

Balancing Antagonistic Time and Resource Utilization Constraints in Over-Subscribed Scheduling Problems

Stephen F. Smith and Dhiraj K. Pathak

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

Previous research has demonstrated the utility of constraint-based problem structuring and multi-perspective scheduling as a means for effectively balancing conflicting objectives in "constraint-relaxable" scheduling problems. In this paper, we report initial work aimed at applying these concepts to "over-subscribed" scheduling problems. We focus specifically on the problem of constructing short term schedules for the Hubble Space Telescope (HST), where complex state-dependent constraints on resource utilization interact antagonistically with the ability to satisfy absolute time constraints on execution of goal activities. We first define two time-bounded heuristic scheduling strategies and experimentally analyze the performance of each in balancing this tradeoff. Then, on the basis of results indicating differential superiority of each strategy, a composite multi-perspective strategy based on recognition of characteristics of the current problem solving state is defined. Its experimental performance indicates the benefits of constraint-based problem structuring in over-subscribed domains.

AI TOPIC: Constraint-Based Scheduling

DOMAIN: Space Mission Scheduling

LANGUAGE: Commonlisp, CRL

STATUS: Experimental Prototype

EFFORT: .5 person years

IMPACT: Better exploitation of the unprecedented opportunities provided for mankind by general-purpose space technology.

1 Introduction

Most scheduling problems can be characterized as problems of allocating resources to the activities of multiple processes over time subject to two broad types of constraints: (1) utilization constraints associated with the resources to be allocated, and (2) temporal constraints on the execution of goal activities. Resource utilization constraints dictate the circumstances under which required resources are available to support a given goal activity and are often complex functions of the projected current state (e.g. resource reconfiguration or "setup" activities are often necessary to transition from one application of a resource

to the next). Temporal constraints on activity execution further restrict when an activity can be executed, and typically reflect a confluence of constraints relating to the physics of the processes being coordinated (e.g. activity precedence relations, activity durations) as well as externally imposed requirements (e.g. process release times and deadlines).

Prior research [11, 14] has demonstrated the profitability of viewing scheduling from a constraint satisfaction perspective. At the same time, characterization of scheduling problems as "pure" constraint satisfaction problem is typically misleading, as it implies that the goal is to derive a solution that satisfies all constraints (or determine that one does not exist). The complexity of many important scheduling problems derives in large part from the fact that a solution that satisfies all constraints typically does not exist. The stated problem constraints define a potentially unattainable ideal, and the goal of scheduling is to determine the best (or a satisfactory) overall compromise. Thus, the problem is actually a combinatorial optimization problem within the much larger space of possible compromises.

We can distinguish two broad types of scheduling problems based on the nature of the space of possible compromises (many actual scheduling problems have elements of both). In many cases, specific types of constraints are not truly rigid and are more appropriately considered as choice sets over which optimization criteria can be defined. For example, in factory scheduling, deadlines are typically viewed as relaxable and minimizing order tardiness is a common objective criterion. Similarly, it may be possible to acquire additional resource capacity (e.g. sub-contract manufacturing orders to another facility) but at additional cost. We refer to this class of scheduling problems as *constraint relaxable*.

A second class of scheduling problems is what we refer to as *over-subscribed* problems. Here the time and resource utilization constraints associated with goal activities are not themselves relaxable, but over any finite planning horizon there are always more demands for resources than can be accommodated. The aim of scheduling in over-subscribed problems is to achieve as many goal activities as possible (or alternatively to reject as few potential goal activities as possible).

The problem considered in this paper, the construction of short-term observation schedules for the Hubble Space Telescope (HST), constitutes a representative example of this second class of scheduling prob-

lem. Astronomer demands for viewing time far exceed its capabilities, and maximization of overall telescope utilization is a principal objective. This objective is confounded by complex, state dependent constraints on telescope reconfiguration (which dictate variable delays between the execution of different observations). These constraints interact antagonistically with user-imposed temporal constraints that require specific observations to be executed during specific time periods.

2 Problem solving perspective

Our overall perspective on approaching the HST scheduling problem is based on concepts of constraint-directed scheduling developed in our previous research [11, 16]. This work demonstrated the utility of these concepts as a means for effectively balancing conflicting objectives in constraint-relaxable scheduling problems, and our goal here is to examine their usefulness in the context of over-subscribed problems. Before considering the specific solutions investigated, we first identify the essential characteristics of this problem solving perspective and distinguish our approach from other research in constraint-based scheduling.

