventional goal-driven system. The incorporation of the TÆMS

⁵⁴ For example if task *A* in agent 1 *enabled* a task *B* in agent 2, then agent 1 would initiate the establishment of a commitment. Agent 1 in deciding when to do task *A* would call its local scheduler to find out what was the earliest time that task *A* could be done while still achieving reasonable utility for its other local activities. The results of the scheduling process would then define the completion time of task *A* in the proposed commitment that was transmitted to agent 2. This proposed commitment by agent 1 to complete task *A* by a specified time could be either accepted or rejected by agent 2 depending on whether it was still worthwhile for agent 2 to execute task *B* given that its earliest start time was the completion time of task *A* plus some time for message transmission.

relationships allows the system to model in advance the contributions of the actions of the remote agent, to reason about the nature and magnitude of the coordination relationships required, and to parameterize the request according to this information, as well as deadlines for transmitting the information represented by the coordination relationships.

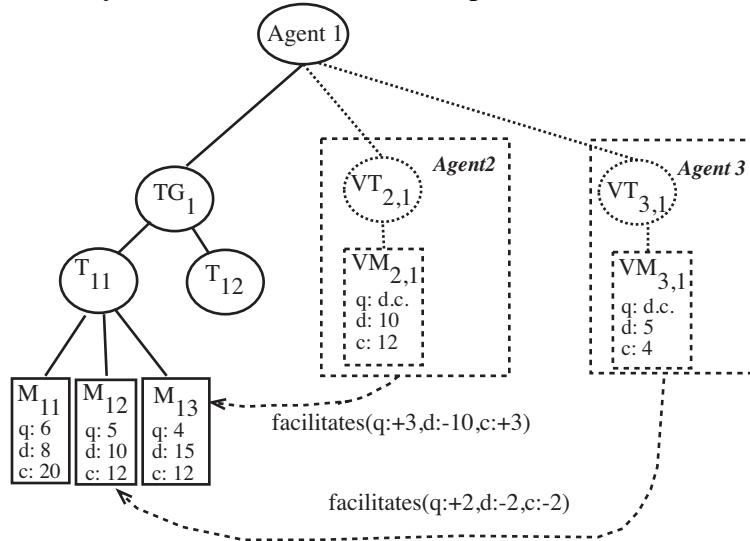


Figure 17: Example of agent reasoning about how best to contract for needed information from another agent

Consider the idealized situation depicted in Figure 17. *Agent 1* has a choice of three methods for achieving task T_{11} . Two of these methods can be facilitated by the actions of other agents. Unlike the original GPGP, there is no presupposition that the tasks will actually be performed unless the agent specifically requests them. Each local method has different performance characteristics. Method M_{11} , for example, has a potential quality of 6, a duration of 8, and a cost of 20. Based purely on a local perspective, and weighing quality of solution most highly, this would be the method of choice. All other methods produce lower quality solutions, although with lower costs and/or durations. However, by understanding in detail how other agents' activities may interact with its own, the agent is able to predict that, were other agents to execute the appropriate tasks, methods M_{12} and M_{13} would be facilitated with quality/cost/duration tuples of (7,8,10) and (7,5,15) respectively. Knowing that higher quality can be achieved, *Agent 1* can now reason about the coordination issues involved in ensuring that another agent executes the appropriate task in a timely fashion. First, it can use secondary criteria such as agent and duration to choose which task should be executed by a remote agent. Duration is of consideration both in scheduling later activities and in ensuring that sufficient time exists for a request to be transmitted, processed, and for a response to be sent. The agent may or may not have information regarding the cost of executing the tasks by the remote agents and the duration of these tasks— information that might be useful in determining *which* agent to ask for assistance.

Once a remote agent and remote task have been selected, the agent must transmit a request that the task be executed. This request can be annotated with information provided by the local scheduler, indicating the desired parameters under which the task should be executed, including but not limited to the quality, cost, and duration modifiers of the transmitted information, and the preferred finish time of the task. We next describe how this single-shot task allocation coordination mechanism can be extended to allow for a cooperative search process among agents to find an appropriate commitment.

