Automatic Generation of Heuristics for Scheduling

Robert A. Morris
Florida Institute of Technology
150 W. University Blvd.
Melbourne, FL 32901, U.S.A.
morris@cs.fit.edu

John L. Bresina
Recom Technologies
NASA Ames Research Center
Mail Stop 269-2
MofFett Field, CA 94035, U.S.A
bresina@ptolemy.arc.nasa.gov

Stuart M. Rodgers
Computer Science Department
University of West Florida
Pensacola, FL 32514, U.S.A
sturodgers@sprintmail.com

Abstract

This paper presents a technique, called GENH, that automatically generates search heuristics for scheduling problems. The impetus for developing this technique is the growing consensus that heuristics encode advice that is, at best, useful in solving most, or typical, problem instances, and, at worst, useful in solving only a narrowly defined set of instances. In either case, heuristic problem solvers, to be broadly applicable, should have a means of automatically adjusting to the idiosyncrasies of each problem instance. GENH generates a search heuristic for a given problem instance by hill-climbing in the space of possible multi-attribute heuristics, where the evaluation of a candidate heuristic is based on the quality of the solution found under its guidance. We present empirical results obtained by applying GENH to the real world problem of telescope observation scheduling. These results demonstrate that GENH is a simple and effective way of improving the performance of an heuristic scheduler.

1 Introduction

Employing heuristic methods to solve intractable constrained optimization problems like scheduling often suffers from narrowness in the range of problems to which they can be effectively applied. To overcome these limitations, researchers have sought to develop more adaptive approaches to problem solving; e.g., [Gratch and Chien, 1996]. Adaptive heuristic problem solving delays the selection of an heuristic strategy until some information can be obtained with respect to which strategy is expected to perform most effectively in solving a problem instance or class of instances. For problems which require a solution to satisfy a set of constraints, the "best" heuristic is typically one which is expected to allow the problem solver to most efficiently converge on a solution. For example, the approach exemplified by SOAR [Laird, *et al.*, 1986] uses traces of past problem-solving efforts to refine heuristics in order to speed up the search for a solution. However, other criteria besides problem solving efficiency for heuristic selection are possible. For scheduling and other constrained optimization problems, quality of solution may be a more crucial metric with which to compare and select heuristics.

This paper proposes a technique, GENH, for automatically generating heuristics for solving scheduling problems. GENH generates a search heuristic for a given problem instance by hill-climbing in the space of possible multi-attribute heuristics, where the evaluation of a candidate heuristic is based on the quality of the solution found under its guidance. GENH has been successfully applied to the problem of scheduling telescope observations using the *Associate Principal Astronomer, or APA* [Drurnmond, *et al.*, 1994], a system developed at NASA Ames Research Center. Given a set of telescope observation requests supplied by the user of the APA, GENH solves the problem instance using several "versions" of the APA scheduler, where the versions differ only in the search heuristic employed. Adding GENH to the scheduling process incurs acceptable computational overhead, is *accurate* (i.e., converges to a better solution more often than previously employed techniques), *robust* (solves a wide range of problem instances), *simple* (is based on a simple algorithm easily integrated into the scheduling process), and potentially *generalizable* to other problem domains. Following a brief summary of the problem domain and the APA scheduler in Section 2, we present a description of the GENH method (Section 3) and a summary of experimental results (Section 4). We then discuss related research (Section 5), future research (Section 5) and conclude (Section 6).

2 Telescope observation scheduling

The input to the APA observation scheduler is a set of requests, expressed using the Automatic Telescope Instruction Set, or ATis, [Boyd, *et al.*, 1993]. Each request is composed of a sequence of telescope movements and instrument commands, as well as scheduling constraints and preferences. An observation request is said to be *enabled* on a given night if all its constraints are met. The *enablement intervalis* the duration of enabled time for a request. The enablement interval is determined by the observation season, as well as factors such as position of the moon on a particular night. For further details on this domain, see [Bresina, *et al.*, 1994; Drummond, *et al.*, 1995].