The fundamental notion underlying our approach is that characteristics of evolving problem structure (e.g. projected resource contention, tightness of activity execution constraints) can be profitably exploited to dynamically (opportunisticly) focus the scheduler on the subproblems (and tradeoffs) most critical to overall schedule quality. In this regard, our approach has much in common with the perspective of recent work in extending constraint satisfaction problem solving (CSP) frameworks to address scheduling problems [3, 4, 6, 14]. However, where this work emphasizes heuristics for ordering and making individual decisions, our use of dynamic problem decomposition is somewhat different: to provide a basis for mediating the use of a set of heuristic scheduling methods, each differentially attending to various optimization objectives. Problem decomposition heuristics are thus oriented toward identifying (potentially larger granularity) subproblems where specific optimization concerns dominate and should thus drive decision-making. We refer to this approach to subproblem formulation as *constraint-based problem structuring* and the selective use of a set of heuristic scheduling methods as *multi-perspective scheduling* [15].

A second point of departure from CSP frameworks is that, while we advocate the use of consistency labeling and lookahead constraint analysis techniques to uncover evolving problem structure, our approach does not subscribe to the CSP computational paradigm of a complete (and in this case exponential) backtracking search in the worst case. We instead advocate an incomplete search paradigm in which problem constraints and goals are heuristically relaxed as scheduling proceeds, and problems (or opportunities) that are subsequently encountered are dealt with in the context in which they arise. In this regard, our perspective has more in common with the other recent work in the area of mission scheduling [1, 2, 5, 7, 18].

There are several reasons for adopting this perspec-

tive. First, we are interested in guaranteeing solutions to realistically sized problems in a bounded time frame. As suggested in [18], subsequent effort can always be put into improving a generated schedule if time permits. Moreover, since scheduling problems are rarely static in nature (i.e. unanticipated events will repeatedly force changes to the schedule over time), there is always a tradeoff between the expected lifetime of a given solution, its utility and the effort spent obtaining it. A second reason for adopting an incomplete search perspective is that we are interested in addressing problems whose characteristics do not naturally fit into CSP frameworks. Oversubscribed scheduling problems like HST require dynamic determination of the final set of “variables” (i.e. which set of goal activities should be included in the schedule), and cannot be cast strictly as a search for variable assignments. Recent work in the area of design [8] has proposed a framework for integrating variable determination into a CSP framework, but it would appear that the effectiveness of this approach depends heavily on the presence of rich constraints on variable co-dependence. In mission scheduling problems like HST, in contrast, there is much less static structure in the interactions among candidate goal activities (i.e. their principal interactions stem from the fact that they are all competing for the same resources).

In the remainder of this paper, we report work aimed at applying concepts of constraint-based problem structuring and multi-perspective scheduling to over-subscribed scheduling problems. We first provide a brief description the representative HST scheduling problem. Next, we define and experimentally assess the performance of two “basic” heuristic scheduling strategies in balancing the tradeoff between resource setup time minimization and satisfaction of absolute time constraints. The first strategy considered is motivated by “dispatch-based” manufacturing scheduling research (e.g. [13]), and employs a problem decomposition strategy that concentrates local search on minimizing resource idle time due to “setup” activities. The second is motivated by research in micro-opportunistic scheduling (e.g. [14]) and advocates a problem decomposition that focuses attention on the goal activities that have the tightest temporal constraints. Analysis of experimental results gives predictable evidence of the differential superiority of each strategy in different problem solving circumstances. Finally, to examine the potential benefits of constraint-based problem structuring and multi-perspective scheduling, a composite strategy that monitors evolving problem structure and makes selective use of both basic strategies is defined and tested.

3 The HST scheduling problem

The overall objective of HST scheduling is to efficiently allocate viewing time to competing candidate observations (goals) in the presence of complex operational constraints. In brief, a candidate observation represents a user request for an exposure of a certain duration of a particular celestial object (target) using a particular viewing instrument in a par-

ticular operational configuration. Since HST is in low earth orbit, most targets are periodically occulted by the earth, and thus visible only for a portion of each orbit. Over longer periods targets may be similarly occulted by the moon and the sun. Thus, execution possibilities are limited by target “visibility windows” (which are known with certainty over short term horizons). Each of the remaining requirements variably affects when the observation can be executed depending on the prior state of the telescope. It takes time to repoint the telescope toward a different target (an activity referred to as slewing). Similarly, it takes time to reconfigure viewing instruments. There are 6 viewing instruments onboard, and each is capable of being used in a variety of different configurations. Spacecraft power constraints limit the number of instruments that can be operational at any point, which requires execution of complex power-up/power-down sequences as changeovers are made from one instrument to another. Further constraints may also be placed on the execution of an observation for scientific reasons by the user. Candidate observations are typically components of larger observing programs, which may designate partial orderings among observations, separation constraints, absolute execution bounds, and priorities.