Multi-Step Coordination Protocols

Before discussing how to structure a cooperative search process among agents, it is crucial to understand the criteria/objective being maximized in the search process; the criteria chosen is very strongly related to the scope and nature (different degrees) of cooperation among agents. Having appropriate criteria allows agents to reason more accurately about how a specific task allocation commitment will affect overall system performance. The most encompassing form of cooperation is what has been called “global cooperation” [34], which occurs when an agent, while making its local decision, always tries to maximize the global utility function that takes into account the activities of all agents in the system. Global cooperation is unachievable in most realistic situations because of the number of agents and bounds on computational power and bandwidth. Thus we focus our approach on organizing a cooperative distributed search among agents to find an acceptable commitment based on “local cooperation” [34]. This type of cooperation occurs when two or more agents, while negotiating over an issue, try to find a solution that increases the sum of their local utilities, without taking into account the rest of the agents in the system. As mentioned briefly in the section on coordination based on commitment uncertainty, we introduce the notion of marginal utility gain and cost (loss) as a way of computing the effect of a specific commitment on the two agents involved in the negotiation.

In [73], we developed a multi-dimensional, multi-step negotiation mechanism for task allocation that uses marginal utility gain and marginal utility cost to structure this search process to find a solution that maximizes the agents’ combined utility⁵⁵. These two utility values together with temporal constraints generated as a result of an agent’s local search process summarize the agents’ local information used in the negotiation process, and thus reduce the communication bandwidth necessary to implement the negotiation. This is a multi-step negotiation process since agents engage in a series of proposals and counter-offers to decide whether the contractee agent will perform a task for the contractor agent. The proposals and counter-offers specify ranges of acceptable times for what is the earliest time a contracted task can be started and when it needs to be completed. Furthermore, this negotiation is over multiple dimensions (by a specified time with a certain minimum quality) rather than over a single dimension. For example, agent *A* wants agent *B* to do task *T* for it by time 10 and starting it no earlier than time 3, and requests the minimum quality of 8 for the task to be achieved. Agent *B* replies that it can do task *T* by time 10 but only with the quality of 6; however, if agent *A* can wait until time 15, it can get a quality of 12. Agent *A* will select the alternative it believes is better for both agents. Since the negotiation relates to both the completion time and achieved quality of the task, the scope of the search space for the negotiation is increased and thus improves the agents’ chance of finding a solution that increases the combined utility.

This mechanism is *anytime* in character: by investing more time, the agents increase the likelihood of getting a better solution. A parameterized version of this coordination mechanism allows control over the amount of search done among agents—from a single-shot mechanism, all the way to an unlimited search for the best solution. As part of this effort, a measure of negotiation complexity that is based on local measures of agent task complexity was developed that can be used by an agent to choose the appropriate protocol through the exchange of limited meta-level information. This measure allows agents to explicitly balance the gain from the negotiation and the resource usage of the negotiation⁵⁶.

⁵⁵ A very similar cooperative search process, which we have implemented, can also be used to establish coordination for non-local relationships such as *enables* and *facilitates*. It is also our suspicion that a variant of this same search process can be used to implement coordination over *qaf*s such as *min*, *max* and *sum*.

⁵⁶ An interesting research question that we are just beginning to address is how this cooperative search process can be extended to more than two agents. The need for this type of a more encompassing search process occurs in a variety of different forms of multi-linking of activities among a set of agents [74].

Conclusions and Future Directions

The experiences that we have had in applying GPGP/TÆMS to a number of different applications and the ability to extend GPGP/TÆMS as described in this paper have led us to believe that our basic approach to coordination, based on quantitative coordination relationships represented in TÆMS and the generation of commitments among agents that then constrain local agent scheduling, is a powerful and general framework for implementing sophisticated, domain-independent coordination mechanisms. Further, GPGP/TÆMS can be naturally extended to be highly situation-specific where the overhead for coordination can be adjusted for the specific coordination situation. This can be accomplished by substituting *a priori* knowledge for dynamically acquired knowledge and dynamically generated commitments and the use of a “conditioned” TÆMS task structure. Additionally, GPGP has been able to be easily extended to allow the implementation of more top-down and contracting types of coordination mechanisms. These extensions are what make GPGP applicable to a wider range of multi-agent domains that require both top-down and bottom-up coordination regimes and more easily integratable into a multi-layered control architecture in which organizational design rules and/or domain-specific coordination protocols define policies for multi-agent coordination [37].

Our plans for further extending GPGP fall into the following major categories: (1) to explore the interplay between local, environment-centered coordination mechanisms and coordination of large agent organizations; (2) to expand the range of possible coordination mechanisms, including those that might make sense only in a large organizational context, (3) to experiment with different coordination protocols and determine their contextually dependent efficacy, (4) to make GPGP operate in a soft real-time environment where the true costs of coordination are factored into dynamically made decisions about what and how much effort to put into coordination, and when, (5) to allow for multiple coordination mechanisms to operate concurrently, (6) to explore the interplay between domain specific coordination policies and other approaches to coordination with GPGP, and (7) to connect the GPGP approach to coordination with more formal models of coordination.

