Reasoning about Uncertainty in Agent Control *

Anita Raja
Computer Science Department
University of Massachusetts
Amherst, MA 01003, USA
araja@cs.umass.edu

Thomas Wagner
Computer Science Department
University of Maine
Orono, ME 04469, USA
wagner@umcs.maine.edu

Victor Lesser
Computer Science Department
University of Massachusetts
Amherst, MA 01003, USA
lesser@cs.umass.edu

Abstract

It is paramount for agent-based systems to adapt to the dynamics of open environments. The agents need to adapt their processing to available resources, deadlines, the goal criteria specified by the clients as well their current problem solving context in order to survive in these environments. Design-to-Criteria scheduling is the soft real-time process of custom building a schedule to meet real-time performance goals specified by dynamic client goal criteria (including real-time deadlines). Problem solving tasks are represented using a task model that lays out alternate ways to achieve tasks and subtasks. Recent extensions to this technology has spawned a post-scheduling contingency analysis step that can be employed in deadline critical situations where the added computational cost is worth the expense. This is based on the evaluation of available schedules that can be used to recover from a situation in which partially executed schedules cannot be completed successfully. We describe how this analysis improves the schedule evaluation from the uncertainty perspective, and provide empirical evidence to support our claims.

Keywords: Selection and Planning, Real-time performance, Designing Agent Systems

1 Introduction

Agents need to adapt their processing to available resources, deadlines, the goal criteria specified by the clients as well their current problem solving context in order to survive in open environments. Design-to-Criteria (DTC) scheduling[8] is the soft real-time process of finding an execution path through a hierarchical task network such that the resultant schedule meets certain design criteria, such as real-time deadlines, cost limits, and quality preferences. It is the heart of agent control in agent-based systems such as the resource-Bounded Information Gathering agent BIG [3] and the multi-agent Intelligent Home [2] agent environment. Casting the language into an action-selecting-sequencing problem, the process is to select a subset of primitive actions from a set of candidate actions, and sequence them, so that the end result is an end-to-end schedule of an agent's activities that meets situation specific design criteria.

The general Design-to-Criteria scheduling process is designed to cope with exponential combinatorics and to produce results in soft real-time. However, its somewhat myopic approximation and localization methodologies do not consider the existence of recovery options or their value to the client. In the general case, explicit contingency analysis is not required. In the event of a failure, the scheduler is reinvoked and it plans a new course of action based on the current context (taking into consideration the successes as well as the failures and the value of results that been produced to the particular point). In hard deadline situations, such as those in mission critical systems [5, 4] however, the scheduler may not be able to recover and employ an alternative solution path because valuable time has been spent traversing a solution path that cannot lead to a final solution. Our uncertainty based contingency analysis tools can help in this situation by pre-evaluating the likelihood of recovery from a particular path and factoring that into the utility associated with a particular schedule. The improved estimates (based on the possibility of recovery options) can result in the selection of a different schedule, possibly one that leads to higher quality results with greater frequency resulting in improved real-time performance.

This paper is structured as follows. Section 2 discusses how uncertainty is integrated and leveraged in the main Design-to-Criteria scheduling process. In Section 3 we step outside of the main scheduling process and discuss secondary contingency analysis methodology that uses Design-to-Criteria to explore uncertainty and the ramifications of schedule failure. Experimental results illustrating the strength of contingency analysis, relative to Design-to-Criteria's myopic view, for certain classes of task structures are provided in Section 4.

2 Integrating Uncertainty Into Design-to-Criteria

The Design-to-Criteria scheduling problem is framed in terms of a TÆMS [1] task network, which imposes structure on the primitive actions and defines how they are related. The most notable features of TÆMS are its domain independence, the explicit modeling of alternative ways to perform tasks, the explicit and quantified modeling of interactions between tasks, and the characterization of primitive actions in terms of qual-

* Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Air Force Materiel Command, USAF, under agreement number F30602-97-1-0249 and number F30602-99-2-0525 and by the National Science Foundation under Grant number IIS-9812755 and number IRI-9634938. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, NSF, USAF, or the U.S. Government.