For purposes of this paper, we adopt a simplified but characteristic model of the HST operating environment. Specifically, we have chosen not to explicitly model the detailed dynamics of telescope reconfiguration (which would be required to determine and schedule the reconfiguration activities required for schedule executability), but instead account for their existence in terms of temporal delays. In fact, these modeling assumptions correspond precisely to those employed in the abstract layer of the model currently defined within the HSTS observation scheduler [10], a system which integrates scheduling and planning processes to produce executable observation schedules. The work reported below constitutes one component of the larger HSTS project.

In more detail, We assume that the duration of telescope reconfiguration activities (slewing and instrument/configuration changeovers) can be known from the prior observing state of the telescope, and that telescope reconfiguration activities can proceed in parallel. Further we assume that only one instrument can be operational at a time, and model the reconfiguration time from one instrument to another as the maximum of the power down and power up sequence durations. Figure 1 graphically illustrates how these operating constraints might affect the scheduling of observation ob_b after ob_a .

The scheduling problem addressed below consists of constructing a feasible schedule over a finite horizon from a given pool of candidate observing goals. Some portion of the candidate observations are absolutely constrained to occur within a specific sub-interval of the horizon (or become lost opportunities) while the remaining candidates are only absolutely constrained by target visibility. The objective is to maximize telescope utilization within the scheduling horizon while minimizing the number of lost opportunities.




Figure 1: Idle time incurred in scheduling ob_b after ob_a

4 Sequencing heuristics

In this section, we define variants of two types of time-bounded sequencing strategies, differing principally in the relative emphasis placed on the dual (and antagonistic) objectives of maximizing telescope utilization and rejecting as few absolutely constrained observations as possible. In particular, we consider strategies that adopt one or the other of the above two objectives as a driver for structuring the search for a solution. We first consider maximizing telescope utilization as the basis for ordering candidate goals for placement into the schedule, which naturally suggests a dispatch-based (or forward simulation) scheduling strategy. We then consider the relative tightness of the temporal constraints on candidate goals as the basis of goal ordering, which leads to a more opportunistic framework for placing observations on the timeline. In section 5, we report experiments which characterize the differential benefits of each approach.

4.1 Dispatch-based scheduling

If we view the problem strictly as one of maximizing resource utilization (which it is not), then the HST scheduling problem can be seen as a variant of the traveling salesman problem. A simple and reasonably effective heuristic relative to this objective is **nearest neighbor(NN)**. The NN strategy proceeds as follows:

1. Problem Initialization: For each observation ob_i in the pool of unscheduled candidates, the set of possible start time intervals STI_i is computed by
 - (a) first intersecting the visibility windows of the target required by ob_i with the absolute time bounds on ob_i derived from user-imposed timing and ordering constraints (if any), and then
 - (b) removing from this set of start time intervals any start times that are disallowed by the time required to reconfigure the telescope to ob_i 's specification from the initial telescope state.
2. Goal Selection: The candidate $ob_{min-idle}$ having the earliest start time is selected (ties are broken randomly).
3. Reservation Selection: $ob_{min-idle}$ is scheduled at its earliest start time and removed from the pool of unscheduled candidates.

4. Propagation of Resource Unavailability: For each ob_i remaining in the unscheduled pool, STI_i is trimmed to reflect the newly established reservation for $ob_{min-idle}$. Specifically, all start times in STI_i now known to be infeasible due to the confluence of the telescope's new earliest available time, and ob_i 's ordering and telescope reconfiguration time constraints are removed from STI_i .
5. Goal Rejection: Any unscheduled ob_i for which STI_i now contains no start times is marked rejected and removed from the unscheduled candidates pool. If unscheduled candidates remain go to step 2; else stop.

As is evident from the above description, NN emphasizes minimization of telescope dead time (and hence maximization of telescope utilization) to the exclusion of satisfying any absolute temporal constraints on observation execution. Absolutely constrained observations will be summarily rejected unless they serendipitously represent the minimum dead time choice at some point during the schedule development process.