An important emphasis in further extending GPGP is to allow it to operate effectively in large agent organizations involving hundreds of agents. One aspect of our approach to these organizational extensions to GPGP is to view an organization as something that provides limits, constraints, and other bounds on the local agent coordination process (i.e., making coordination situation-specific) [5, 6]. An organizational approach might strictly limit or pre-specify the specific relationships that are “important” to coordinate over. Organizational limits might also extend to how various commitments (including pre-established ones) or their associated tasks should be valued. Organizational information might pre-specify certain relationships, or allow for the partial specification of relationships. Finally, the organization might bound the total amount of effort to spend on the meta-task of coordination reasoning and control. A key question then arises of whether we should represent these organizational directives as part of an extended version of the TÆMS representation or develop a different representation for this type of organizational knowledge. The other aspect of our organizationally centered future work is the integration of organizational roles and knowledge structures into an agent’s decision process. Our work on the MQ framework [64, 67] represents a step in this direction. In this work, the value of local tasks, coordination actions, and commitment satisfaction (or decommitment) is evaluated using a unified framework, and the values of these items are influenced by their relationship to organizations (and the agent’s relationship to said organizations).

As part of this work on making GPGP more appropriate for use in an organizational context, we are also trying to extend GPGP’s coordination mechanisms so that they will work in an agent organization where all agents are not totally cooperative. We want to explore that introduction of mechanisms to support coordination among self-interested agents as well as cooperative agents

[74]. Additionally, as alluded to in our introduction, we feel that the semantics of GPGP mechanisms need to be augmented to more fully operationalize the semantics of the joint-intention framework. This will also require us to develop mechanisms for implementing the relationships among different agent tasks that occur because of quality-accumulation-function relationships. As discussed previously, we see this as a complex distributed search process that can span multiple agents. Likewise, the GPGP mechanisms have been oriented around coordinated interactions between two agents. The question then arises: Do we need to develop special mechanisms that deal explicitly with multi-linked interactions among agents that span more than two agents? Our intuition is that, in fact, there is a definite need for more complex mechanisms to deal with these multi-linked interactions and the need for these mechanisms will begin to surface as we build more complex multi-agent systems in the future. We also see the need for more contingency analysis and explicit reasoning about uncertainty as part of the decision process in establishing commitments among agents. The work we have done on commitment uncertainty [70], described previously, is only a first step in understanding how to effectively deal with uncertainty inherent in coordination.

Experimental work continues to be an important emphasis in our research orientation. This is especially important to be undertaken now before the next generation of sophisticated and large multi-agent applications are constructed. From the perspective of GPGP, we want to understand the context within which certain approaches to coordination are effective. GPGP is a heuristic approach to coordination and makes numerous assumptions about the level of satisficing coordination that is acceptable. We have, as a result of the extensions to GPGP described here, gone a long way in being able to control the degree of satisficing. What we really don't have is a theory about the nature of multi-agent interactions that would tell us the appropriate level of satisficing given the characteristic of a particular application⁵⁷. This is a long-term goal of our research but at this stage of our understanding, we feel it is still very important to generate more empirical studies that will provide the intuitions about the necessary components of such a formal model that can predict the performance of specific types of coordination for a given application domain.

Another important future direction in the evolution of GPGP, closely allied to the previous discussion, is to make GPGP operate in a soft real-time and dynamically evolving environment. There are many requirements necessary for this to occur. Probably most important is the ability to represent and reason about the interleaving (at times concurrently) of domain problem-solving activities and coordination (and control) activities which in part involve being able to make trade-offs between spending resources on either coordination or domain activities. We feel that this meta-level control for multi-agent systems is a very important issue that is just now beginning to be addressed [49]. A necessary aspect of this approach is the existence of a family of coordination protocols that trade off the level of coordination optimality for the end-to-end delay in establishing coordination and also the amount of computing and communication resources required. We also foresee that numerous heuristics will need to be employed in the agent architecture to reduce the need for re-coordination in case of unexpected events. In recent work of ours [58], we have developed the beginning of such an agent architecture that hopefully can be extended to support a soft real-time version of GPGP. More generally, even if the costs of re-coordination can be decreased they will often still be significant which implies that meta-level control reasoning needs to deal with not only the cost of the initial coordination but also potential subsequent costs that will arise out of re-coordination.