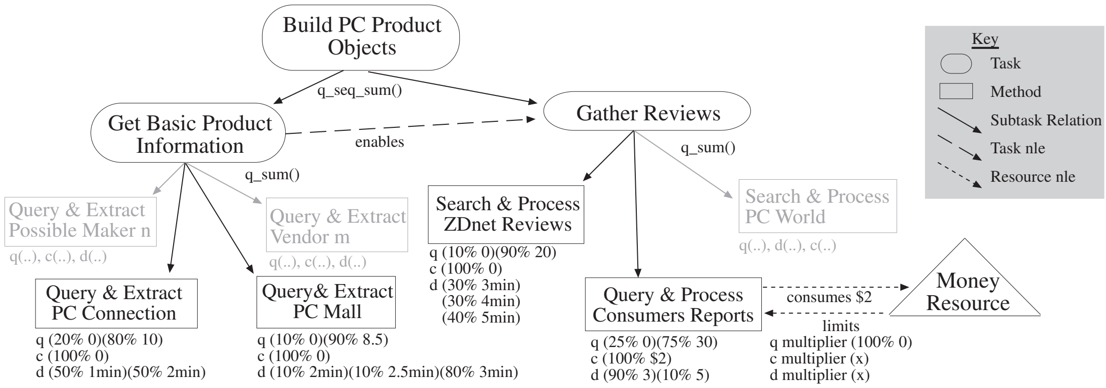

Figure 1: Simplified Subset of an Information Gathering Task Structure

ity, cost, and duration. The issue of task modeling and the associated assumptions are discussed in more detail in [1, 8]. To ground further discussion consider the TÆMS task structure shown in Figure 1. The task structure is a conceptual, simplified sub-graph of a task structure emitted by the BIG information gathering agent; it describes a portion of the information gathering process. The top-level task is to construct product models of retail PC systems. It has two subtasks, *Get-Basic* and *Gather-Reviews*, both of which are decomposed into primitive actions, called *methods*, that are described in terms of their expected quality, cost, and duration. The *enables* arc between *Get-Basic* and *Gather* is a non-local-effect (nle) or task interaction; it models the fact that the review gathering methods need the names of products in order to gather reviews for them. *Get-Basic* has two methods, joined under the *sum()* quality-accumulation-function (*qaf*), which defines how performing the subtasks relate to performing the parent task. In this case, either method or both may be employed to achieve *Get-Basic*. The same is true for *Gather-Reviews*. The qaf for *Build-PC-Product-Objects* is a *seq_sum()* which indicates that the two subtasks must be performed, in order, and that their resultant qualities are summed to determine the quality of the parent task; thus there are nine alternative ways to achieve the top-level goal in this particular sub-structure.

Primitive actions are characterized statistically via discrete probability distributions rather than expected quality values. The quality distributions model the probability of obtaining different quality results and the possibility of failure (indicated by a zero quality result).

The schedules shown in Figure 2 illustrate the value of uncertainty in this model from a scheduling perspective. Schedule A' is constructed for a client who needs a high quality solution, requires the solution in seven minutes or less, and who is willing to pay for it. Note that the quality distribution for Schedule A' includes a 20% chance of failure. Schedule O (Figure 2) is the optimal schedule for the given criteria. Even though the *PC-Connection* method has a higher expected value, the *PC-Mall* method has a lower probability of failure. Since a failure in one of these methods precludes the execution of *Query-Consumers-Reports* (via the task interaction), the issue of failure is not local to the methods but instead impacts the schedule as a whole. Thus, when uncertainty is modeled and propagated during the scheduling process, Schedule O is the optimal schedule as it has the highest net expected quality value and it still meets the client's deadline constraint.

3 Uncertainty-based Contingency Analysis

In the previous sections we explored uncertainty as it is integrated into the standard Design-to-Criteria scheduling methodology. However, in situations where hard deadlines exist, a mid-schedule failure may preclude recovery via rescheduling because sufficient time does not remain to explore a different solution path. In these situations, a stronger analysis that considers the existence of possible recovery options may lead to a better choice of schedules. To address such situations, we have developed a contingency analysis methodology that functions as an optional back-end on the Design-to-Criteria scheduler.