To provide some basis for balancing the need to satisfy absolute constraints with minimization of telescope dead time within a dispatch-based control regime, a second heuristic strategy, **nearest neighbor with lookahead** (NNLA) is defined. NNLA differs from NN as defined above only in the second "Goal Selection" step. Here instead of unilaterally selecting $ob_{min-idle}$ as the next goal to be scheduled, look-ahead is performed to determine whether this choice will cause rejection of other remaining unscheduled candidates. This look-ahead procedure is precisely defined as follows:

1. steps 3, and 4 of the above NN algorithm are simulated, and the set $rej_{min-idle}$ of observations ob_i whose STI_i now contains no start times is identified.
2. If $rej_{min-idle}$ is non-empty, then the set $alt_{min-idle} = rej_{min-idle} \cup ends - before_{min-idle}$ is determined, where $ends - before_{min-idle}$ is the set of all unscheduled ob_j with an earliest end time $<$ the earliest end time of $ob_{min-idle}$. $alt_{min-idle}$ is the set of potentially less disruptive alternatives.
3. For each $ob_a \in alt_{min-idle}$, steps 3 and 4 of the above NN algorithm are simulated, and rej_a computed.
4. The $ob_j \in alt_{min-idle} \cup \{ob_{min-idle}\}$ for which rej_j is smallest is chosen to be scheduled. In case of ties, the observation with the earliest start time is selected.

In simulating step 3 of the NN algorithm, both in lookahead steps 1 and 3, the telescope's new earliest available time is set to the earliest end time of the scheduled observation + *lookahead-period*, where *lookahead-period* is a parameter used to vary the extent to which absolutely constrained observations are

considered in advance of their deadlines. A *lookahead-period* setting of 0 implies they will only be brought into consideration at the last possible moment.

4.2 Most Constrained First Scheduling

The NNLA heuristic places primary emphasis on minimizing resource setup time, giving preference to temporal execution constraints only as necessary to avoid goal rejection. We expect this heuristic to be strong from the standpoint of maximizing the rate of utilization of the resource but be less effective in minimizing the number of rejected goals. In this section, we consider a heuristic strategy designed from a complementary perspective, attending principally to the satisfaction of temporal execution constraints and secondarily to minimization of resource setup time. This is accomplished by dropping the dispatch-based goal ordering strategy in favor of a goal ordering strategy that is instead focused by the relative "tightness" of the execution constraints of unscheduled candidates. Resource setup time constraints are factored into the computation of goal tightness and taken into account during reservation selection. The intuition behind this approach is to provide stronger emphasis on minimizing goal rejection, at the possible expense of lowering the rate of utilization of the resource over the scheduling horizon.

More precisely, this **most-constrained first** (MCF) procedure proceeds as follows:

1. Problem Initialization: For each observation ob_i in the pool of unscheduled candidates, the set of possible start time intervals STI_i is computed by
 - (a) first intersecting the visibility windows of the target required by ob_i with the absolute time bounds on ob_i derived from user-imposed timing and ordering constraints (if any), and then
 - (b) removing from this set of start time intervals any start times that are disallowed by the time required to reconfigure the telescope to ob_i 's specification from the initial telescope state.
2. Goal Selection: The candidate ob_{mc} having the fewest allowable start times is selected for scheduling (ties are broken randomly).
3. Reservation Selection: On any given iteration, the telescope's timeline can be seen as having some number of "holes" into which ob_{mc} can be placed (see Figure 2). A hole is defined to be available for ob_{mc} if there is at least one start time in STI_{mc} that is contained in the hole. Each hole (except for the "rightmost" hole which extends to the end of the horizon) is seen to provide two alternatives:
 - a "leftmost" placement, defined as the earliest start time in STI_{mc} that is contained in the hole, and
 - a "rightmost" placement, defined as the latest start time in STI_{mc} that is contained in the hole.

Only a leftmost placement is possible in the hole that extends to the end of the horizon. For each leftmost (resp. rightmost) placement p_j available to ob_{mc} , the distance from the hole boundary to the start (resp. end) time of p_j , $dist(p_j)$, is computed. The placement p_j for which $dist(p_j)$ is smallest is chosen for ob_{mc} . On the first iteration, when there is only one hole on the timeline, ob_{mc} is placed at its earliest start time.