The ATIS standard also specifies an heuristic dispatch policy which can be used to select the next observation to execute. The policy is expressed as four selection rules: *priority, number-of-observations-remaining, nearest-to-end-window,* and *file-position.* The rules are applied in the sequence given to the currently enabled observations; each rule is used to break ties that remain from the application of those that preceded it. If the result of applying any rule is that there is only one group remaining, that group is selected for execution and no further rules are applied. Since there can be no file-position ties, the dispatch policy is deterministic.

ATIS dispatch is a robust scheduling method that has been used fairly successfully for several years to schedule automatic photoelectric telescopes at Fairborn Observatory before the development of the APA. (See [Henry, 1996] for a performance evaluation of ATIS dispatch.) The dispatch decisions are determined purely locally, without look-ahead; by contrast, the APA uses a search-based scheduler (for APA scheduler details, see [Bresina, *et al,* 1996; Edgington, *et ai,* 1996]).

The scheduler's search space is organized chronologically as a tree, where the root node consists of the world model state at the beginning of the night. Each arc out of a search tree node represents an enabled request. The purpose of the search heuristic in the APA scheduler is to determine which of the enabled requests to select to extend a partial schedule. In the experiments conducted here, selection is "greedy"; i.e., the schedule extension with the lowest (= best) heuristic score is chosen. Because each extension of a partial schedule is feasible, there are no failure nodes. The task of the APA scheduler is to find a sequence of observations that achieves a good score according to the user-defined objective function.

Because the APA scheduler does not employ backtracking in its search for a schedule, it is not relevant in evaluating heuristics to compare the computational cost incurred (i.e., number of states expanded); differences in the cost of applying different heuristics are marginal. Hence, in our experiments the sole performance metric is in terms of schedule quality according to the domain-specific objective function. In collaboration with astronomers, we defined the following three objective function attributes which describe preferred characteristics of telescope observation schedules.

1. *Airmass Quality.* Measured in terms of where in the sky the observations were taken. Closer to the meridian means lower airmass, hence better observation conditions; closer to the horizon means higher airmass.
2. *Conformity to "Season Track".* Roughly, a season track is an ideal "path" of movements a telescope should take through the night from West to East. This is expressed as a mapping from local time to telescope pointing angle.
3. *Priority.* It is usually impossible to observe all the requests on a given night, so prefer observing the more important request (i.e., those with lower priority numbers).

In the experiments reported here, the objective function used to evaluate complete schedules is the weighted summation of these three attributes, with each attribute assigned an equal weight of 1/3. The attribute values were first scaled, so that equal weighting does imply equal importance within the objective score.

As with the objective function, a search heuristic used by the APA scheduler is based on a set of attributes, $A_h = \{a_1, \ldots, a_n\}$. For the experiments conducted here, this set contains the three objective function attributes listed above, as well as the following two attributes which were derived from the second and third rules of the ATIS dispatch policy.

4. *Fewest Runs.* Prefer requests that have been observed the least number of times - a fairness issue.
5. *Least time left in enablement window.* Prefer requests whose enablement interval is smallest, in particular, those requests that are about to become unobservable for the rest of the night.

For each attribute $a_i \in A_h$, there is a domain w_a, of possible weight assignments. In this application, each w_a, is the real unit interval [0,1]. An heuristic is the weighted summation $\sum_{i=1}^{n} w_{a_i} a_i$; as in the objective function, the attribute values are first scaled. The *space* of possible heuristics can be viewed as the (uncountably infinite) set $S_h = w_{a_1} \times \ldots \times w_{a_m}, a_i \in A_h$; hence, exhaustive search of this space is impossible.