Another area for extension, which is also closely allied to developing a GPGP that operates in a soft real-time environment, centers around the problems that occur when multiple sessions of coordination protocols are concurrently executing; this problem can occur since there may be more

⁵⁷ Based on both older work of ours [8] as well as more recent work [71], it seems that the level of uncertainty in the behavior of agents should be explicitly modeled in such a theory. We have also observed many other factors such as the relative utility among different goals, the loading of agents, communication bandwidth, etc., that relate to the effectiveness of different coordination protocols [10, 11, 45].

than one subtask at any one time that requires coordination commitments to be established. Additionally, as we have observed⁵⁸ in our recent work on a distributed situation assessment application [26, 43], the coordination protocol session may need to be terminated prematurely because the domain problem-solving component has changed the domain tasks it now wants executed, or the results of another concurrently executing coordination session invalidates the current resource context being using by this session. This leads to the need for being able to recognize and resolve complex interactions occurring among concurrently executing coordination sessions. We made an initial attempt at solving this problem in earlier work [11, 65] but the framework we established for building protocols was not rich enough to handle the complex interactions that we now understand can occur among coordination sessions. Creating a richer framework for building coordination protocols will be an important focus of future work. We have also done work on separating resource contexts used for concurrent coordination sessions so that they are mostly independent [74], which does not solve the problem but reduces the frequency of interaction among sessions.

In order to make GPGP more generally applicable, we feel that it is important for it to effectively inter-operate with other coordination approaches. As described in Figure 16, we view GPGP as being in a layered relationship with other non-GPGP coordination protocols. The key problem in building such a layered approach is determining how information will be exchanged among the layers and the type of directives that non-GPGP coordination protocols can give to GPGP coordination mechanisms. Research described in [1, 65] in terms of monitoring conditions give us some clues about how to approach this interface problem. We will probably need to monitor the intermediate states of the coordination protocol, as well as the performance and resource usage of tasks. This interface issue is also critical to how an organization layer of control can be integrated with GPGP.

Our final research direction concerns how to relate the GPGP approach to coordination to more formal models of coordination. There are two basic questions that we would like to understand from a more theoretical perspective. One question is to understand for a particular coordination situation how difficult this situation is in terms of the complexity of coordination/communication policy necessary to reach a specific level of coherent agent behavior; we see answering this question in terms of the generated utility of a “satisficing” policy versus the maximum utility of an optimal policy as related to the amount of communication necessary to implement the “satisficing” policy. A second closely related question is to understand for a particular situation how much distributed search is necessary to find a “satisficing” coordination policy. We feel that distributed MDPs provide a rich theoretical framework [4] that may be the right one to understand GPGP since it deals explicitly with uncertainty of agent activities, though it has problems modeling communication costs [3]. We have already begun to investigate the connection between MDPs and TÆMS and our local agent scheduler [48], and have also been using MDPs to study coordination [71], and are currently extending this work to relate more directly to GPGP coordination [72].

In closing, it would be remiss of us not to mention the significance that we attach to learning. As we build multi-agent systems that operate in open and evolving environments, the system must be able to adapt its heuristics/assumptions (and there are many in GPGP), to the current structure of the environment—not some predefined view of the environment. We have done some work in this area [27, 31, 44] but we need to do a lot more. Especially interesting is building up, either implicitly or explicitly, models of task arrival so that opportunity costs [49, 67] can be factored into coordination decisions at both a protocol level and a meta-level.

⁵⁸ Private communication with R. Mailler, the developer of a distributed resource allocation protocol for this domain.

Acknowledgements

As is obvious by the author list, the development and evolution of GPGP has been a long process over the last ten years that has been the focus of much of the intellectual directions of the Multi-Agent System Laboratory at University of Massachusetts at Amherst. Where appropriate, we have referenced papers by specific members of the laboratory who took major responsibility in developing a research idea explained in this paper. However, it should go without saying that there were many interchanges with, and building on ideas developed by, other laboratory members who are not directly recorded in article citations. Though they did not directly contribute to the development of GPGP, we would also like to acknowledge some long-term visitors to our laboratory, other laboratory members, and associated faculty who were part of the intellectual ferment that generated the ideas expressed here. They are: Michael Atighetchi, Ana Bazzan, Brett Benyo, Satoru Fujita, David Jensen, Qiegang Long, Roger Mailler, Lee Osterweil, Tuomas Sandholm, Jiaying Shen, Toshiharu Sugawara, Robert Whitehair and Shlomo Zilberstein. We would also like to thank Sherief Abdallah for his thoughtful comments on this paper. Finally, it almost goes without saying the intellectual debt this paper owes to the pioneering work of our early laboratory members, Daniel Corkill and Edmund Durfee.