4. Propagation of Resource Unavailability: For each ob_i still in the unscheduled candidates pool, STI_i is updated to reflect the newly established reservation for ob_{mc} . Specifically, all start times in STI_i now known to be infeasible due to the confluence of telescope's newly established period of unavailability, and ob_i 's duration, ordering and instrument reconfiguration time constraints are removed from STI_i .
5. Goal Rejection: Any unscheduled ob_i for which STI_i now contains no start times is marked rejected and removed from the unscheduled candidates pool. If unscheduled candidates remain go to step 2; else stop.



Figure 2: Telescope availability over time

It is interesting to note that MCF bears considerable similarity to the scheduling strategy currently in place at the Space Telescope Science Institute (STScI) for short term scheduling of HST [17]. There appear to be two principal differences.¹ First, goal selection appears to be based on static “scores” that reflect both pre-assigned observation priorities and a priori tightness of observation time constraints. MCF does not consider priorities, given that our interest in this paper is in the more basic tradeoff presented by antagonistic time and resource constraints under equal goal priorities, but we recognize the future need to incorporate such influences. With respect to tightness of time constraints, MCF operates with dynamic (as opposed to fixed) “criticality” metrics (the current size of each STI_i). Our belief, supported by recent research in micro-opportunistic scheduling [14], is that a dynamic goal ordering scheme will provide greater leverage in minimizing the number of goals rejected.

The second principal difference in comparison to MCF is that a subset of goals are selected on each cycle (as opposed to a single goal). A variant of Step 3 of the MCF procedure is applied to each candidate ob_i in this selected subset, which computes a score for each possible placement p_j of ob_i by reducing ob_i 's fixed

¹ Available documentation of this algorithm is quite sparse and this characterization reflects our current understanding. We have not seen the code.

Heuristic Strategy	Goals Rej.	% Abs. Sched.	Major Setups	% Util.
NNLA	4.5	72	4.33	21.59
MCF	1.5	93	8.17	17.20

Table 1: Performance of NNLA and MCF

score in proportion to $dist(p_j)$ (as defined above). The highest scoring candidate/placement pair is the reservation that is selected. This evaluation of alternatives for multiple candidates during each cycle raises the overall complexity of the procedure, but may provide some compensation for the absence of dynamic goal ordering.

5 Experimental analysis

In this section we present the results of a set of experiments designed to assess the comparative performance of the NNLA and MCF scheduling strategies. The experiments were performed relative to a partial model of the HST. In particular two instruments were modeled: the Wide Field/Planetary Camera (WF/PC) and the Faint Object Spectrograph (FOS), the former having 4 operational configurations and the latter having 2. Actual reconfiguration times were used, with changeover times ranging from 0 (no configuration change) to 1680 seconds (shortest configuration change on same instrument) to 11740 seconds (longest instrument changeover time). Slew time from one target to another ranged from 0 (same target) to 1500 (approximately 180 degree slew) depending on the angular distance between targets. A data base of 76 target locations provided by the Space Telescope Science Institute (STScI) was used as a basis for problem generation. An orbital event generator also provided by STScI was used to derive visibility windows for each target. For each experiment performed, an initial (7th) configuration state where both instruments are turned off was assumed.

A series of 6 problems were generated, each consisting of 30 randomly generated observations (estimated to be schedulable over a 2 day horizon). For each candidate observation, the required target was selected randomly from the target data base and the required instrument configuration was selected randomly from the 6 possibilities according to a probability distribution assuming equal demand for each instrument. Observation durations were randomly drawn from the interval [200 seconds, 2000 seconds]. In each problem, a subset of the candidate observations (between 10 and 20) where absolutely constrained to occur at some sub-interval within a 1.5 day horizon. Execution intervals were randomly chosen from the interval [30000 seconds, 50000 seconds] and uniformly distributed over the 1.5 day horizon.

Performance results obtained for both NNLA (with *lookahead-period* = 0) and MCF averaged over all 6 runs are given in Table 1. Both strategies exhibited the strengths and weaknesses that would have been expected, given the differential emphasis placed on

the two problem objectives in their operation. NNLA achieved a 4.39% better utilization percentage than did MCF, where percentage utilization is measured as amount of actual time spent observing over the 1.5 day interval. This advantage is directly attributable to NNLA's ability to minimize the number of major setups required (major setups are considered to be telescope reconfigurations lasting longer than 1 hour). With respect to goal rejection, alternatively, MCF produced clearly superior performance, on average rejecting only 1.5 observations. With respect to the average percentage of absolutely constrained observations scheduled, MCF achieved 93% to the 72% achieved by NNLA. MCF acts to maximize the number of absolutely constrained goals satisfied and therefore pays a penalty in its ability to cluster observations with similar configurations. MCF schedules contained, on average, roughly twice as many major setups.