3 The GENH technique

GENH adds to the APA the capability of selecting an effective search heuristic based on the specific characteristics of a problem instance. The input to GENH consists of an objective function, the set *Ah* of heuristic attributes, an observing night, and a set of requests that are enabled sometime during that night. The output from GENH is the best heuristic found as the result of a search through the space of heuristics defined by *Ah*- GENH employs a hill-climbing approach to conduct a focussed exploration of S_h the following factors control this search:

1. a *seed* to initialize the search;
2. a *tuning procedure,* based on a weight adjustment-function; and
3. a *termination condition.*

The pseudo-code for a generic GENH heuristic selection algorithm for observation scheduling based on these factors is displayed in Figure 1. First, GENH initializes a seed heuristic as a "head start" for hill-climbing. The algorithm then iteratively tunes the heuristic by conducting a hill-climbing search through S_h, always selecting the best improvement to the currently best heuristic, *bestH.* The procedure *selectAdjustments* generates candidate improvements, and the procedure *bestScore* selects the best improvement from these candidates.

```
input: R, set of observation requests;
       Ah, set of heuristic attributes;
       OB, objective function.
output: bestH, best heuristic found
begin
bestH = seed(OB,A,R)\
bestScore({bestH}, R, OB, bestH, score);
initialize boundary variables;
until some boundary condition is exceeded do
    H = select Adjustments(bestH);
    bestScore(H, R, OB, h, bestscore);
    if bestscore < score then
       score = bestscore,
       bestH = h; endif
    reset boundary variables;
return bestH
end
```

Figure 1: GENH heuristic selection algorithm

The result of a single call to *selectAdjustments* is a set H of candidate heuristics formed by adjusting the weights of different subsets of *bestH's* attributes. The procedure *bestScore* has input variables comprised of H, R, a set of requests (problem instance), and *OB*, a user-defined objective function. For each candidate heuristic, *bestScore* invokes the APA scheduler to find a schedule via greedy heuristic search. Each resulting schedule is scored using the objective function, *OB*, and *bestScore* returns the best scoring heuristic, h, among the set of candidates along with its score, *score*. If the returned h is an improvement over the current best heuristic (i.e., *score* is better than *bestScore*), then h becomes the new *bestH* and the improvement process is repeated.

The termination condition for GENH is determined by boundary conditions, chosen empirically, for a set of global variables. One boundary variable establishes the size of the weight adjustment in the tuning algorithm. This value is decremented until some limit is reached. Another pair of variables establish the upper and lower bounds of possible weight values of attributes. The interval defined by these bounds shrinks during GENH'S search. Eventually, a condition is reached in which tuning produces no changes to the value of *bestH*, at which time boundary variables are reset to exceed their limits, whereupon the algorithm terminates.

In the remainder of this section, we discuss the seed selection and weight adjustment modules of this algorithm in more detail. It has been noted that the ability of hill-climbing algorithms to produce good solutions seems to depend upon the ability to provide a good "head start" on hill climbing. This is achieved using a pre-processing stage to produce an initial assignment which is "close to a solution" [Morris, 1992]. We isolated the seed selection component of GENH and empirically investigated different methods for selecting a seed heuristic. Here we only discuss the one method used in the experiments reported in Section 4.

This seed method roughly "mirrors" the weights assigned in the given objective function. The intuition behind this approach is that applying the objective function to a partial schedule is a reasonable predictor of the score of the complete schedule. One potential drawback is that attributes which are more "global" (e.g., schedule fairness) tend to not perform as well a local heuristic; however, such attributes should get weeded out during the weight adjustment phase.

In this seed method, those candidate heuristic attributes in the objective function are initially assigned the same weights as in the objective and the additional attributes are initially assigned a very low weight. These initial weight assignments are then normalized so that they sum to 1. This requirement that weights sum to 1 is also maintained during the weight adjustment phase. This constraint is not imposed by the APA scheduler; rather it was applied here simply to avoid redundancy in the exploration of S_h.

A weight adjustment method systematically adjusts values to a weight vector $w \in S_h$ in the direction that most improves the performance of the scheduler. The adjustment mechanisms that we have explored are inspired by optimization methods in Operations Research, in particular, by methods for selecting multipliers for solving Lagrangian relaxation problems [Fisher, 1985]. As with seed selection, we empirically investigated a range of adjustment methods which vary in the size of the heuristic search space explored; however, here, we only define the one used in the reported experiments.

