


**Estimating Information Value in  
Collaborative Multi-Agent Planning  
Systems**

David Sarne and Barbara Grosz

TR-16-06



Computer Science Group  
Harvard University  
Cambridge, Massachusetts

# Estimating Information Value in Collaborative Multi-Agent Planning Systems

David Sarne and Barbara J. Grosz

Division of Engineering and Applied Sciences  
Harvard University, Cambridge MA 02138 USA  
{sarned, grosz}@eecs.harvard.edu

## ABSTRACT

This paper addresses the problem of identifying the value of information held by a teammate on a distributed, multi-agent team. It focuses on a distributed scheduling task in which computer agents support people who are carrying out complex tasks in a dynamic environment. The paper presents a decision-theoretic algorithm for determining the value of information that is potentially relevant to scheduling revisions, but is directly available only to the person and not the computer-agent. The design of a “coordination autonomy” (CA) module within a coordination-manager system provided the empirical setting for this work. By design, the CA module depends on an external scheduler module to determine the specific effect of additional information on overall system performance. The paper describes two methods for reducing the number of queries the CA issues to the scheduler, enabling it to satisfy computational resource constraints placed on it. Experimental results indicate the algorithm improves system performance and establish the exceptional efficiency—measured in terms of the number of queries required for estimating the value of information—that can be achieved by the query-reducing methods.

## 1. INTRODUCTION

Incorporating autonomous-agents capabilities into planning and scheduling systems can significantly increase their power, especially in highly dynamic environments and settings in which multiple agents’ schedules are being coordinated [19; 23, *inter alia*]. Agents can suggest alternative courses of action as well as negotiate conflicts. They can help achieve team objectives more effectively in such highly dynamic environments as first-response situations [22, 12], exploration of remote planets, cleanup of hazardous sites, and military conflicts [10], because they have the ability to efficiently process in real-time the rich set of information affecting the problem. By taking responsibility for changes in task schedules, they can free people involved in doing tasks from being unnecessarily overburdened, especially in times of crisis when plans need to be adapted to new circumstances [12, 10]. Situations characterized by agents being separated geographically; coordinating groups having limited communication bandwidth; and changes in events and physical environment being complex, dynamic and stochastic in nature typically preclude a complete global view of the problem. They thus require coordinated, but localized, planning and scheduling decisions [20, 22]. In the remainder of this paper, we use the term “automated scheduling agent” (ASA) to refer to autonomous-agent systems that support multi-agent coordinated scheduling.

ASAs must take into account various external facts that may influence choice of scheduling changes significantly, including changes in the environment in which tasks are being carried out and the state of the people or other agents working on a distributed, coordinated task. Information beyond the system’s limited knowledge of the environment may affect constraints, either relaxing them or imposing new ones that better reflect the actual situation. For example, a driver will see changes in weather conditions that may af-

fect route selection when they occur, while an automated navigation system does not. A scientist can induce that a colleague will be going to the school play and thus be unavailable to meet simply because their children go to the same school. Human-agent interaction must be managed appropriately for the system to obtain such information without overburdening other participants, especially people [12, 19]. Obtaining system-external information is associated with costs, including communication costs and the degradation in task performance of the person being interrupted [6, 9, 7, 23]. An ASA needs to plan its interactions to maximize the difference between the value of the information being obtained (taking into account the probability the person being queried has this information) and the costs, deploying appropriate decision-making mechanisms [8].

This paper addresses the problem of identifying the value of information held by a teammate on a distributed, multi-agent, mixed human and computer team. We focus on a scheduling task in which the computer agents are ASAs supporting people who are carrying out complex tasks in a dynamic environment. The ASA is responsible for coordinating modifications to the task schedule as the environment changes with as few interruptions of the person as possible, thus enabling the person to focus on actual task performance.

Although coordination itself is domain independent, determining the value of information requires scheduling knowledge; information is valuable to the extent it influences schedule changes, but that influence cannot be determined without task and schedule knowledge. The design of a “coordination autonomy” (CA) module within a coordination-manager system (as part of the DARPA Coordinators project<sup>1</sup>) provided the empirical setting for this work, including both constraints on system design and a testbed environment. The overall coordination-system design separates the CA module from the *scheduler* which is the locus of scheduling-related expertise.<sup>2</sup> As a result, the CA must consider not only interruptions of the person, but also resource demands it makes on the scheduler through querying for schedule-expertise. Thus, a key challenge to CA-module design was to develop methods that would obtain information from the (human) participants in a distributed, coordinated task without interrupting them too frequently or at the wrong times, while simultaneously not consuming too many scheduler, communication or other coordination system’s resources.

<sup>1</sup>The DARPA *Coordinators* research program aims to demonstrate a computational framework for distributed scheduling from which intelligent cognitive software agents can evolve that will help fielded military units adapt their mission plans and schedules as the situation around them changes [21].

<sup>2</sup>This separation of the scheduler allows it to be used by other ASA modules. Furthermore, unlike such user-focused information as estimating the probability a user knows relevant information, the user’s confidence level in the information, and the costs associated user-interactions, which naturally fit in the CA, the scheduling functionality required by the CA is system-focused. Nonetheless, the mechanisms described in this paper are useful even in environments in which a CA would have scheduling expertise, though in such settings the CA could also bring to bear heuristics that took advantage of its knowledge of problem structure.

The paper makes two major contributions. First, it presents a decision-theoretic algorithm for determining the value of information in a multi-agent context in which a person rather than a computer agent has the information in question and in a systems context in which the calculation of information-value is complicated by the need to invoke a separate systems module (in this case, the scheduler).<sup>3</sup> Second, it describes two novel methods for dealing with the computational resource constraints on the CA. The first method filters queries depending on whether there has been a change in “relevant” tasks in the schedule. The second translates the problem of generating queries into a game and then attempts to minimize the score. Experimental results demonstrate the effectiveness of the algorithm and that these two methods increase efficiency by significantly decreasing the number of queries required for estimating the value of information.

To provide context for the empirical results as well as for the methods we designed for determining information value, the next section of the paper briefly describes the *Coordinators* problem and the basic architecture of the CA module. Section 3 discusses the principles on which the process of estimating the value of information is based, and Section 4 presents mechanisms for improving the efficiency of the scheduler-querying process. Experimental results are given in Section 5. Section 6 discusses related work, and Section 7 contains our conclusions.