The results described above support the hypothesis that different heuristics must be devised to exploit different characteristics of the sequencing problem. NNLA attempts to minimize setup cost by clustering observations with similar configurations. It delays absolutely constrained observations until the very last moment. The result is that NNLA achieves good utilization at the expense of the ability to satisfy all goals. MCF acts to maximize the number of absolutely constrained goals satisfied by focusing first on the most constrained ones. As a result, MCF is able to satisfy a greater proportion of goals at the expense of lower utilization. In general, NNLA is best suited for a situation in which all goals have similar criticality while MCF is best suited for a situation in which a subset of goals are more severely constrained.

The differential performance of NNLA and MCF in this experiment suggests the potential advantage of a composite, multi-perspective scheduling strategy. Such a strategy is defined and evaluated in the following section.

6 A composite sequencing strategy

One property of the MCF strategy is that once the tightness of unscheduled candidates becomes relatively equal, there is no selective pressure (other than the opportunities provided by current hole boundaries on the timeline) toward minimizing idle times. This lack of focus was highlighted in the experiments reported above, due to MCF's failure to establish multiple holes in the timeline during the placement of absolutely constrained goals. The desired pressure at this point in the schedule development process is exactly that which is provided by NNLA.

To test this hypothesis, a composite strategy (MCF/NNLA) was defined. The MCF procedure was modified to include a monitoring step at the beginning of each cycle. During this monitoring step, the number of start time intervals of each unscheduled candidate was determined, and both the minimum and maximum were computed. If, on any given cycle, the difference between these computed maximum and minimum numbers of start time intervals fell below a specified threshold, the [goal selection] and [reservation selec-

Heuristic Strategy	Goals Rej.	% Abs. Sched.	Major Setups	% Util.
MCF/NNLA	1.5	93	7.17	20.54

Table 2: Performance results for MCF/NNLA

tion] steps of the NNLA procedure were substituted for those of MCF.

Performance results obtained with the composite MCF/NNLA strategy on the problem set defined above are given in Table 2. As can be seen, the differential capabilities of both individual strategies are productively combined. The average percentage of telescope utilization over the scheduling horizon achieved comes close to that achieved by NNLA alone. At the same time, the percentage of absolutely constrained goals remains the same as that achieved by MCF alone.

7 Discussion

Our objective in conducting the above study was to explore the utility in over-subscribed scheduling domains like HST of constraint-based techniques that selectively exploit heuristics with differential optimization capabilities according to the structure of the problem's constraints. We believe the results obtained demonstrate the viability of this approach and provide initial evidence of the leverage to be gained in attending to conflicting goal satisfaction and resource utilization objectives.

At the same time, it is not our intent to argue for or conclude about the general merits of the MCF/NNLA strategy as a basis for balancing antagonistic time and resource constraints in over-subscribed scheduling problems. In this regard, the experimentation performed to date must be considered preliminary. The problem set used to produce the reported results exhibits a particular set of characteristics, whose influences can be seen in the behavior of each base strategy. For example, the average tightness of associated absolute constraints is fairly large (approximately 10 hours) relative to the overall 1.5 day scheduling horizon. This problem structure would appear to work against the early establishment of multiple "holes" on the timeline by MCF (which results in less opportunities to minimize idle time in subsequent placements). This phenomenon was actually observed in tracing the development of the schedule in the MCF runs. Under different problem circumstances, where some observations are very tightly constrained relative to the length of scheduling horizon and there is higher variance in tightness of absolute constraints, we expect there would be a larger probability of establishing more useful islands from which to expand. Similarly, The goal rejection rate observed in NNLA's solutions can also be related to characteristics of the test problems. Obviously, as the number of absolutely constrained goals is reduced, the likelihood of simultaneously encountering the end of execution possibilities for multiple candidates is reduced. In fact, other ex-

periments performed but not reported here, indicate the dominance of NNLA in problems where no absolutely constrained goals are present.