This robust tuning method generates a candidate improvement to *bestH* by adjusting the weights on a pair of attributes, increasing one and decreasing the other. Hence, with this method, a single call to *select Adjustment* returns a set of $a \times (a - 1)$ candidates. The size of each weight adjustment is determined by the values of the boundary variables, which diminish monotonically as the method is applied. The result is that as better heuristics are found, smaller regions of S_h around *bestH* are explored. Eventually, no significant changes are made to the currently best heuristic, at which point the search terminates.

There is no guarantee that this weight adjustment method ever eventually converges to an optimal heuristic; however, the next section demonstrates that this tuning method works well in practice over a wide range of problem instances.

4 Experiments with GENH

Experiments were conducted to test the hypothesis that automatically adapting the heuristic to the problem instance with GENH will result in higher quality solutions than scheduling with a fixed heuristic approach. GENH'S heuristics were pitted against the ATIS dispatch policy, and they were pitted against using the objective function as the search heuristic. The reason for the former comparison is that for this application, ATIS dispatch is the gold standard; the reason for the latter comparison is

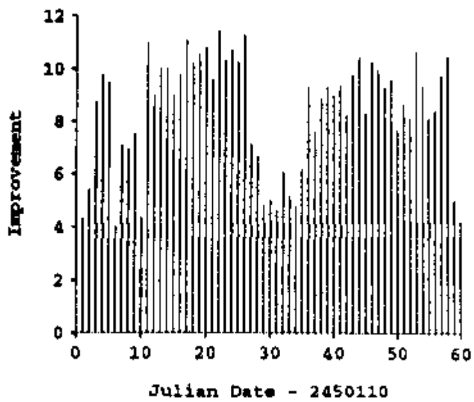

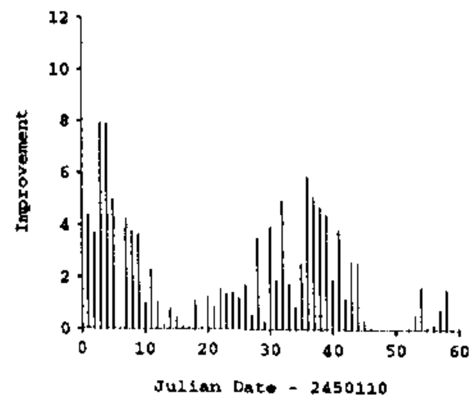

Figure 2: Performance improvement of GENH over ATIS dispatch (at left) and over the objective function as a greedy search heuristic (at right). Improvement is measured in terms of the standard deviation of each JD's QDF.

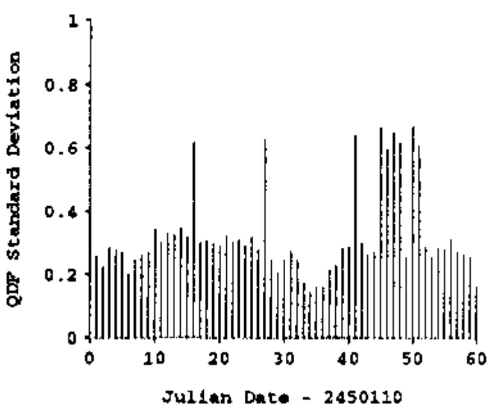

Figure 3: The standard deviation of each JD's QDF, based on the objective function scores of 200 randomly generated schedules; these are used in Figure 2 to express the performance improvement.

that using the objective function as the greedy heuristic is a standard approach to try in a search-based system.

The APA system maintains an extensive log of previous scheduling problem instances; this database was used for empirical evaluations of different versions of GENH. The experiments reported here test the one GENH version described in the previous section and use a set of sixty problem instances over the interval of days, expressed as Julian Dates (JDs), [2450111,2450170]. This allows for a range of problem characteristics in the test suite due to changes in the stars' relative positions, as well as in the phase and location of the moon; hence, they are sufficient to challenge the adaptive capabilities of GENH.