## 2. IMPLEMENTATION BACKGROUND

The CA module and algorithms we describe in this paper were developed and tested in the “Coordinators” application domain [22]. In this domain, the ASAs, called “Coordinators”, operate in a rapidly changing environment and are intended to help maximize an overall team-objective which is measured using a hierarchical objective function. Each ASA operates on behalf of its *owner* (e.g., the team leader of a first-response team or a unit commander) whose schedule it manages. The actual tasks being scheduled are executed either by owners or by units they oversee, and the ASAs’ responsibility is limited to scheduling these tasks.

The quality of ASAs’ schedules depends on knowledge of the environments in which their owners are operating, and in many cases the owners of ASAs may have more accurate information about the state of the environment. Owners in this domain may acquire relevant external information through various communication channels (formal and informal) they maintain with other owners. Typical sources of such information are occasional coordination meetings (e.g., used for reporting status of task execution), open communication the owner passively listens to (e.g. when leaving the radio open, one gets to listen to messages associated with other teams in the area) and direct communication used for coordination with other people acting on a joint task (throughout which an individual often learns informally about the status and existence of actions being executed by others). An additional source of task-outcome related information is an owner’s accumulated experience.

In this domain, scheduling problems are distributed in nature. Each ASA receives a different, typically only partial, local view of the different tasks and structures that constitute the full multi-agent problem. Thus, the intended context of use of Coordinators precludes central management. Scheduling problems must be solved distributively. As a result, each ASA needs to reason about changes in the timings of tasks with regard to their correlation with other ASAs’ tasks and to adapt its owner’s schedule accordingly. The time pressure of the intended applications (e.g., emergency response) re-

<sup>3</sup>The algorithm focuses on the “benefits” side of the decision making process that determines when to interact with a user. Specifically, we focus on estimating the expected value of the information the user may have. A complete decision-making mechanism would combine the work described in this paper with a method for determining the cost of interrupting the user and the probability the user has the necessary information [17].

quires that ASAs make decisions in real time concurrently with their owner’s (or their units) execution of tasks.

The ASA-owner relationship is a collaborative one, with the ASA needing to interact with its owner to obtain task- and environment information relevant to scheduling. The CA module is responsible for deciding intelligently when and how to interact with the owner for improving the ASA’s distributed scheduling. It interacts primarily with a *meta-cognition* component of the ASA, which controls computational resources allocated to the CA and other ASA modules, and with the scheduler model of the ASA.<sup>4</sup> The scheduler, which is the locus of the ASA’s scheduling expertise, is responsible for task analysis, and thus is a resource whenever scheduling reasoning is required for evaluating the effect of changes in scheduling-problem settings on the system’s expected performance. The CA initiates interaction with the ASA when it requires scheduler resources or after interacting with the owner (for the purpose of transferring the information received to the ASA).

Our basic CA architecture is given in Figure 1. The *user profiling* and *user’s state monitoring* components track the owner’s interruptibility preferences and costs as well as managing interactions with the owner. These owner-interaction focused components are described elsewhere [17, 18]. The *Event Handler* manages interactions with the meta-cognition module and determines (future) times at which there is a high probability the owner will have new information that may affect system performance. For example, an owner is more likely to know the probability of success of a task to be done in the future shortly prior to the task’s execution than further from its execution time. The *Interaction Evaluator* evaluates the benefit of interacting with the owner at a particular time. If the expected benefit is positive, then an interaction with the owner will be initiated and monitored by the *Interaction Handler* component. The mechanism and algorithms presented below enable the *Interaction Evaluator* to efficiently activate the scheduler to estimate the value of a specific piece of task- or environment information the owner has. The *Interaction Evaluator* computes the benefit of interrupting the owner based on a comparison of this latter value multiplied by the (estimated) probability an owner has this specific information with the CA’s estimation of the cost such an interaction will incur.

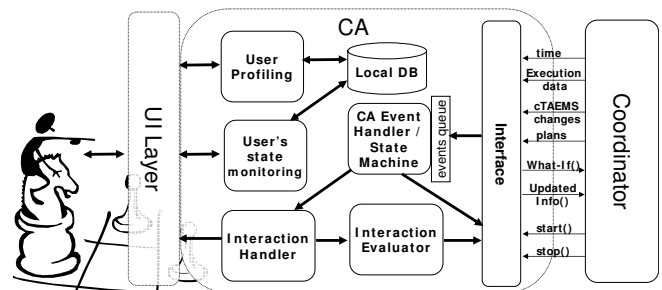


Figure 1: The CA architecture

Coordinators ASAs represent planning and scheduling problems in cTAEMS structures [2] defining multi-agent hierarchical tasks with probabilistic expectations on their outcomes. For this paper the relevant elements of cTAEMS are methods, quality measures and inter-action dependency links, which we describe briefly.

Methods, which are actions executed by a single agent, are the atomic elements in cTAEMS. Figure 2 provides an example of a method definition. As shown, a method has multiple possible discrete outcomes that dictate its duration and quality. Methods are usually constrained by “release times” (earliest possible start) and deadlines that are set by tasks higher in the cTAEMS hierarchy to

<sup>4</sup>Several different architectures have been suggested for ASAs in the Coordinators domain [13, 11, 16, 20]. All of them share the same core set of modules, including the scheduler and meta-cognition components.

which a method belongs, and inherited hierarchically. If a method violates its temporal constraint, it is considered to have failed and yields a zero quality. The probability of zero quality and the duration distributions are given in the “failure” portion of the method definition, as shown in Figure 2. The quality of execution of a method contributes to the quality of its parent task by means of a quality accumulation function (QAF), which defines how the qualities are aggregated up the hierarchy. For example, a  $q_{min}$  QAF for a task results in a quality equal to the minimum achieved by its subtasks. The functions  $q_{max}$ , and  $q_{sum}$  are defined analogously. Only methods that are executed within their constraints accrue quality, and the overall objective is to maximize quality. The non-local effects links (NLEs) in cTAEMS indicate inter-node relationships such as enablement and facilitation.

The hierarchical outcomes of cTAEMS structures may be translated into a simple tabular representation of outcomes and their probability, as given in the table to the right of Figure 2. Each cell in the table contains the probability associated with the outcome described by its row quality and column duration. This tabular presentation is particularly useful when we investigate methods for increasing the CA mechanism’s efficiency in Section 4.