In sum, we believe that NNLA and MCF constitute useful initial building blocks for multi-perspective scheduling in over-subscribed domains and that MCF/NNLA provides a useful starting point in exploiting problem constraints to arbitrate their use, but we do not yet have a full understanding of the behavior of these base strategies. Further experimentation under a range of different problem characteristics is underway, and we expect that this work will produce additional insights with respect to both base strategy refinement and heuristics for effectively exploiting the structure of problem constraints to direct the overall scheduling effort.

We are also interested in exploring the comparative behavior of other base strategies. First, we believe variants of MCF that exploit focus of attention heuristics based on resource contention (e.g. [9, 14]) could prove to provide a more effective basis for establishing initial islands on the timeline under many problem circumstances. Second, we believe that repair-based methods (e.g. [12]) can be profitably exploited to improve the generated schedule when time permits. The investigation and integration of such additional scheduling methods is a second focus of our current work.

Finally, the strategies described in this paper have ignored certain types of constraints that are present both in the HST problem and generally in other over-subscribed scheduling problems (e.g. priorities, restrictions/ preferences on goal completion levels). A third focus of our current work concerns extension of the approach to cover these additional types of constraints.

Acknowledgements

This research has been sponsored in part by the National Aeronautics and Space Administration under contract NCC 2-531 and the CMU Robotics Institute.

References

- [1] E. Beifeld and L. Cooper. Operations mission planner: Final report. Publication 90-16, Jet Propulsion Laboratory, March 1990.
- [2] C.E. Bell. Scheduling experiments in an environment with complex temporal and resource constraints. Department of management sciences working paper, University of Iowa, 1990.
- [3] P.M. Berry. Satisfying conflicting objectives in factory scheduling. In *Proceedings 6th IEEE Conference on AI Applications*, pages 101-106, Santa Barbara, CA, March 1990.
- [4] M.S. Fox, N. Sadeh, and C. Baycan. Constrained heuristic search. In *Proceedings IJCAI-89*, pages 309-315, Detroit, MI, August 1989.
- [5] M. Johnston. Spike: Ai scheduling for nasa's hubble space telescope. In *Proceedings 6th IEEE Conference on AI Applications*, pages 184-190, Santa Barbara, CA, March 1990.
- [6] N.P. Keng and D.Y.Y. Yun. A planning/scheduling methodology for the constrained resource problem. In *Proceedings IJCAI-89*, pages 998-1003, Detroit, Michigan, August 1989.
- [7] S. Minton, M.D. Johnston, A.B. Philips, and P. Laird. Solving large-scale constraint satisfaction and scheduling problems using a heuristic repair method. In *Proceedings AAAI-90*, pages 17-24, Boston, MA, July 1990.
- [8] S. Mittal and B. Falkenhainer. Dynamic constraint satisfaction problems. In *Proceedings AAAI-90*, pages 25-32, Boston, MA, July 1990.
- [9] N. Muscettola and S.F. Smith. A probabilistic framework for resource-constrained multi-agent planning. In *Proceedings IJCAI-87*, Milano, Italy, August 1987.
- [10] N. Muscettola, S.F. Smith, A. Cesta, and D. D'Aloisi. Coordinating space telescope operations in an integrated planning and scheduling framework. In *Proceedings IEEE Conference on Robotics and Automation*, April 1991.
- [11] P.S. Ow and S.F. Smith. Viewing scheduling as an opportunistic problem solving process. *Annals of Operations Research*, 12:85-108, 1988.
- [12] P.S. Ow, S.F. Smith, and A. Thiriez. Reactive plan revision. In *Proceedings AAAI-88*, St. Paul, Minn., August 1988.
- [13] S.S. Panwalker and W. Iskander. A survey of scheduling rules. *Operations Research*, 25:45-61, 1977.
- [14] N. Sadeh. *Look-Ahead Techniques for Micro-Opportunistic Job Shop Scheduling*. PhD thesis, School of Computer Science, Carnegie Mellon University, May 1991.
- [15] S.F. Smith and P.S. Ow. The use of multiple problem decompositions in time-constrained planning tasks. In *Proceedings IJCAI-85*, Los Angeles, CA, August 1985.
- [16] S.F. Smith, P.S. Ow, J.Y. Potvin, N. Muscettola, and D. Matthys. An integrated framework for generating and revising factory schedules. *Journal of the Operational Research Society*, 41(6):539-552, June 1990.
- [17] STScI. Space telescope planning and scheduling software review meeting. November 1987.
- [18] M. Zweben, M. Deale, and R. Gargan. Anytime rescheduling. In *Proc. 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA, November 1990.