The results of these comparative analyses are shown in Figure 2. If we express performance improvement in terms of the difference between the objective function scores, then it would be difficult to interpret the significance of the improvement. To overcome this problem, we employ the *expected solution quality* (ESQj methodology [Bresina, *et al,* 1995] and express the improvement on a JD in terms of the standard deviation of that JD's *quality density function (QDF).* A QDF is a statistical estimate of the expected density of schedules within different quality ranges and is based on the objective function scores obtained via iterative sampling. The QDF standard deviation measure better indicates the significance of the improvement, and also takes into account the varying (relative) problem difficulty over the set of instances. For our experiment, each JD's QDF is based on 200 randomly generated schedules; Figure 3 plots the standard deviations of the sixty QDFs.

From the left plot in Figure 2, it is clear that greedy search with GENH's heuristics outperforms ATIS dispatch. The right plot shows the comparison between using a fixed heuristic (the objective function in this case) versus adapting the heuristic for each problem. Furthermore, since the seeding method used in these experiments starts GENH'S search with an heuristic very close to the objective function, the plot also illustrates the amount of improvement that the tuning method was able to achieve. Only on five of the sixty problems was GENH unable to make any improvement upon the objective function as a greedy heuristic: the consecutive JDs [2450157,2450161]. On ten of the problems, GENH makes a substantial improvement of over 4.0 QDF standard deviations.

Figure 4 illustrates the heuristics generated by GENH by plotting each attribute's assigned weight for the sixty problems. As is clear from the five plots, there is significant variance in the heuristics considered best by GENH, indicating that GENH did find relevant differences in the characteristics of the problem instances.

Although these results cannot be used to infer that GENH produces optimal schedules for the telescope observation problem, the results confirm that automatic adaptation of the scheduling heuristic for each problem improves the quality of schedules over non-adaptive ap-