REFERENCES

1. S. Abdallah, et al. "Monitoring and Synchronization for Teamwork," *Proceedings of the 17th ACM Symposium on Applied Computing (SAC2002)*, 2002. (To appear.)
2. M. Barbuceanu and M. S. Fox, "COOL: A Language for Describing Coordination in Multi-Agent Systems," in *Proceedings of the First International Conference on Multi-Agent Systems*, Menlo Park, CA: AAAI Press, pp. 17-24, 1995.
3. D.S. Bernstein, S. Zilberstein, and N. Immerman, "The Complexity of Decentralized Control of Markov Decision Processes," *Proceedings of the 16th International Conference on Uncertainty in Artificial Intelligence*, Stanford, CA, pp. 32-37, 2000.
4. C. Boutilier, "Sequential Optimality and Coordination in Multiagent Systems," *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI99)*, 1999.
5. D.D. Corkill and V.R. Lesser, "The Use of Meta-Level Control for Coordination in a Distributed Problem-solving Network," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pp. 748-756, 1983.
6. Daniel Corkill, "A Framework for Organizational Self-Design in Distributed Problem-Solving Networks," Ph.D. thesis, University of Massachusetts/Amherst, 1983.
7. K. Decker and V. Lesser, "Generalizing The Partial Global Planning Algorithm," *International Journal on Intelligent Cooperative Information Systems*, 1(2):319-346, 1992.
8. K. Decker and V. Lesser, "An Approach to Analyzing the Need for Meta-Level Communication," *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, Vol. 1, 1993.
9. K. Decker and V. Lesser, "Quantitative Modeling of Complex Environments," *International Journal of Intelligent Systems in Accounting, Finance, and Management*, 2(4):215-234, Special issue on Mathematical and Computational Models of Organizations: Models and Characteristics of Agent Behavior, 1993.
10. K. Decker and V. Lesser, "Designing a Family of Coordination Algorithms," *Proceedings of the First International Conference on Multi-Agent Systems*, San Francisco: AAAI Press, pp. 73-80, 1995.

11. Keith S. Decker, "Environment Centered Analysis and Design of Coordination Mechanisms," Ph.D. thesis, University of Massachusetts/Amherst, 1995.
12. Keith S. Decker, "Task Environment Centered Simulation," in *Simulating Organizations: Computational Models of Institutions and Groups* (M. Prietula, K. Carley, and L. Gasser, editors), AAAI Press/MIT Press, 1997.
13. K. Decker, M. Williamson, and K. Sycara, "Intelligent Adaptive Information Agents," *Journal of Intelligent Information Systems*, 9:239-260, 1997.
14. K. Decker and J. Li, "Coordinating Mutually Exclusive Resources Using GPGP," *Autonomous Agents and Multi-Agent Systems*, Special Issue: Best of ICMAS'98 – Part II, Volume 3(2): 133–158, 2000.
15. E.H. Durfee and V.R. Lesser, "Negotiating Task Decomposition and Allocation Using Partial Global Planning," *Distributed Artificial Intelligence* (M. Huhns and L. Gasser, eds.), Vol. 2, Pitman Publishing Ltd.: London, pp. 229–244, 1989.
16. 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, 1991.
17. S. Fujita and V.R. Lesser, "Centralized Task Distribution in the Presence of Uncertainty and Time Deadlines," *Proceedings of the Second International Conference on Multi-Agent Systems*, AAAI Press: CA, pp. 87–94, 1996.
18. A. Garvey, K.S. Decker and V.R. Lesser, "A Negotiation-based Interface Between a Real-time Scheduler and a Decision-Maker," *Proceedings of Workshop on Models of Conflict Management in Cooperative Problem Solving*, Seattle: AAAI Press, 1994.
19. A. Garvey and V.R. Lesser, "Design-to-time Scheduling and Anytime Algorithms," *SIGART Bulletin*, Vol. 7, No. 3, 1996.
20. J. Graham and K.S. Decker, "Towards a Distributed, Environment-Centered Agent Framework," in *Intelligent Agents VI: Agent Theories, Architectures, and Languages* (N. Jennings and Y. Lesperance, editors), LNAI #1757, Springer, pp. 290–304, 2000.
21. J. Graham, D. McHugh, M. Mersic, F. McGeary, M. Windley, D. Cleaver, K.S. Decker, "Tools for Developing and Monitoring Agents in Distributed Multi-Agent Systems," in *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems. Lecture Notes in Computer Science #1887*, (T. Wagner and O. Rana, eds.), Springer, 2001.
22. J. Graham, "Real-Time Scheduling in Distributed Multi-Agent Systems," Ph.D. Dissertation, University of Delaware, January 2001.
23. J. Graham, K.S. Decker, and M. Mersic, "DECAF: A Flexible Multi-Agent System Architecture," *Autonomous Agents and Multi-Agent Systems*. Accepted for publication.
24. B.J. Grosz and S. Kraus, "Collaborative Plans for Complex Group Action," *Artificial Intelligence*, 86(2): 269–357, 1996.
25. Bryan Horling, et al. "The TÆMS White Paper," Technical Notes of Multi-Agent Systems Lab, Department of Computer Science, University of Massachusetts, 1999.
26. B. Horling, R. Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins, V.R. Lesser, "Distributed Sensor Network for Real-Time Tracking," *Proceedings of the Fifth International Conference on Autonomous Agents*, Montreal, pp. 417–424, 2001.
27. B. Horling, B. Benyo, V.R. Lesser, "Using Self-Diagnosis to Adapt Organizational Structures," *Proceedings of the Fifth International Conference on Autonomous Agents*, Montreal, pp. 529–536, 2001.