In this section we discuss contingency scheduling issues and formalize four different measures of schedule *robustness*, where robustness describes the quantity of recovery options available for a given schedule. In Section 4 we then present experiments comparing the use of the contingency algorithms to the standard Design-to-Criteria scheduling approach.

This work in contingency analysis of schedules is closely related to recent work in conditional planning. A detailed comparison of our work to other planning-centric research is provided in [7]

Our contingency scheduling research differs from previous work in the following ways:

1. The contingency analysis algorithms use the Design-to-Criteria scheduler to explore mainly the "good" portions of the schedule solution space – that is those schedules that best address the client's design criteria. This serves to constrain the computation and reduces the combinatorics from their general upper bounds. More importantly, the algorithm presented here is amenable to future research in bounding the algorithm directly, which would enable the contingency analysis approach to operate in interactive time, as does the underlying Design-to-Criteria scheduler. Construction of contingent schedules in our analysis is done in interactive time even as the problem is being solved and hence we have real duration and cost constraints in evaluating the entire search space.
2. Contingency analysis takes place in the context of the multi-dimensional goal criteria mechanism used in Design-to-Criteria. Thus the analysis approach is fully targetable to

Schedule A′

PC-Connection	Consumers-Reports

Quality distribution (sum of TGs): (0.20 0.0)(0.20 10.0)(0.60 40.0)
  Expected value: 26.00
  Probability q or greater: 0.60
Cost distribution (sum of methods costs): (1.00 2.0)
  Expected value: 2.00
  Probability c or lower: 1.00
Finish time distribution (finish time of last method): (0.45 4.0)(0.45 5.0)(0.05 6.0)(0.05 7.0)
  Expected value: 4.70
  Probability d or lower: 0.45

Schedule O - Optimal Schedule

PC-Mall	Consumers-Reports

Quality distribution (sum of TGs): (0.10 0.0)(0.22 8.5)(0.67 38.5)
  Expected value: 27.90
  Probability q or greater: 0.67
Cost distribution (sum of methods costs): (1.00 2.0)
  Expected value: 2.00
  Probability c or lower: 1.00
Finish time distribution (finish time of last method): (0.09 5.0)(0.09 5.5)(0.72 6.0)
(0.01 7.0)(0.01 7.5)(0.08 8.0)
  Expected value: 6.05
  Probability d or lower: 0.90

Figure 2: Uncertainty Representation in Schedules

different applications, e.g., situations where quality should be traded-off to obtain lower cost accompanied by a hard deadline, or situations in which quality should be maximized within a hard deadline.

3. Our algorithm takes advantage of the structural properties of the input problem. The TÆMS task structure representation is used to constrain the analysis process and to help limit the exploration of the search used to locate recovery options. This is in contrast to a simple exploration of all primitive actions without regards for interactions or for how the actions relate to achieving the overall goal.

3.1 Performance Measures

In this section we formalize a general theory relating to the contingency planning concepts discussed in the previous section. The question we strive to answer formally here is the following: *What performance measure is the most appropriate estimator of the actual execution behavior of a schedule given the possibility of failure?* Our basic approach is to analyze the uncertainty in the set of candidate schedules to understand whether a better schedule can be selected or an existing schedule can be slightly modified such that its statistical performance profile would be better than that normally chosen by the Design-to-Criteria scheduler. We accomplish this analysis through the use of several performance measures. Prior to presenting the measures, a few basic definitions are needed:

1. A schedule s is defined as a sequence of methods $(m_1, m_{2,}..m_{n-1}, m_n)$.
2. Each method has multiple possible outcomes, denoted m_{ij}, where j denotes the j'th outcome of method m_i. This is part of the TÆMS definition of methods or primitive actions. Though the examples generally present methods as having quality, cost, and duration distributions, methods actually may have sets of these distributions where each set is one possible outcome. For example, if method m may produce two classes of results, one class that is useful by method m_1, and one class that is useful by method m_2, method m will have two different possible outcomes, each of which is modeled via its own quality, cost, and duration distributions. Additionally, these different outcomes will have different nles leading from them to the client methods, m_1 and m_2 respectively.
3. Each outcome is characterized in terms of quality, cost, and duration, via a discrete probability distribution for each of these dimensions and each outcome has some probability of occurrence.
4. m_{ij}^{cr} is a Critical Task Execution Region(*CTER*) when the execution of m_i results in outcome j which has a value or set of values characterized by a high likelihood that the schedule as a whole will not meet its performance objectives. For instance, m_{ij} is a CTER if the probability of the quality of m_{ij} being zero is non-zero.
5. A schedule s could have zero, one or more CTER's in it. A general representation of such schedule with at least one CTER would be $s^{cr} = (m_1, m_{2,}..m_{ij}^{cr}..m_{kl}^{cr}...m_{no}^{cr}..m_{n-1}, m_n)$.
6. f_{ij}^{cr} is the frequency of occurrence of m_i's , j'th outcome where m_{ij} is a CTER.
7. $\overline{m_i^{cr}}$ is m_{ij}^{cr} with its current distribution being redistributed and normalized after the removal of its critical outcome. In other words, the criticality of m_{ij}^{cr} is removed and the new distribution is called $\overline{m_i^{cr}}$.
8. If $s^{cr} = (m_1.., m_{i-1}, m_{ij}^{cr}, m_{i+1}, .m_{kl}^{cr}..m_{no}^{cr}..m_{n-1}, m_n.)$, then
$\overline{s_i^{cr}} = (m_1.., m_{i-1}, \overline{m_i^{cr}}, m_{i+1}, ..m_{kl}^{cr}..m_{no}^{cr}..m_{n-1}, m_n.)$
$\overline{s_k^{cr}} = (m_1.., m_{i-1}, \overline{m_i^{cr}}, m_{i+1}, .\overline{m_k^{cr}}..m_{no}^{cr}..m_{n-1}, m_n.)$ and
$\overline{s^{cr}} = (m_1.., m_{i-1}, \overline{m_i^{cr}}, m_{i+1}, .\overline{m_k^{cr}}..\overline{m_n^{cr}}..m_{n-1}, m_n.)$

The four statistical measures that aide in detailed schedule evaluation are:

Expected Lower Bound (ELB) The expected lower bound rating, of a schedule s, is the performance measure of a schedule execution without taking rescheduling into consideration [8]. It is an expected rating because it is computed on a statistical basis taking quality, cost and duration distributions into account, but ignoring the possibility of rescheduling.

Approximate Expected Upper Bound (AEUB) The AEUB is the statistical schedule rating after eliminating all regions where rescheduling could occur. The assumption is that there are no failure regions and hence the schedule will proceed without any failures and hence no rescheduling will be necessary. The following is a formal definition of AEUB:

Suppose m_{ij}^{cr} is a CTER in the schedule $s = (m_1..m_n)$ and it occurs with frequency f_{ij}^{cr}. Let $\overline{s_i^{cr}}$=$(m_1, m_2..\overline{m_i^{cr}}..m_n)$. If $\frac{ELB(\overline{s_i^{cr}})-ELB(s)}{ELB(s)} \geq \alpha$, then m_{ij} is a *CTER*, where α is a percentage value that determines when a region should be classified a *CTER* and thus a candidate for more detailed analysis. The value of α is contextually dependent and should be specified by a scheduler client. For instance, if saving on computational expense is more important to the client than high certainty, α should be high, and thus the threshold for *CTER* classification is also high. However, if certainty is paramount, then α should be low, indicating that any significant change in the ELB should be explored. When this computation is done on an entire schedule for all of its *CTER's*, we call it the Approximate Expected Upper Bound. Generalizing this formula for k *CTER's* $m_{i_1j_1}...m_{i_kj_k}$, $AEUB(s) = ELB(m_1...m_{i_1-1}, \overline{m_{i_1}^{cr}}..\overline{m_{i_2}^{cr}}.......\overline{m_{i_k}^{cr}}...m_n)$. The AEUB is thus the best rating of a schedule on an expected value basis without any rescheduling.

Optimal Expected Bound (OEB) The OEB is the schedule rating if rescheduling were to take place after each method execution. So the first method is executed, a new scheduling subproblem which includes the effects of the method completion is constructed and the scheduler is re-invoked. The first method in this new schedule is executed and the steps described above are repeated. Hence the optimal schedule is chosen at each rescheduling region. For complex task structures, the calculation would require a tremendous amount of computational power and it is unrealistic to use it for measuring schedule performance in a real system. In most situations, $ELB(s) \leq OEB(s) \leq AEUB(s)$, since the $OEB(s)$ is based on recovery from a failure while $AEUB(s)$ assumes no failure.