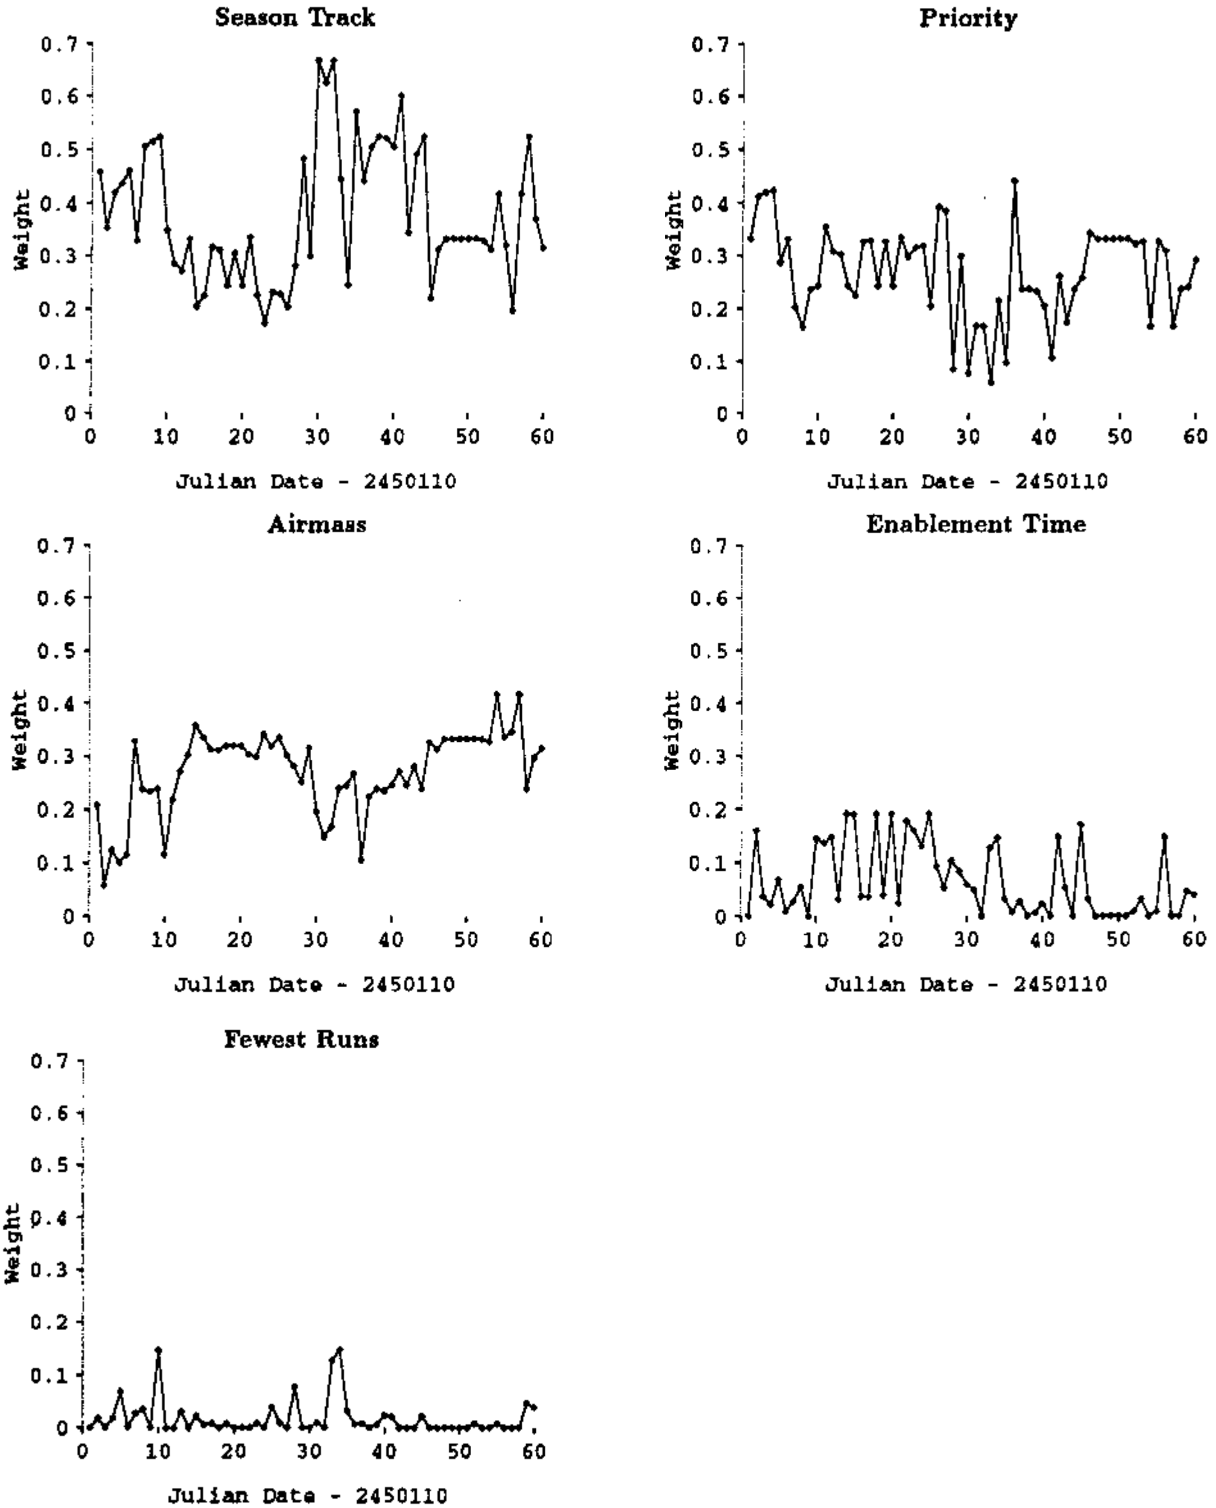

Figure 4: The heuristics generated by GEN H illustrated by the weights assigned to each of the five candidate attributes.

proaches for this domain, which was the primary purpose of these tests.

GENH is easily integrated into the APA system as a preprocessing module. Furthermore, adding GENH to the scheduling process is not prohibitively costly; we found that the particular version of GENH presented here took five to ten minutes to terminate.

5 Related Research

GENH resembles systems such as PALO [Greiner and Jurisica, 1992] which contain a mechanism for modifying a problem solver based on experience obtained from previous activity. The PALO algorithm incrementally produces a series of problem solvers (called *problem elements* or PEs) such that each element of the series is statistically likely to be an improvement over its predecessors over an ensemble of problems. PE selection is accomplished via a hill-climbing search in which candidates are evaluated using a performance cost function.