28. F.F. Ingrand, M.P. Georgeff, A.S. Rao, "An Architecture for Real-Time Reasoning and System Control," *IEEE Expert*, 7(6), 1992.
29. N. Jennings, "Commitments and Conventions: The foundation of coordination in multi-agent systems," *The Knowledge Engineering Review*, Vol. 8(3), pp. 223–250, 1993.
30. N. Jennings, "Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems Using Joint Intentions," *Artificial Intelligence*, Vol. 75, No. 2, 1995.
31. D. Jensen, M. Atighetchi, R. Vincent, V.R. Lesser, "Learning Quantitative Knowledge for Multiagent Coordination," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, 1999.
32. K. Kuwabara, T. Ishida, N. Osato, "AgentTalk: Coordination Protocol Description for Multiagent Systems," in *Proceedings of the First International Conference on Multi-Agent Systems*, Menlo Park, CA: AAAI Press, pp. 455, 1995.
33. B. L  asri, H. L  asri, S. Lander, and V.R. Lesser, "A Generic Model for Intelligent Negotiating Agents," *International Journal on Intelligent Cooperative Information Systems*, 1(2): 291-317, 1992.
34. S. Lander and V.R. Lesser, "Sharing Meta-Information to Guide Cooperative Search among Heterogeneous Reusable Agents," in *IEEE Transactions on Knowledge and Data Engineering*, 9(2), pp. 193-208, 1997.
35. V.R. Lesser and D.D. Corkill, "The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem-solving Networks," *AI Magazine*, 4(3), pp. 15–33, 1983.
36. V.R. Lesser (1991). "A Retrospective View of FA/C Distributed Problem Solving," *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6): 1347–1362, 1983.
37. V.R. Lesser, "Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture," in *Autonomous Agents and Multi-Agent Systems*, Volume 1, pp. 89-111, Kluwer Academic Publishers, 1998.
38. V.R. Lesser, K.S. Decker, N. Carver, A. Garvey, D. Neiman, M. Nagendra Prasad, and T. Wagner, "Evolution of the GPGP Domain-Independent Coordination Framework," in University of Massachusetts/Amherst CMPSCI Technical Report 98-05, 1998.
39. V.R. Lesser, M. Atighetchi, B. Benyo, B. Horling, A. Raja, R. Vincent, T. Wagner, P. Xuan, and S.X.Q. Zhang, "The Intelligent Home Testbed," *Proceedings of the Autonomy Control Software Workshop (Autonomous Agent Workshop)*, 1999.
40. V.R. Lesser, M. Atighetchi, B. Benyo, B. Horling, A. Raja, R. Vincent, T. Wagner, P. Xuan, and S.X.Q. Zhang, "A Multi-Agent System for Intelligent Environment Control." *Proceedings of the Third International Conference on Autonomous Agents*, ACM Press: New York, NY, pp. 291-298, 1999.
41. V.R. Lesser, B. Horling, F. Klassner, A. Raja, T. Wagner, S.X.Q. Zhang, "BIG: An Agent for Resource-Bounded Information Gathering and Decision Making," in *Artificial Intelligence, Special Issue on Internet Information Agents*, 118(1-2), Elsevier Science, pp. 197-244, 2000.
42. H.J. Levesque, P.R. Cohen, and H. T. Nunes, "On Acting Together," in *Proceedings of the Eighth National Conference on Artificial Intelligence*, pp. 94–99, 1990.
43. R. Mailler, R. Vincent, V.R. Lesser, T. Middelkoop, and J. Shen, "Soft Real-Time, Cooperative Negotiation for Distributed Resource Allocation," in *AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems*, 2001.
44. M.V. NagendraPrasad and V.R. Lesser, "Learning Situation Specific Coordination in Cooperative Multi-Agent Systems," *Autonomous Agents and Multi-Agent Systems*, Vol. 2, Kluwer Academic Publishers, pp. 173-207, 1999.