Approximate Expected Bound (AEB) It is the schedule rating with rescheduling only at *CTER's* and using expected lower bound of the new stable schedule for methods following the *CTER*. This is limited contingency analysis at *CTER's*. Consider a schedule s of n methods s=$(m_1, m_2..m_i..m_n)$. Now suppose m_{ij} is a *CTER* with a frequency of occurrence of f_{ij}. In order to compute the AEB of the schedule, we replace the portion of the schedule succeeding m_{ij}^{cr}, which is $m_{i+1}, m_{i+2}, ...m_n$ by $l_{i+1}, l_{i+2}......l_k$ if there exists a $l_{i+1}, l_{i+2}......l_k$ such that $ELB(m_1...m_{ij}^{cr}, l_{i+1}...l_k) \geq ELB(m_1...\overline{m_i^{cr}}, m_{i+1}...m_n)$.

The Approximate Expected Bound for this instance is computed as follows:
$AEB_{ij}(m_1, ...m_n)$=$ELB(m_1...\overline{m_i^{cr}}, m_{i+1}..m_n) * (1 - f_{ij}) + ELB(m_1..m_{ij}^{cr}, l_{i+1}..l_k) * f_{ij}$. The new schedule rating thus includes the rating from the original part of the schedule as well the ELB of the new portion of the schedule. This is basically the calculation described when the AEB was introduced in a previous section.

Let $m_1, m_2, m_3, ...m_i...m_n$ be a schedule s of n methods with k *CTER's* named $m_{i_1j_1}^{cr}, m_{i_2j_2}^{cr}...m_{i_kj_k}^{cr}$. Let the recovery path available at each *CTER* m_{ij}^{cr} be s_{ij}^{r} and each m_{ij}^{cr} occurs with frequency f_{ij}^{cr}. The AEB of the entire schedule is described recursively as $AEB = ELB(m_1...m_{ij}^{cr}, l_1, ...l_k) * f_{ij}^{cr} + AEB(m_1...\overline{m_i^{cr}}, m_{i+1}, ...m_n) * (1 - f_{ij}^{cr})$ which can be expanded out as follows:

$$\begin{aligned}
AEB = {} & f_{i_1j_1}^{cr} * ELB(m_1...m_{i_1-1}, m_{i_1j_1}^{cr}, l_{a1}...l_{b1}) \\
& + (1 - f_{i_1j_1}^{cr}) * f_{i_2j_2}^{cr} * ELB(m_1...\overline{m_{i_1}^{cr}}...m_{i_2j_2}^{cr}, l_{a2}...l_{b2}) \\
& + ...(1 - f_{i_1j_1}^{cr}) * ... * (1 - f_{i_{k-1}j_{k-1}}^{cr}) * f_{i_kj_k}^{cr} * \\
& ELB(m_1...\overline{m_{i_1j_1}^{cr}}...\overline{m_{i_2}^{cr}}...\overline{m_{i_3}^{cr}}...m_{i_kj_k}^{cr}, l_{ak}...l_{bk}) + \\
& (1 - f^{cr}{i_1j_1}) * (1 - f_{i_2j_2}^{cr}) * ... * (1 - f_{i_kj_k}^{cr}) * \\
& \underbrace{ELB(m_1...\overline{m_{i_1}^{cr}}...\overline{m_{i_2}^{cr}}...\overline{m_{i_k}^{cr}}...m_n)}_{AEUB}
\end{aligned}$$

The above computation produces an approximate measure since we use the $ELB(m_1..m_{ij}, l_{i+1}..l_k)$. The deeper the recursion in the analysis of *CTER's*, the better the schedule performance measure and the closer it is to the actual performance measure when rescheduling occurs. This describes the potential anytime nature of the AEB computation. Thus, in most situations, $AEB(s) \geq ELB(s)$ by definition.