Figure 2: A typical method in cTAEMS

cTAEMS problems are known to be difficult scheduling problems and are not easily modeled and solved by traditional solvers for planning and scheduling [4]. Coordinators’ ASAs have used a variety of techniques for scheduler implementations [4, 11, 20]. The value-of-information methods we present do not depend on the particular scheduler technology. The only functionality they require is the ability to produce a feasible schedule given changes in the cTAEMS problem and to be consistent (i.e., if the same problem arrives again then the same solution is supplied).

### 3. INFORMATION VALUE

In cTAEMS, the outcome of a method after it has been executed (or has failed) represents the contribution of performance of that action (its “quality”) to the overall team effort (after weighing it according to the NLEs and QAFs associated with this method). Because execution outcome is probabilistic, if the CA can obtain a more accurate information about the probability distribution of a method’s outcomes before method execution is completed, it can help to significantly increase solution quality. External environmental information relevant to method outcomes may be helpful in deciding among possible schedules. For instance, if an owner can tell ahead of time that the method its ASA has scheduled to be executed next will definitely fail (e.g., because the weather has changed and will prevent success), early notification of the ASA will allow it to change the schedule before the owner even begins executing this method. If the ASA notifies other owners’ ASAs of this failure, then the schedules of any other owners whose tasks are affected by this method’s outcome (because their methods either are part of common QAFs or are enabled by NLE links originating from this method) also can be changed earlier, rather than waiting for the failure to occur. In the absence of such early external information,

the ASAs may lose flexibility in choosing alternative schedules as well as valuable task-execution time, because they will only know method outcome at the end of execution.

In this section, we describe a Bayesian-decision-theoretic [15, 1, 3] algorithm for estimating the value of information an owner may have. The algorithm calculates possible method outcomes and associated consequences resulting from getting this information at different times; it factors each possibility’s contribution by the probability it will occur. The algorithm queries the scheduler iteratively to determine the value of information an owner may have. Unlike prior approaches to using decision theory to calculate value of information [6, 9, 7, 23], the algorithm must compute the value of information without knowing the actual piece of information (i.e., without knowing a-priori what value the owner will provide), because the CA does not know what information the owner has, only that the owner has (may have) relevant information. Furthermore, the algorithm handles the overall changes in schedule that may result when new information is obtained (essentially emulating the re-scheduling that would result), and it considers all possible outcomes of a method, not simply success or failure. This functionality is not included in the scheduler and is essential for updating the value-of-information appropriately.

### 3.1 Querying Principles

Throughout this section we will consider a scenario in which the ASA has scheduled method  $M$  to be executed next, and the potential outcomes of this  $M$  are given by the duration vector  $D^M = \{d_1^M, \dots, d_k^M\}$  and the quality vector  $Q^M = \{q_1^M, \dots, q_l^M\}$ . The owner may be able to reveal (with some probability) that the actual outcome of  $M$  is  $(d_i^M, q_j^M)$ , information not directly available to the CA. That is, prior to interacting with the owner, the CA has only the a-priori probability  $P(d_i^M, q_j^M)$  for each potential outcome as given in the method’s definition (as illustrated in Figure 2). The value of the information  $(d_i^M, q_j^M)$  is the marginal improvement in the quality from the schedule after it is changed as a result of obtaining this information (at query time). That is, the value of information  $(d_i^M, q_j^M)$  is the performance of the ASA with this information “a-priori” (i.e., when it queries the owner) minus the performance of the system when it obtains this information only at the end of execution of a method. The value is positive only if the system changes its action depending on  $(d_i^M, q_j^M)$ , but is always non-negative.

In the ASA setting, to estimate the value of an interaction with the owner, the CA needs to weigh the value of each possible specific outcome according to the probability of this outcome, because (as noted above) even when the CA knows that the owner has certain information about a method’s outcome, it does not know the outcome-value the owner knows. The algorithm presented below uses hypothetical (“what-if”) queries, which include the original cTAEMS problem settings and the potential new information from the owner, to the scheduler to determine the value of a specific piece of information.

We will use the notation  $S_t(T, I)$  to denote the best schedule the scheduler can produce at time  $t$ , given a cTAEMS problem  $T$  and the additional information  $I$ . The structure  $I$  may be empty, indicate a specific outcome  $(d_i^M, q_j^M)$  of the method  $M$ , or eliminate a specific outcome from the distribution of possible outcomes. We denote the value (quality) of a schedule  $S_t(T, I)$  by  $S_t(T, I).quality$ . The expected quality of a schedule given updated-method-status information from an interaction with the owner is given by

$$E[Q/access\_owner] = \sum_{i=1}^k \sum_{j=1}^l S_t(T, d_i^M, q_j^M).quality \quad (1)$$

The value of interacting with the owner at a particular time is thus  $E[Q/access\_owner]$  minus the expected quality obtained if the system does not have information about the method’s actual out-

come ahead of time, which we will denote by  $E[Q/\neg access\_owner]$ .

Calculating the value  $E[Q/\neg access\_owner]$  is a non-trivial task. Although the owner is operating according to some schedule, the quality associated with this schedule cannot be used as a baseline, because it does not reflect the re-scheduling that can occur after a method's actual outcomes are known (or other outcomes are eliminated). This rescheduling can affect significantly the actual quality that is achieved. For example, consider a method  $M$  planned to be executed at time  $t$ , that a-priori has four duration outcomes  $\{d_1^M, \dots, d_4^M\}$ , with probability  $P(d_i^M) = 0.25 \forall 1 \leq i \leq 4$ . Now assume that the actual duration outcome of the method is  $d_3^M$ . At time  $t + d_1^M$  the ASA can deduce that outcome  $d_1^M$  is no longer valid for this method, because the ASA did not receive an execution completion (EC) message. As a result, the ASA and all other relevant ASAs can update their expectations for method  $M$  using Bayesian update; the method will now have three possible duration outcomes  $\{d_2^M, \dots, d_4^M\}$  with probability  $P(d_i^M) = 0.33 \forall 2 \leq i \leq 4$ . The ASAs can then re-analyze the problem, based on the elimination of  $d_1^M$ , and re-schedule methods that have not started execution prior to time  $t + d_1^M$ . Similar re-scheduling can be performed at time  $t + d_2^M$ . Finally, at time  $t + d_3^M$  the ASA receives the EC message and all agents can re-schedule any methods that have not yet (i.e., prior to time  $t + d_3^M$ ) started execution. This final rescheduling is based on the actual duration and quality outcome of method  $M$ .