45. M.V. NagendraPrasad, K.S. Decker, A. Garvey, V.R. Lesser, "Exploring Organizational Designs with TÆMS: A case study of distributed data processing," in *Proceedings of the Second International Conference on Multi-Agent Systems*, 1997.
46. T. Oates, M.V. NagendraPrasad, V.R. Lesser, "Cooperative Information Gathering: A Distributed Problem-Solving Approach," in *IEEE Proceedings on Software Engineering*, Special Issue on Agent-based Systems, Volume 144, No. 1, 1997.
47. D. Pynadath and M. Tambe, "An Automated Teamwork Infrastructure for Heterogeneous Software Agents and Humans," *Journal of Autonomous Agents and Multiagent Systems (JAAMAS)*, 2002. (To appear.)
48. A. Raja, V.R. Lesser, T. Wagner, "Toward Robust Agent Control in Open Environments," in *Proceedings of the Fourth International Conference on Autonomous Agents*, Barcelona, pp. 84-92, 2000.
49. A. Raja and V. Lesser, "Towards Bounded-Rationality in Multi-Agent Systems: Reinforcement-Learning-Based Approach," in University of Massachusetts Computer Science Technical Report, #2001-34, 2001.
50. A. Raja, T. Wagner, V.R. Lesser, "Reasoning about Uncertainty in Agent Control," in *Proceedings of the 5th International Conference on Information Systems, Analysis, and Synthesis*, Computer Science and Engineering: Part 1, Volume VII, pp. 156-161, 2001.
51. A.S. Rao and M.P. Georgeff, "BDI agents: From theory to practice," in *Proceedings of the First International Conference on Multi-Agent Systems*, AAAI Press: San Francisco, pp. 312-319, 1995.
52. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, MIT Press: Cambridge, MA, 1994.
53. Z. Rubinstein, "The Hurrier I Go, The Behinder I Get," in *Proceedings of Artificial Intelligence and Manufacturing*, IJCAI & AAAI, pp. 97-102, 2001.
54. R.G. Smith and R. Davis, "Frameworks for Cooperation in Distributed Problem Solving," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-11, pp. 61-70, 1981.
55. T. Sugawara and V.R. Lesser, "Learning to Improve Coordinated Actions in Cooperative Distributed Problem-Solving Environments," in *Machine Learning*, Vol. 33, Nos. 2/3, Kluwer Academic Publishers, 1998.
56. M. Tambe, "Agent Architectures for Flexible, Practical Teamwork," in *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 1997.
57. M. Tambe, "Towards Flexible Teamwork," *Journal of Artificial Intelligence Research*, Volume 7, pp. 83-124, 1997.
58. R. Vincent, B. Horling, V.R. Lesser, T. Wagner "Implementing Soft Real-Time Agent Control," *Proceedings of the Fifth International Conference on Autonomous Agents*, pp. 355-362, 2001.
59. T. Wagner and V.R. Lesser, "Design-to-Criteria Scheduling for Intermittent Processing," *University of Massachusetts/Amherst Computer Science Technical Report 1996-81*, 1996.
60. T. Wagner, A. Garvey, V.R. Lesser, "Complex Goal Criteria and Its Application in Design-to-Criteria Scheduling," in *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 1997.
61. T. Wagner, A. Garvey, V.R. Lesser, "Leveraging Uncertainty in Design-to-Criteria Scheduling," UMass Computer Science Technical Report 1997-11, 1997.