Here we would like to add that all computations above are based on heuristics and hence are approximations. We could define AEUB',OEB', AEB' and ELB' which would involve complete analysis of all paths by the scheduler. The resulting schedules would display higher performance characteristics and meet goal criteria better but will also be computationally infeasible to generate [6, 8].

4 Experimental Results

Using the measures described above, effective contingency planning is a complex process. It involves taking into account a number of factors, including task relationships, deadlines, the availability of alternatives, and client design criteria (i.e., quality, cost, duration, and certainty trade-offs). In this section, we evaluate the performance of the contingency analysis tools by comparing them to the standard Design-to-Criteria scheduler. Comparison is done by examining the ELB (standard scheduler metric) and the AEB (contingency analysis metric) and comparing schedules selected on the basis of these metrics to the actual results obtained by executing the schedules in a simulation environment.

The experiments in this section were conducted by randomly generating task structures while varying certain characteristics. Intuitions of which characteristics would lead to structures that are amenable to contingency analysis were used to seed the search for interesting test cases. Since method failure is a crucial factor for the contingency analysis argument, the generation of task structures was designed to concentrate on the variance of two factors, namely, the effects of failure location and failure intensity (probability of failure) within a task structure. Ten randomly generated task structure classes were then modified to varying degrees with respect to these two factors. The design criteria in these experiments is to maximize quality given a hard deadline on the overall schedule. This simple design criteria setting is one that lends itself to contingency analysis as the existence of a hard deadline (in contrast to a soft preference, e.g., soft deadline) may preclude recovery via rescheduling in certain circumstances.

The results for the experiments are shown in Figure 3. For each task structure instance, 100 simulated executions were performed using the schedule with the highest ELB and with the schedule having the highest AEB. Each row in the table indicates a different *(failure location, failure probability)* parameter setting for the ten task structures; each row is also an aggregation of results for the ten task structure instances. Of the two factors used to differentiate the task structures in each row, failure location (Lo) (found in the first column of the table) refers to the position of critical method(s) in a task structure and hence in the schedule. Failure intensity (In) (second column) refers to the probability of a method failing. Three different classifications of failure location are used in the experiments: early(E), medium(M), and late(La). Similarly, three different settings for failure intensity are used in the experiments, namely, low(L), medium(M) and high(H) where low is 1%-10% probability of failure, medium is 11%-40%, and high is 41%-90%.

For each problem instance, the execution results produced by the AEB selected schedule were compared to the results for the ELB selected schedule via statistical significance testing. The third column, *N.H. valid count*, identifies the number of problem instances for which the null hypothesis of equivalence could not be rejected at the .05 level via a one-tailed t-test. In other words, *N.H. valid count* identifies the number of experiments for which the results produced via AEB are not statistically significantly different from the results produced by the ELB. These experiments are omitted from subsequent performance measures.

The fourth column indicates the number of task structures of the ten possible whose data is compared. These are task structures that led to schedules for the ELB case and the AEB case that produced execution results that are statistically significantly different, i.e., the null hypothesis of equivalence was rejected at the .05 level. The remaining columns compare the AEB and ELB selected schedules' execution results for the these task structures from an aggregate perspective.

Columns five and eight, titled *Contingency A.Q* and *Normal A.Q.* respectively, show the mean, normalized quality that was produced by the AEB and ELB selected schedules respectively. In other words, the best schedule per the AEB metric was selected and executed in an unbiased simulation environment, when failure occurred the scheduler and contingency-analysis tools were reinvoked and a new schedule generated that attempted to complete the task. The resultant quality was measured and recorded and the experiment repeated 100 times. The same procedure was done for the ELB selected schedule, though when rescheduling occurred, the contingency analysis tools were not invoked (nor were they invoked in the production of the initial schedule). The overall maximum quality produced by either the AEB or the ELB simulation runs was recorded and all resultant quality then normalized over the maximum, resulting a quality value that expresses the percentage of the maximum observed quality that a given trial produced. This procedure was then repeated for the other task structure that produced statistically significantly different results, and the normalized quality values averaged. Thus, the 0.73512 A.Q. from the first row of Table 3, column four, indicates that contingency analysis yielded schedules that produced approximately 74% of the maximum observed quality on average. Column seven indicates that the standard Design-to-Criteria scheduler produced approximately 63% of the maximum observed quality, on average, for the same set of task structures. Thus, contingency analysis yielded a 14.24% percentage increase in resultant quality over the standard Design-to-Criteria scheduler, as shown in column 11.