The quality achieved at the end of this process may be very different from the expected quality of the schedule the ASAs had at the beginning of the process. Thus, to calculate the expected quality, it is necessary to capture the effect of the re-scheduling of events that occur whenever a method's execution time exceeds one of the possible durations defined for the method (with no EC message being sent) or when the method completes execution. Generally, if a method is still executing at time  $t + d_j^M$ , then its outcome distribution can be updated to  $\{d_{j+1}^M, \dots, d_{|D^M|}^M\}$  where  $P'(d_i^M) = P(d_i^M)/(1 - \sum_{k=1}^j P(d_k^M)) \forall j < |D^M|, j < i \leq |D^M|$ .

Therefore, any schedule that needs to be produced at a future time  $t'$  ( $t' > t$ ) should be constructed incrementally, based on the last re-scheduling event that would have occurred. For keeping the exposition simple, we assume that method  $M$  is supposed to start execution at time  $t$ . Therefore,

$$S_{t'}(T, (d_i^M, q_j^M)) = S_{t'}(T_c(S_{t+d_i^M}, \neg d_i^M), \neg d_i^M) \quad , t' < t + d_i^M \quad (2)$$

where  $T_c$  is the constrained problem generated by the re-scheduling process that occurs at the time  $d_{i'}$  ( $i' = \text{argmax}_i \{d_i^M < t'\}$ ) when the ASA learns that duration  $d_{i'}^M$  is no longer a valid outcome. Finally, the expected value of the information the owner may have can be calculated as:

$$V = \sum_{i=1}^k \left( \sum_{j=1}^l S_t(T, (d_i^M, q_j^M)).quality - S_{t+d_i^M}(T, (d_i^M, q_j^M)).quality \right) \quad (3)$$

Thus, to estimate the improvement in quality that results from learning each possible outcome two types of queries of the scheduler must be made, (1) a non-constrained query using the new information, and (2) a constrained query that incorporates the results of re-scheduling processes that occur until an EC message is received.

## 3.2 The Querying Algorithm

Algorithm 1, given in pseudo-code below, embodies the principles just described to estimate the value of getting the actual outcome of a method from the owner (before the method completes execution). The core functionality it supports is: (a) the calculation of the differences in the expected quality with and without the owner's information for the outcomes associated with the first (shortest) duration; and, (b) if execution exceeds the first possible duration outcome, recursive execution of the algorithm on the new schedule to obtain the revised quality for the reduced problem.

---

**Algorithm 1** *GetValue*( $T, M$ ) - Evaluating the benefit in interacting with the owner

---

**Input:**  $T$  - a cTAEMS problem;  $M$  - The method we are working on;  $t$  - the time when method  $M$  starts executing;  
**Output:**  $V$  - the expected additional utility if receiving the actual outcome from the owner.

- 1: **if**  $D^M.length = 0$  **then**
- 2:     **return** 0
- 3: **end if**
- 4: Set  $S = \text{GetSchedule}(T)$ ;  $Value = 0$ ;
- 5: Remove schedule  $S$  from cTAEMS  $T$
- 6: **for**  $j = 1$  to  $Q^M.length$  **do**
- 7:     Set  $T_c = T.clone()$ ;  $T_{nc} = T.clone()$ ;  $M' = M.clone()$
- 8:     Set  $D^{M'} = (d_1^M)$ ;  $Q^{M'} = (q_j^M)$ ;  $P(d_1^M) = P(q_j^M) = 1.0$
- 9:     Replace method  $M$  in  $T_c$  and  $T_{nc}$  with method  $M'$
- 10:     **for**  $i = 1$  to  $S.length$  **do**
- 11:         **if**  $S[i].start\_time < t + d_1^M$  **then**
- 12:             Add  $S[i]$  to  $T_c$
- 13:         **end if**
- 14:     **end for**
- 15:     Set  $Value = Value + P(d_1^M)P(q_j^M)(\text{GetSchedule}(T_{nc}).quality - \text{GetSchedule}(T_c).quality)$
- 16: **end for**
- 17: **for**  $j = 2$  to  $D^M.length$  **do**
- 18:     Set  $P(d_j^M) = P(d_j^M)/(1 - P(d_1^M))$
- 19: **end for**
- 20: Set  $P_{remain} = 1 - P(d_1^M)$
- 21: Remove  $d_1^M$  from  $D^M$
- 22: Replace method  $M'$  from  $T_c$  with method  $M$
- 23: **return**  $Value + P_{remain} * \text{GetValue}(T_c, M)$

---

The input to the algorithm is a cTAEMS problem,  $T$ , which includes a partial schedule and a method  $M$  (with an updated outcome distribution).  $M$ 's actual outcome is what the CA is considering asking the user about. The recursive implementation (step 23) allows the algorithm to emulate the re-scheduling process used by the ASA, as described above. The algorithm first queries the scheduler for the best schedule that can be produced given the initial problem settings (represented by  $S$  (in step 4)). The variable *value* accumulates value over the recursion. The main loop (steps 6-16) weighs the differences in expected quality of the schedules for the first possible duration outcome, of the method  $M$ , by creating two new cTAEMS problems (step 7) in which method  $M$  has that duration outcome and one of the quality possible outcomes with probability 1.0 (steps 8-9). In the first problem,  $T_{nc}$ , no scheduling constraints are imposed, whereas in the second problem,  $T_c$  all the ASAs' methods that were scheduled to start executing before the ASA learns the actual outcome of method  $M$  are included in the problem settings as an initial schedule that cannot be changed (steps 10-14). By removing any re-scheduling flexibility up to time  $t + d_1^M$ , the cTAEMS problem  $T_c$  emulates the scenario in which the ASAs adjust to information about method  $M$ 's outcomes only when the actual outcome is determined by the environment (i.e., when the method has completed execution or failed). Similarly, the cTAEMS problem  $T_{nc}$  emulates the scenario in which the information is received a-priori and the ASAs are free to re-plan at this earlier time using this information.