COMPOSER [Gratch and Chien, 1996] uses statistical hill-climbing to explore a space of heuristics. A candidate heuristic is adopted if it increases the expected performance of solving problems over a suite of problems. The adaptive version of the hybrid scheduler LR-26, incorporating both integer programming and constraint propagation, settles on a particular combination of heuristics after a training period. A utility function assigns positive values to heuristics that minimize computational effort.

The GENH algorithm also bears a resemblance to the family of local search algorithms (also known as repair search or iterative improvement) for solving Constraint Satisfaction Problems (CSPs) and optimization problems [Minton, *et* al., 1990; Morris, 1990]. Local search algorithms explore the space of candidate solutions (sometimes including invalid solutions) by performing local modifications to the current solution. At each step, the local solution modification that results in the greatest cost reduction, with respect to the given cost function, is selected. This process continues until a solution is reached in which no local changes can reduce the cost further. With GENH, of course, the goal is not to solve the problem, but rather to select an heuristic with which to subsequently solve the problem.

6 Future Research

A number of refinements to GENH are planned. At the time this work was carried out, the APA scheduler used a deterministic greedy search; this same scheduler search strategy was used in GENH to evaluate candidate heuristics. However, more recently, the scheduler's search strategy has been modified to incorporate *heuristic-biased stochastic sampling,* or HBSS, [Bresina, 1996]. HBSS stochastically explores the search space in the "neighborhood" of the greedy solution. The balance between heuristic adherence and exploration in HBSS is parameterized by specifying a ranking function and a bias function. At each decision point during search, these functions are used to assign selection probabilities to the different choices; a weighted stochastic selection is then made according to these probabilities. By varying these functions, HBSS encompasses a family of search algorithms of which greedy and random search are extreme members.

The current version of the APA scheduler first generates a schedule via ATIS dispatch, then performs a number of samples of HBSS; each sample generates a schedule. The best schedule found during these two search phases is then communicated to the telescope controller for execution. The evaluation metric for candidate heuristics used by GENH should be correlated with how the heuristic is used during solution search. We plan to compare, in terms of APA scheduler performance, GENH'S current evaluation metric based on greedy search with one based instead on "greedy sampling". *Greedy sampling* is an instance of HBSS in which all the choices tied with respect to the best heuristic score are assigned equal selection probabilities and all other inferior choices are assigned a zero selection probability. Hence, at each decision point, a random selection is made among the equally best choices.

7 Concluding Remarks

This paper has presented GENH, an approach to adaptive problem solving in the telescope observation scheduling domain. GENH adapts to the idiosyncrasies of a given problem instance by generating an effective scheduling heuristic. GENH performs a hill-climbing search through the space of possible heuristics using a series of local modifications to a seed heuristic. An innovative aspect of this work is the attention paid to user-defined, solution quality-based measures of heuristic performance, rather than performance measures based on search time to a solution. During GENH's search, each candidate heuristic is evaluated by employing the heuristic to find a schedule and then scoring the resulting schedule with the user-defined objective function.

The research that led to GENH was initially motivated by the idea that the role *of* a domain expert should be limited to constructing the domain-specific objective function; by contrast, the task of constructing the search heuristic is more appropriately performed by the system that generates the schedules. Before GENH, it was up to the user to carry out a "generate and test" exploration in the space of heuristics in order to find a satisfactory one. This is a time-consuming and difficult process, one that the user was not likely to repeat very often. However, as evidenced by the empirical results presented above, the heuristic that performs well for one day's scheduling problem may not do so well for the next one.

GENH represents progress in the automation of the schedule optimization process. Experiments conducted using GENH in the telescope observation scheduling domain demonstrate significant improvements over scheduling without adaptive heuristic selection, with little computational overhead incurred.

Acknowledgments

Thanks to Mark Drummond, Will Edgington, and Keith Swanson for their assistance with the APA and for their support of the work reported here. Thanks also to John Allen and Lina Khatib for helpful comments.