Columns six and nine show the number of times a given selected schedule failed to produce any result, that is, recovery before the deadline was not possible, for the AEB and ELB cases respectively. It is interesting to note that the contingency selected schedule failed to produce a result with somewhat greater frequency for rows two and five. This is because both the contingency selected schedule and its recovery option had some probability of failure, though, we do not actually consider the failure rate in these cases to be statistically significant. The failure rate in row three illustrates the classic case in which recovery before the deadline is often not possible for the schedules chosen by the standard Design-to-Criteria scheduler, whereas it is more often possible for the schedules selected by contingency analysis.

Columns seven and ten show the number of times rescheduling was necessary during execution. These results are somewhat counter intuitive as the contingency analysis selected schedules generally resulted in more rescheduling during execution due to failure. This is because the contingency analysis tools explore the possibility of recovery and do not seek to avoid the failure in the first place. Relatedly, because the contingency analysis considers the existence of recovery options, it may actually select a schedule more prone to initial failure than the standard Design-to-Criteria scheduler because the schedule has a higher potential quality. For example, say two schedules s_1 and s_2 have the following respective quality distributions: $q_1 = (25\%\ 0)(75\%\ 10)$ and $q_2 = (50\%\ 0)(50\%\ 14)$. The expected value of s_1 is 7.5 whereas the expected value of s_2 is 7. The standard scheduler will prefer s_1 over s_2 because it has a higher expected quality value (assuming that the goal is to maximize quality within a given deadline). However, the contingency analysis tools might actually prefer s_2 over s_1 if there are recovery options, e.g., s_3 for s_2, because s_2 has the potential for a higher quality result than s_1. If s_3 has a quality distribution like $q_3 = (100\%\ 7)$, then the s_2 / s_3 recovery scenario has a higher joint expected quality than does s_1 alone. Associating a cost with rescheduling in the contingency algorithms could modulate this opportunistic risk-taking type of behavior. If a cost were associated with rescheduling, the utility of a recovery option could be weighted to reflect such a cost.

The last column shows the mean normalized OEB of the AEB selected schedule. This is the measure where rescheduling is invoked after every method execution irrespective of the execution outcome. It describes the optimal performance of a schedule since the best possible path is selected every step of the way. The quality value shown is the average of 100 executions of the OEB schedule, normalized by the maximum observed quality over all the AEB selected and ELB selected schedules' executions. The OEB is higher than both *Contingency A.Q.* as well as *Normal A.Q.* for each class of task structures. This is as it should be, as the OEB is a computationally intensive performance measure which strives to obtain the optimal schedule at every point of the plan.

Irrespective of rescheduling, in general, for the task structures that lead to statistically significantly different results, contingency analysis produced schedules that yielded higher average quality than did the standard Design-to-Criteria scheduler. However, as illustrated by the large number of task structures that lead to results that were not statistically significantly different, very few of the candidate task structures were suitable for contingency analysis (about 20%).

Based on the results presented here and other similar results, it is possible to characterize the types of task structures that are amenable to contingency analysis, i.e., those for which analysis of recovery options is beneficial from a cost/benefit perspective. The general characteristics include:

1. Methods in task structures should have a possibility of failure in their distribution.
2. Task structures should contain alternate paths. The absence of of possible recovery paths in the face of failure also makes contingency analysis dispensable.
3. Task structures should contain alternate paths with some overlapping structure(preferably in the initial stage) and with significant performance differences.
4. Dependence of methods with good average performance on critical methods (enables non-local effect from a critical method to a non-critical method).

5 Conclusions and Future Work

Dealing with uncertainty as a first class object both within the scheduling process and via the secondary analysis is beneficial. The addition of uncertainty to the TÆMS modeling framework increases the accuracy of TÆMS models. Including explicit models of uncertainty improves the scheduling process not simply by increasing modeling power, but also by increasing the representational power of all the computations in the scheduling process.