Step 15 uses the function *GetSchedule*( $T$ ) to query the scheduler so that it returns the optimal schedule that can be generated for the problem  $T$  (taking the schedule elements it contains as a constraint for scheduling). A comparison of the quality of the schedule received for the non-constrained problem  $T_{nc}$  and that for the constrained problem  $T_c$  yields the incremental quality improvement that can be obtained by knowing the outcome  $(d_i^M, q_j^M)$  a-priori. The result is multiplied by the probability of getting this outcome from the owner (step 15).

Therefore, upon the completion of steps 6-16 the algorithm has the weighted expected utility of outcomes associated with duration  $d_1^M$ . To complete the process (i.e., for adding the weighted utility for methods with  $d_i^M > d_1^M$ ), the algorithm creates a new problem in which the outcome  $d_1^M$  is removed from  $M$ , and it updates the probabilities of the remaining outcomes (steps 17-22). The new problem ( $T_c$ ) contains the scheduling constraints imposed by the inability to change any of the methods that are scheduled to start their execution prior to time  $t + d_1^M$ . This problem is the input for the recursive execution of the algorithm, which stops when there are no more duration outcomes for  $M$ .

Once the CA module calculates this gain, it can multiply it by the probability representing its estimation that the owner knows the actual outcome (as detailed in [17]) and if the result is greater than its estimated cost associated with the interruption then the interaction is being initiated.

The total number of queries sent to the scheduler throughout the recursive executions of the algorithm is  $(2 * |Q^M| + 1) * |D^M|$ , which is derived as follows: (a) two queries (constrained and non-constrained) for each outcome (making a total of  $|Q^M| * |D^M|$ ), used for calculating the outcome's marginal utility (see step 15); and (b) one query for generating the schedule evolution due to realizing the elimination of a duration outcome, as time goes by (step 4) (i.e., a total of  $|D^M|$ ).

#### 4. IMPROVING EFFICIENCY

The  $(2 * |Q^M| + 1) * |D^M|$  scheduler queries required by Algorithm 1 is a significant overhead, because querying the scheduler is a resource-intensive task. As a result, techniques that enable minimizing the number of queries the CA uses for its benefit estimation process can significantly improve ASA performance. In this section, we show how to reduce significantly the number of queries used for estimating the information value without compromising the accuracy of the obtained result. This decrease is achieved by an intelligent selection of the sequence of outcomes for which queries are generated.

As shown in Equation 3, the expected value of information can be represented as a weighted sum of the differences between the quality associated with the non-constrained problem,  $T_{nc}$ , and the constrained one,  $T_c$ . By identifying outcomes for which the difference is zero, it is possible to eliminate the need for generating the two queries in such cases. Information has value at a specific time only if receipt of that information at that time leads an agent to change its owner's schedule in a way that it could not have if it only received this information in an EC message or induced it from not receiving an EC message at the end of a possible method duration. Thus a zero difference occurs when no agent in the environment changes its owner's schedule in comparison to EC-messages based schedule. Formally, we define "relevant change", which denote the type of changes that indicates the existence of value if the ASA can learn about a method's outcome before it finished executing, as follows:

**DEFINITION 1.** Consider an initial problem  $T$  and a specific outcome  $(d_i^M, q_j^M)$  of method  $M$ . Any change between the schedules  $S_t(T, (d_i^M, q_j^M))$  and  $S_{t+d_i^M}(T)$  over the interval  $(t, t + d_i^M)$  is a "relevant" change.


The basic idea of the first query-reduction method is that the existence of "relevant" changes is the only thing that matters for identifying zero value differences, because these changes determine whether knowing the outcome  $(d_i^M, q_j^M)$  before the EC message occurs contributes value or not. The difference in quality due knowing this information prior to time  $t + d_i^M$  is necessarily positive if "relevant" changes between the non-constrained schedule and the EC-based schedule are found at time prior to  $t + d_i^M$ . Otherwise, the difference is necessarily zero.

	Self changes	Changes in other coordinators
$D \uparrow$	Method $M$ exceeds deadline thus replaced by method $M'$	Method $\overline{M}_1$ that (enabled by $M$ ) cannot start on time thus is replaced with $\overline{M}_2$
$D \downarrow$	The execution of method $M$ can be now delayed and a new method $\overline{M}'$ can be scheduled before it	Method $\overline{M}_1$ enabled by $M$ is added to the schedule, and consequently method $\overline{M}_2$ that facilitates $\overline{M}_1$ is added at time $t' < d^M + t$
$Q \uparrow$	Non-applicable	Method $\overline{M}_1$ facilitated by $M$ is added to the schedule, and consequently method $\overline{M}_2$ that enables $\overline{M}_1$ is added at time $t' < d^M + t$
$Q \downarrow$	Method $M$ is replaced by method $\overline{M}_1$ (that has a better quality)	A method $\overline{M}_1$ under a $q_{min}$ QAF along with $M$ is replaced by method $\overline{M}_2$

**Table 1: "Relevant" schedule changes due to a new information concerning method  $M$ 's quality and duration**

Thus, the CA should first send the scheduler all its non-constrained queries ( $(T_{nc})$ -based), and initiate a constrained query ( $(T_c)$ -based) only for those outcomes  $(q_i^M, d_j^M)$ ,  $q_i^M \in Q^M, d_j^M \in D^M$ , for which a "relevant" change was identified by inspecting the schedule sent back. If the number of outcomes that result with "relevant" changes is  $N < |Q^M| * |D^M|$ , then the number of queries required using this technique for choosing scheduler queries is  $(|Q^M| + 1) * |D^M| + N$ , derived as follows: (a)  $|D^M|$  queries for generating the schedule changes resulting from the elimination of duration outcomes as in the original Algorithm 1; (b)  $|Q^M| * |D^M|$  for obtaining the non-constrained schedules; and (c)  $N$  constrained queries for calculating the quality differences for those scenarios which yield "relevant" changes. Although this approach significantly lessens the number of queries, it is still far from the theoretical minimum number of queries that need to be used (which is  $2 * N + |D^M|$ , where the  $|D^M|$  additional constrained queries are required for following the re-planning procedure and the  $2 * N$  queries ( $N$  constrained and  $N$  non-constrained) are required for calculating the quality differences for those scenarios which yield "relevant" changes).