The first author was supported by the NASA-ASEE Summer Faculty Fellowship Program. The third author was supported by a NASA-Ames Research Center summer student fellowship.

References

[Boyd, *et a/.*, 1993] Louis Boyd, Donald Epand, John L. Bresina, Mark Drummond, Keith Swanson, D. Crawford, D. Genet, R. Genet, Greg Henry, G. McCook, W. Neely, P. Schmidtke, D. Smith, and M. Trublood. Automatic Telescope Instruction Set 1993. *LA.P.P.P. Communications,* 52.

[Bresina, 1996] John L. Bresina. Heuristic-Biased Stochastic Scheduling. *Proceedings of AAA1-96,* Portland, OR.

[Bresina, *et a/.*, 1995] John L. Bresina, Mark Drummond, and Keith Swanson. Expected Solution Quality. *Proceedings of IJCAI-95,* Montreal, Quebec.

[Bresina, *et a/.*, 1994] John L. Bresina, Mark Drummond, Keith Swanson, and Will Edgington. Automated Management and Scheduling of Remote Automatic Telescopes. *Optical Astronomy from the Earth and Moon,* ASP Conference Series, 55.

[Bresina, *et a/.*, 1996] John L. Bresina, Will Edgington, Keith Swanson, and Mark Drummond. Operational Closed-Loop Observation Scheduling and Execution. *Proceedings of the AAAI Fall Symposium on Plan Execution.*

[Drummond, *et al.;* 1995] Mark Drummond, John L. Bresina, Will Edgington, Keith Swanson, Greg Henry, and Ellen Drascher. Flexible Scheduling of Automatic Telescopes over the Internet. Henry & Eaton (eds). *Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy.*

[Drummond, *et a/.*, 1994] Mark Drummond, John L. Bresina, Keith Swanson, Will Edgington, and Greg Henry. The Associate Principal Astronomer Telescope Operations Model. *Proceedings of i-SAIRAS,* 351-354.

[Edgington, *et a/.*, 1996] Will Edgington, Mark Drummond, John L. Bresina, Greg Henry, and Ellen Drascher. Improved Scheduling of Robotic Telescopes. T. Boroson, J. Davies, and I. Robson (eds). *New Observing Modes for the Next Century.* ASP Conference Series, Vol. 87.

[Fisher, 1985] Marshall L. Fisher. An Applications Oriented Guide to Lagrangian Relaxation. *Interfaces* 15, 10-21.

[Gratch and Chien, 1996] Jonathan Gratch and Steve Chien. Adaptive Problem-Solving for Large-Scale Scheduling Problems: A Case Study. *J AIR* 4: 365-396.

[Greiner and Jurisica, 1992] Russell Greiner and Igor Jurisica. A Statistical Approach to Solving the EBL Utility Problem. *Proceedings of AAAI-92,* 241-248. Seattle, WA.

[Henry, 1996] Greg Henry. ATIS Dispatch Scheduling of Robotic Telescopes. T. Boroson, J. Davies, and I. Robson (eds). *New Observing Modes for the Next Century.* ASP Conference Series, Vol. 87.

[Laird, *et a/.*, 1986] John E. Laird, Paul S. Rosenbloom and Allan Newell. Universal Subgoaling and Chunking: The Automatic Generation and Learning of Goal Heuristics. Norwell, MA: Klewer Academic Publishers.

[Minton, *et a/.*, 1990] Steven Minton, Mark D. Johnston, Andrew B. Phillips, Philip Laird. Solving Large Scale Constraint Satisfaction and Scheduling Problems Using A Heuristic Repair Method. *Proceedings of AAAI-90,* 17-24. Boston, MA,

[Morris, 1990] Paul Morris. Solutions Without Exhaustive Search. *Proceedings of AAAI-90 Workshop on Constraint-Directed Reasoning.* Boston, MA.

[Morris, 1992] Paul Morris. On the Density of Solutions in Equilibrium Points for the Queens Problem. *Proceedings of AAAI-92,* San Jose, CA.