	Fail		N.H valid	T.S.	Contingency			Normal			Perf.	OEB
	Lo	In	count	count	A.Q.	F.R.	R.C.	A.Q.	F.R.	R.C	Impr.	
	E	M	8	2	0.73512	0/200	72	0.63041	0/200	0	14.24%	0.75227
	M	M	8	2	0.70125	2/200	64	0.63883	0/200	0	8.89%	0.71222
	La	M	8	2	0.79936	21/200	100	0.66246	38/200	48	17.12%	0.84531
	M	L	10	0	0	0	0	0	0	0	0%	0
	M	M	8	2	0.70125	3/200	64	0.63883	0/200	0	8.89%	0.71222
	M	H	10	0	0	0	0	0	0	0	0%	0
Col. #	1	2	3	4	5	6	7	8	9	10	11	12

Figure 3: Experimental Results

The secondary contingency analysis procedures presented in Section 3 step outside of this context to perform a more detailed analysis of schedule performance based on the existence of recovery options. Since the algorithms explore the schedule recovery space using the Design-to-Criteria scheduler, they still exhibit a satisficing, approximate, resource conservative nature. They are more appropriate for mission-critical situations which also allow for the post-scheduling off-line contingency analysis. It is interesting to note that even the coarse analysis performed in the AEB and AEUB computations is beneficial in certain circumstances. In [6], we compare the performance of the heuristic-based contingency analysis to that of the policy generated by an optimal controller. Future efforts in contingency analysis will involve explicitly bounding and controlling the complexity of the contingency analysis process. Intertwined with this research objective is the ability to classify particular problem solving instances.

Another area of future exploration in contingency analysis lies in the area of determining critical regions, *CTER*s, within schedules. One aspect of this is determining *CTER* status based on the existence and types of task interactions.

Another area to be explored involves leveraging the uncertainty-enhanced TÆMS models in multi-agent scheduling and coordination. In multi-agent systems the scheduler is typically coupled with a multi-agent coordination module that forms commitments to perform work with other agents; local concerns are thus modulated by non-local problem solving.

Other, more general, future efforts in Design-to-Criteria include using organizational knowledge to guide the scheduler decision process when operating in multi-agent environments and to support negotiation between the scheduler and its clients, which may be other AI problem solvers or humans.

6 Acknowledgments

We would like to thank Professor Shlomo Zilberstein for his contribution in the presentation of the contingency analysis algorithm as well as general comments on improving the quality of the paper.

References

[1] Keith S. Decker. *Environment Centered Analysis and Design of Coordination Mechanisms*. PhD thesis, University of Massachusetts, 1995.

[2] Victor Lesser, Michael Atighetchi, Bryan Horling, Brett Benyo, Anita Raja, Regis Vincent, Thomas Wagner, Ping Xuan, and Shelley XQ. Zhang. A Multi-Agent System for Intelligent Environment Control. In *Proceedings of the Third International Conference on Autonomous Agents (Agents99)*, 1999.

[3] Victor Lesser, Bryan Horling, Frank Klassner, Anita Raja, Thomas Wagner, and Shelley XQ. Zhang. BIG: A resource-bounded information gathering system. In *Artificial Intelligence Journal, Special Issue on Internet Applications*, 1999.

[4] D. J. Musliner, J. A. Hendler, A. K. Agrawala, E. H. Durfee, J. K. Strosnider, and C. J. Paul. The Challenges of Real-Time AI. In *IEEE Computer*, pages 28(1):58–66., 1995.

[5] David J. Musliner. Plan Execution in Mission-Critical Domains. In *Working Notes of the AAAI Fall Symposium on Plan Execution - Problems and Issues*, 1996.

[6] Anita Raja, Victor Lesser, and Thomas Wagner. Toward Robust Agent Control in Open Environments. In *Proceedings of the Fourth International Conference on Autonomous Agents, Barcelona, Spain*, 2000.

[7] Anita Raja, Thomas Wagner, and Victor Lesser. Reasoning about Uncertainty in Design-to-Criteria Scheduling. Technical report, March 2000.

[8] Thomas Wagner, Alan Garvey, and Victor Lesser. Criteria-Directed Heuristic Task Scheduling. *International Journal of Approximate Reasoning, Special Issue on Scheduling*, 19:91–118, 1998. A version also available as UMASS CS TR-97-59.