A further, significant reduction in the number of queries sent to the scheduler may be obtained by analyzing the possible-outcomes space. Figure 3 illustrates this space of outcomes (and is an analogue of the table given in figure 2) for a method  $M$  with 5 possible duration outcomes and 4 quality outcomes. We begin our analysis with an outcome  $o = (d_i, q_j)$  that is not associated with "relevant" changes. Even though this outcome is not associated with "relevant" changes, its neighboring outcomes in the outcomes space (i.e.,  $(d_i, q_{j+k})$  or  $(d_{i+k}, q_j)$ ,  $k \in \{+1, -1\}$ ) may be associated with "relevant" schedule changes. Table 1<sup>5</sup> gives examples of ways an increase or a decrease in the duration or quality results in "relevant" schedule changes (both in a particular ASA's and other ASA's' schedules).



**Figure 3: A method's outcome space**

Nevertheless, if we find that method  $o'$ , which is a neighbor of method  $o$  has "relevant" changes, then "relevant" changes appear in any other outcome in that direction of movement in the outcomes space. For example, if no "relevant" changes are found for outcome  $o = (d_i, q_j)$ , and a "relevant" change is found for outcome  $o' = (d_{i-1}, q_j)$ , then "relevant" changes exist also in any outcome

<sup>5</sup>The table uses  $\overline{M}_j$  to denote a method that is currently non-scheduled.

$o^* = (d_{i-k}, q_j), \forall 0 < k < i - 1$ . This property may be explained by the fact that the same arguments used for changing the schedule when moving from outcome  $o$  to a neighbor outcome  $o'$ , remain valid in the transition to  $o^*$ . Therefore, typically, the area representing the outcomes for which “relevant” schedule changes occur partially wrap around the area representing the outcomes for which no “relevant” schedule changes occur. This characteristic is illustrated in Figure 3 where those outcomes in which “relevant” schedule changes occur are marked with a gray background. An immediate implication of this result is that for any fixed duration  $d_i$ , if there are outcomes  $(d_i, q_j), 1 \leq j \leq |Q^M|$  not associated with “relevant” schedule changes, then these outcomes appear in a sequence along the quality axis. A similar result is obtained when setting a fixed quality.

This regularity suggests an algorithm for solving the problem as a “game” defined by the following rules: (a) given a matrix of size  $|Q^M| * |D^M|$  representing the space of possible outcomes, in which each of the elements can be either associated with “relevant” changes or not; (b) the goal is to identify the matrix elements that are associated with “relevant” changes (marked in gray in Figure 3); (c) at each game stage, determine which element of the matrix to query next; (d) for each element queried, receive an indication of whether it is associated with “relevant” schedule changes; (e) if two matrix elements within the same row (alternatively, column) are found to be not associated with “relevant” changes, then all outcomes between them at that row (alternatively, column) are also not associated with schedule changes. The goal is to query all the elements associated with “relevant” changes while querying as few elements as possible not associated with “relevant” changes.

A simple implementation of an efficient algorithm for playing this game would start with the lowest duration and quality outcome (i.e., the lowest left corner in Figure 3) and will check the matrix elements by moving upward. Upon hitting an outcome that does not involve “relevant” schedule changes (or scanning all elements of a column in this direction), the process moves to the next column (next duration value), moving upward again in the same manner. This process will continue until it either has covered all duration columns or the first outcome of a scanned column reveals no “relevant” schedule changes. In this latter case, the process will continue in the same manner, but starting with the rightmost column and switching to lower duration columns whenever it hits an outcome that does not involve “relevant” schedule changes (or it scans all columns in this direction). When the process reaches a column that has already been scanned, or if the first method in the scanned column reveals no “relevant” schedule changes, then the process is repeated from the top of the columns down the quality scale. (Depending on the dimensions of the problem, it may be better to scan rows instead of columns, but that process is analogous.) This method guarantees a maximum overhead of  $2 * \min(|D|, |Q|)$  unnecessary queries to the scheduler, since at most it will go over all columns and will stop once while moving upward and once while moving downward in each column. Therefore, the number of queries generated using this method is bounded by  $2 * N + 2 * \min(|D^M|, |Q^M|) + |D^M|$ .

## 5. EXPERIMENTS AND RESULTS

Our CA module implementation was successfully integrated in a Coordinators ASA and evaluated in a real-time simulation environment created for the *Coordinators* project.

Table 2 shows sample executions obtained in running the Coordinators’ ASA both with and without the CA module for cTAEMS problems that were typical of the *Coordinators*’ test suite. To focus on the issue of estimating the value of information, the experiments used a fixed interruption-cost value and assumed that any interaction with the owner yielded the actual outcome of the method. The results described in the table are from scenarios that used a fixed cost  $c = 1$  of interacting with the owner. Table 2 lists the 18 (out of

Test	With CA (c=1)				No CA
	quality	# of int.	int. cost	total	
1	12.3	1	1	11.3	12.3
3	19.2	2	2	17.2	8.4
4	36	3	3	33	36
5	31	2	2	29	7.8
6	6	1	1	5	6
7	24	3	3	21	24
9	9	1	1	8	9
10	28	2	2	26	8.8
12	9	1	1	8	2.2
13	55	2	2	53	35
14	53	2	2	51	53
15	19	1	1	18	19
20	6	2	2	4	1.4
21	74	2	2	72	74
24	2.4	1	1	1.4	2
25	0	1	1	-1	0
28	28.8	2	2	26.8	28.8
29	8.2	2	2	6.2	7.9
Totals:	420.9	31	31	389.9	335.6

**Table 2: Results with and without the CA module**

a total of 30 in this particular scenario) test cases in which the CA actually initiated an interaction with the owner. Additional experimental runs produced results similar to the ones in this table. The last two columns specify the quality achieved by the system when operated with and without the CA module. Overall, the quality improvement obtained by using the CA module in those test cases in which it initiated an interaction with the owner was 16.2%. (This percentage grew as we decreased the cost of interruption value and vice versa.) For those test cases resulting in lower overall quality with the CA, the entire reduction in quality value is explained by the interaction cost and is relatively small. In these cases, the information provided by the (simulated) owner did not lead to an inferior plan, but rather left the plan unchanged. This scenario is common, since the decision whether to interrupt the user or not is based on the expected benefit which may differ from the actual benefit. In the scenarios in which the system with CA had higher quality, the quality improved tremendously. Therefore, despite the fact that 12 out of 18 scenarios resulted in a quality decrease (due to the interaction cost), the overall improvement is positive. The same usefulness was reflected in experiments in which we continuously increased the cost of interrupting the user (though here the number of interactions initiated gradually decreased).