62. T. Wagner, A. Garvey, V.R. Lesser, "Criteria Directed Task Scheduling," *International Journal of Approximate Processing, Special Issue on Scheduling*, 19:91-118, 1998.
63. T. Wagner, B. Benyo, V. Lesser, A. Raja, P. Xuan, and S. Zhang, "GPGP²: Supporting situation specific coordination protocols," *UMASS Computer Science Technical Report #98-042*, 1998.
64. T. Wagner and V.R. Lesser, "Relating Quantified Motivations for Organizationally Situated Agents," in *Intelligent Agents VI: Agent Theories, Architectures, and Languages* (N.R. Jennings & Y. Lesperance, eds.), Springer Verlag, Volume 1757, pp. 334-349, 2000.
65. T. Wagner, B. Benyo, V.R. Lesser, P. Xuan, "Investigating Interactions Between Agent Conversations and Agent Control Components," in *Issues in Agent Communication*, F. Dignum & M. Greaves (eds.), Springer-Verlag, Vol. 1916, pp. 314-331, 2000.
66. T. Wagner and V.R. Lesser, "Toward Soft Real-Time Agent Control," *Lecture Notes in Computer Science: Infrastructure for Large-Scale Multi-Agent Systems* (Wagner/Rana, editors), Springer-Verlag, 2001.
67. T. Wagner and V.R. Lesser, "Evolving Real-Time Local Agent Control for Large-Scale Multi-Agent Systems," in *Intelligent Agents VIII: Agent Theories, Architectures, and Languages* (J-J. Meyer & M. Tambe, eds.), Springer-Verlag, 2002. (To appear.)
68. M. Wellman and J. Doyle, "Modular Utility Representation for Decision-Theoretic Planning," in *Proceedings of the First International Conference on Artificial Intelligence Planning Systems*, pp. 236-242, 1992.
69. A. Wise, A.G. Cass, B. Staudt-Lerner, E.K. McCall, L.J. Osterweil, and S.M. Sutton, Jr., "Using Little-JIL to Coordinate Agents in Software Engineering," in *Proceedings of the Automated Software Engineering Conference (ASE 2000)*, pp. 155-163, 2000.
70. P. Xuan and V.R. Lesser, "Incorporating Uncertainty in Agent Commitments," in *Intelligent Agents VI: Agent Theories, Architectures, and Languages* (N.R. Jennings & Y. Lesperance, eds.), Springer Verlag, Volume 1757, pp. 57-70, 2000.
71. P. Xuan, V.R. Lesser, and S. Zilberstein, "Communication Decisions in Multi-Agent Cooperation: Model and Experiments," *Proceedings of the Fifth International Conference on Autonomous Agents*, pp. 616-623, 2001.
72. P. Xuan and V.R. Lesser, "Multi-Agent Policies: From Centralized Ones to Decentralized Ones," *University of Massachusetts/Amherst Computer Science Technical Report 2001-48*, 2001. (Also submitted to AAMAS 2002.)
73. X.Q. Zhang, R. Podorozhny, and V.R. Lesser, "Cooperative, MultiStep Negotiation Over a Multi-Dimensional Utility Function," in *Proceedings of the IASTED International Conference, Artificial Intelligence and Soft Computing (ASC 2000)*, IASTED/ACTA Press, pp. 136-142, 2000.
74. X.Q. Zhang, V.R. Lesser, T. Wagner, "A Proposed Approach to Sophisticated Negotiation," *AAAI Fall Symposium on Negotiation Methods for Autonomous Cooperative Systems*, 2001. (See also X.Q. Zhang and V.R. Lesser, *University of Massachusetts/Amherst Computer Science Technical Report 2000-057*, 2000.)
75. X.Q. Zhang, A. Raja, B. Lerner, V.R. Lesser, L.J. Osterweil and T. Wagner, "Integrating High-Level and Detailed Agent Coordination into a Layered Architecture," in *Lecture Notes in Artificial Intelligence 1887: Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems* (Wagner & Rana, eds.), Springer, pp. 72-79, 2001.

76. E. Durfee and V.R. Lesser, "Predictability versus Responsiveness: Coordinating Problem Solvers in Dynamic Domains," *Proceedings of the Seventh National Conference on Artificial Intelligence*, pp. 66-71, 1988.
77. A. Raja, V.R. Lesser, T. Wagner, "Toward Robust Agent Control in Open Environments," in *Proceedings of the Fourth International Conference on Autonomous Agents*, pp. 84-92, 2000.