We also tested the CA module with a constraint-based scheduler [4] which allowed us to execute extensive tests without the overhead of the full *Coordinators* simulation system. It thus enabled us to test query-reduction performance (as described in Section 4) more thoroughly. These experiments used the test suite that the *Coordinators* project used to evaluate different ASA architectures.<sup>6</sup>

For each test case, the CA picked a random snapshot of the schedule being executed and activated its utility estimation mechanism, for the specific method being executed at that time, with the following different techniques (from Section 4): (a) evaluating all outcomes (denoted “Naive”); (b) sending a non-constrained query for each outcome and then sending constrained queries as necessary (denoted “Improved”); and, (c) using the game-based algorithm described in Section 4 (denoted “Improved+”). In addition, we calculated the theoretical minimum number of queries needed ( $2 * N + |D^M|$ ) for each problem. The graphs below depict the average number of queries used by the CA for estimating the value of information using each of these methods, and the average theoretical minimum (denoted “Theoretical”). The results are based on 1841 problems from the test suite. These problems were divided into types following a classification based on parameters such as the scale of required

<sup>6</sup>The only change we made was to augment the outcome distribution (but without changing the mean and edges values) above the small number (3) typical in the test suite.

scheduling task, the number of agents, and the interdependencies between tasks. Each type represents different problem characteristics and complexity.

The left graph in Figure 4 depicts performance by problem type (x-axis) in terms of number of queries (y-axis). For this set of experiments, there were 5 quality outcomes and 3 duration outcomes for each evaluated method. As shown in the graph, the Improved+ approach provides the greatest improvement and is very close to the theoretical minimum. For problem types 5 and 6, both the theoretical minimum and the Improved+ approach produce very few queries. This result is correlated with the relatively low level of dependencies between methods in these problem types, leading to a relatively small effect of external information on the optimal schedule.

The right graph in Figure 4 depicts the additional overhead (in percentages) generated by each of the different methods in comparison to the theoretical minimum number of queries. In this graph, each group (represented on the horizontal axis) is distinguished by the number of quality outcomes used for each method (the number of duration outcomes was set to be 3) and includes all the problems, regardless of their type (i.e., each group includes all the problems used in the left graph). The theoretical minimum number of queries is the baseline (100%) against which the overhead of all other methods is calculated. This graph shows that while the overhead of the “Naive” and “Improved” methods remains relatively constant as the number of possible outcomes increases, the Improved+ algorithm exhibits a constant improvement as a function of the number of outcomes; it has only a 20% average overhead in comparison to the theoretical minimum number of queries for 7 quality outcomes.




Figure 4: # of interactions initiated

## 6. RELATED WORK

The research reported in this paper is within the broad category of AI research that uses decision-theoretic techniques to address problems of agent decision making in dynamic environments with incomplete information [8]. It relates most directly to work on interruption management, autonomous agents to support planning, and information gathering.

Most work on interruption management is concerned with the question of determining when to interrupt users to provide help or information relevant to their work [6, 9]. For instance, the Notification Platform system [6, 7] modulates the flow of messages from multiple sources to devices by performing ongoing decision analysis. As in the CA problem, the agent is uncertain about a user’s workload and focus of attention as well as about the value of the alternative actions a user will undertake as a result of receiving new information. The primary goal of such work is to balance the expected value of information with the attention-related costs of disruption. It typically uses probabilistic models (often Bayesian network models) of a user’s future actions to infer the probability of alternative feasible actions as a function of actions the agent has observed, us-

ing that context to determine the cost of interruption. Although this problem domain has much in common with the ASA domain, work in this area differs from our work in three ways. First, it focuses almost entirely on modeling the user’s attentional state and avoiding unnecessary or unhelpful interruptions and the ensuing frustration they cause [9, inter alia]. Second, it commonly makes two assumptions that do not hold in the ASA domain: (1) that users will change their actions in a pre-determined manner when they receive the new information, and (2) that it is straightforward to calculate the benefit from performing the alternative action. Third, in this domain, it is the agent that has information needed by the person rather than the reverse. As a result, the question is how much the user having this information will improve the user’s performance. Unlike the CA’s situation, there is no need to deal with uncertainty about the information (i.e., to consider different types of information the user has). Furthermore, since the information is not principally probabilistic (as in the CA’s case), the systems do not need to reason about how users will change their actions in situations analogous to not receiving an EC message.

Research on autonomous agents providing assistance to human planners [14; 23, inter alia] by providing information “alerts” at appropriate times also focuses mainly on the person side of the human-computer interaction, considering such “human” aspects as the cognitive load created on the person providing information or alerting the person to “important” events. In this domain, the value of alerting the user with new information derives from its usefulness in making informed decisions.

Research on value-driven information gathering system [5] focuses on the development of autonomous agents operating in an information rich domain under time and monetary resource restrictions. It considers decisions about the information to collect based on an explicit representation of the user’s decision model and a database of information sources. Queries are activated based on the marginal query value, i.e., by considering the current information known by the system and the set of queries that have been sent but have yielded answers. Query-values are derived from a utility function based on the user’s preferences in combination with information provided by the expert who constructed the decision model.

Our work differs from all this previous work in two ways. First, it considers the problem of efficiently obtaining the value of information; in prior efforts, as discussed above, this is not an issue because the value of information either is known prior to making the interruption decision or can be calculated simply. Furthermore, even in cases where the information value-estimation process requires more complex calculation, the system’s goal is assisting the user so the systems-resource problems that we faced do not arise. Second, because schedule-expertise is external to the CA, it needs to interact effectively with another ASA module to calculate information value whereas prior work has presumed all systems-available information related to the information-value problem is internal to the component making this calculation.

## 7. DISCUSSION AND CONCLUSIONS

This paper addresses the problem of computing value of information for multi-agent coordinated scheduling in situations in which the (human) users of an autonomous agent system (ASA) have access to information that is not directly available to the ASA. It considers this problem in the context of developing a coordination autonomy module (CA) within a coordination-manager system. Unlike other problem domains in which value-of-information must be determined, this setting presents the challenge of computing the value of information when the expertise needed for reasoning about alternative schedules is not contained within the CA itself but rather in a separate scheduler module. Furthermore, the fact that scheduling problems require the value of information be computed in an environment in which the schedule is dynamically (and often, con-



tinuously) being revised, poses a great challenge with respect to the computation of the effect of not getting this information a-priori. As earlier sections of this paper elaborate, even when the ASAs do not receive method-related information from the user at the current time, they will sometime later be able to improve their schedules based on the elimination of different duration outcomes of the executed method (as time goes by) or the receipt, eventually, of the execution completion message of this method. The CA thus must emulate those dynamics in order to accurately infer the system performance for the case in which it does not receive this information from the user. These demands combined with strong constraints on ASA resources raise significant efficiency issues, issues which are normally not addressed in applications for estimating the value of information in other domains. The need for systems to have adequate planning and scheduling capabilities permeates autonomous agents and multi-agent systems applications, making mechanisms such as those described in this paper important for a wide-range of multi-agent systems operating in uncertain dynamic environments.

The paper presents an algorithm that determines the value of information that is directly available only to the person being supported by the ASA, but is potentially relevant to scheduling revisions. The empirical results described in Section 5 provide strong support for the usefulness of the algorithm. In addition, the paper presents two methods for reducing queries to the scheduler to address the problem of strong resource constraints on the ASA. These methods were extensively tested, and the results of these experiments establish that a significant decrease can be achieved in terms of the number of queries the CA module sends by adopting these methods. The implementation of the algorithm and query-reduction methods within an ASA developed in the context of the *Coordinators* provided an environment for validating assumptions in system design and enabled testing of the mechanisms we developed based on an externally developed test suite of experiments. Furthermore, these mechanisms we developed for the CA are scheduler-independent and thus can be used for any ASA implementation.

Although the algorithm and query reduction methods developed in this paper were motivated in part by the separation of CA from the scheduler, these mechanisms are also applicable in systems in which scheduling and value-of-information expertise are contained within a single module. Although such a unified module has the potential for using heuristics to filter queries based on knowledge of scheduling constraints, the use of our approach guarantees obtaining the actual value of information, while such heuristic approaches can provide only estimations.

Possible extensions of this work include addressing the problem of determining the methods to focus on in a given schedule, assuming the CA is constrained with respect to the total number of scheduler queries it can initiate. While heuristic approaches to this problem are important, their advantage is in specific scenarios and such techniques do not fully eliminate the need for mechanisms, such as those developed in this paper, that provide an accurate measure of the value of information the owner may have.

## 8. REFERENCES

- [1] A. Ang and W. Tang, editors. *Probability Concepts in Engineering Planning and Design, Volume II - Decision, Risk, and Reliability*. John Wiley & Sons, New York.
- [2] M. Boddy, B. Horling, J. Phelps, R. Goldman, and R. Vincent. ctaems language specification. Unpublished; available from this paper's authors, 2005.
- [3] R. Clemen, editor. *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press, Belmont, CA.
- [4] C. Gomes, W. van Hoeve, and B. Selman. Constraint programming for distributed planning and scheduling. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [5] J. Grass and S. Zilberstein. A value-driven system for autonomous information gathering. *J. Intell. Inf. Syst.*, 14(1):5–27, 2000.
- [6] E. Horvitz, C. Kadie, T. Paek, and D. Hovel. Models of attention in computing and communication: from principles to applications. *Commun. ACM*, 46(3):52–59, 2003.
- [7] E. Horvitz, P. Koch, C. Kadie, and A. Jacobs. Coordinate: Probabilistic forecasting of presence and availability. In *UAI-02*, pages 224–233. Morgan Kaufmann Publishers, 2002.
- [8] E. J. Horvitz, J. S. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247–302, 1988.
- [9] B. Hui and C. Boutilier. Who's asking for help?: a bayesian approach to intelligent assistance. In *IUI '06*, 2006.
- [10] W. McClure. Technology and command: Implications for military operations in the twenty-first century. Maxwell Air Force Base, Center for Strategy and Technology, 2000.
- [11] D. Musliner, E. Durfee, J. Wu, D. Dolgov, R. Goldman, and M. Boddy. Coordinated plan management using multiagent mdps. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [12] K. Myers, P. Jarvis, and T. Lee. Active coordination of distributed human planners. In *AIPS'02*, 2002.
- [13] J. Phelps and J. Rye. Gppp a domain-independent implementation. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [14] M. Pollack. Intelligent technology for an aging population: The use of ai to assist elders with cognitive impairment. *AI Magazine*, 26(2):9–24, 2005.
- [15] H. Raiffa and R. Schlaifer, editors. *Applied Statistical Decision Theory*. Harvard University Press, Cambridge, MA.
- [16] A. Raja, G. Alexander, and V. Mappillai. Leveraging problem classification in online meta-cognition. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [17] D. Sarne and B. Grosz. Sharing experiences to learn user characteristics in dynamic environments with sparse data. Technical Report TR-15-06, Harvard University, Division of Engineering and Applied Sciences, 2006.
- [18] D. Sarne and B. J. Grosz. Timing interruptions for better human-computer coordinated planning. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [19] P. Scerri, D. Pynadath, W. Johnson, P. Rosenbloom, M. Si, N. Schurr, and M. Tambe. A prototype infrastructure for distributed robot-agent-person teams. In *AAMAS*, pages 433–440, 2003.
- [20] S. Smith, A. Gallagher, T. Zimmerman, L. Barbuлесcu, and Z. Rubinstein. Multi-agent management of joint schedules. In *AAAI Spring Symp. on Distributed Plan and Schedule Management*, 2006.
- [21] T. Wagner. The darpa/ipto coordinators program, 2005.
- [22] T. Wagner, J. Phelps, V. Guralnik, and R. VanRiper. An application view of coordinators: Coordination managers for first responders. In *AAAI*, pages 908–915, 2004.
- [23] D. Wilkins, S. Smith, L. Kramer, T. Lee, and T. Rauenbusch. Execution monitoring and replanning with incremental and collaborative scheduling. In *Workshop on Multiagent Planning and Scheduling, ICAPS'05*, June 2005.