

Integrating Planning and Scheduling through Adaptation of Resource Intensity Estimates

Karen L. Myers¹, Stephen F. Smith², David W. Hildum², Peter A. Jarvis¹, Raymond de Lacaze¹

¹AI Center, SRI International

Menlo Park, CA USA

{myers,jarvis,delacaze}@ai.sri.com

²The Robotics Institute

Carnegie Mellon University, Pittsburgh, PA USA

{sfs,hildum}@cs.cmu.edu

Abstract

We describe an incremental and adaptive approach to integrating hierarchical task network planning and constraint-based scheduling. The approach is grounded in the concept of approximating the ‘resource intensity’ of planning options. A given planning problem is decomposed into a sequence of (not necessarily independent) subtasks, which are planned and then scheduled in turn. During planning, operators are rated according to a heuristic estimate of their expected resource requirements. Options are selected that best match a computed ‘target intensity’ for planning. Feedback from the scheduler is used to adapt the target intensity after completion of each subplan, thus guiding the planner toward solutions that are tuned to resource availability. Experimental results from an air operations domain validate the effectiveness of the approach relative to typical waterfall models of planner/scheduler integration.

Introduction

Goal-oriented activity in complex domains typically requires a combination of planning and scheduling. A manufacturing facility must develop process plans for ordered parts that can be cost-effectively integrated with current production operations. Military planners must select courses of actions that achieve strategic objectives, while making the most of available assets. Space observatories must allocate viewing instruments to maximize scientific return under a large and diverse set of causal restrictions and dependencies. Though conceptually decomposable, planning and scheduling processes in such domains can be and often are highly interdependent. Different planning options for achieving a given objective can make quite different demands on system resources; correspondingly, current resource commitments and

availability will impact the feasibility or desirability of various planning options.

The effectiveness of goal-oriented activity is ultimately tied to an ability to keep pace with evolving circumstances, and one recognized obstacle in practice is poor integration of *planning* and *scheduling* processes. In manufacturing organizations, this problem has been characterized as the “wall between engineering and manufacturing”. Similar sorts of barriers can be found in other large-scale enterprises. The crux of the problem is lack of communication. Plans are developed with no visibility of resource availability and operational status, and likewise, schedules are developed and managed without knowledge of objectives and dependencies. Without such information exchange, planning and scheduling processes are forced to each proceed in an uninformed and inherently inefficient manner. In the simplest case, the result is an *iterative waterfall* model of integration, where planning and scheduling are performed in sequential lockstep fashion and any problem encountered during scheduling simply triggers the generation of a new plan.

In this paper, we present a method for improving the overall planning and scheduling process through a tighter integration of these constituent activities. By planning, we refer generally to the process of deciding *what* to do; i.e., the process of transforming strategic objectives into executable activity networks. We use the term scheduling to refer alternatively to the process of deciding *when* and *how*; i.e., which resources to use to execute various activities and over what time frames. Traditionally, AI research has viewed planning and scheduling as distinct activities, and different solution techniques and technologies have emerged for each. Relatively few attempts have been made to combine respective technologies into larger integrated frameworks.

We take as our starting point previously developed technologies for hierarchical task network (HTN) planning and constraint-based scheduling. We describe and evaluate

an approach to their integration based on the idea of approximating the resource requirements (called *resource intensity*) of different planning options, and incrementally exchanging and exploiting information about likely resource shortfalls and excesses to settle on options that best utilize available resources. Finally, we present experimental results that compare an implementation of the method to an iterative waterfall model of integration within the air operations domain. These results show that the intensity-based approach provides plans of comparable quality for greatly reduced computation time.

Technology Foundations

Planning. The CPEF system provides the planning component for our work (Myers 1999). CPEF embodies a philosophy of plans as dynamic, open-ended artifacts that evolve in response to a continuously changing environment. CPEF provides a range of operations required for continuous plan management, including *plan generation, plan execution, monitoring, and plan repair*. Plan generation within CPEF is based on the CHIP system – an HTN planner derived from SIPE-2 (Wilkins 1988).

Scheduling. ACS, a constraint-based scheduler, provides the base scheduling capability. ACS is an air operations scheduler constructed using OZONE (Smith et al. 1996), a customizable constraint-based modeling and search framework for developing incremental scheduling tools. OZONE consolidates the results of application development experiences in a range of complex domains, including one recently deployed system for day-to-day management of airlift resources at the USAF Air Mobility Command (AMC) (Becker and Smith 2000). The ACS scheduler adapts techniques underlying the AMC application to the air operations domain. ACS can be used to generate, incrementally extend and revise assignments of aircraft and munitions to input target demands over time, taking into account priorities, desired levels of damage, time-on-target (TOT) windows, temporal sequencing constraints, feasible resource alternatives, and aircraft/munitions positioning and availability constraints.

Air Operations Domain Characteristics and Model

Applications that require integrated planning and scheduling will have individual characteristics that dictate the relative importance of each of these capabilities. Much of the work to date on combining AI planning and scheduling has focused on *resource-driven* domains (such as satellite observation scheduling (Muscettola et al. 1992)), which emphasize optimization of resource usage in

satisfying a pool of tasks. In contrast, the air operations domain has a more *goal-driven* flavor: while effective resource usage is important, the key motivation is to identify and schedule actions that will ensure attainment of stated objectives.

Objectives within the air operations domain reduce to goals of neutralizing enemy capabilities (e.g., antiaircraft capability, electricity production, communications) modeled as hierarchical networks that ground out at the level of specific targets. We provide several strategies for attacking different network types that vary in their aggressiveness, and hence resource demands. These strategies range from attacking all components in a network, to attacking a coherent subset, or an isolated node (Lee 1998).

Resources (i.e., aircraft, munitions) are assigned to support prosecution of individual targets. A given type of target usually has several possible aircraft/munitions configurations. However, different configurations will have different degrees of effectiveness, and hence the numbers of resources that must be allocated to achieve the desired effect can vary with each choice. Quantities (or capacities) of different types of resources are positioned at various locations nearby or within the geographic region of interest. The set of resources assigned to fly against a given target can vary in type and, depending on availability, may either originate from multiple locations (converging on the target within a particular time interval) or recycle from the same base location (making sufficient sets of consecutive strikes on the target).

The style of planning required for this domain differs markedly from standard AI approaches. Here, the search space is dense with solutions, making it easy to find a plan that satisfies stated goals. The real challenge is to find ‘good’ plans rather than settling for the first available solution. While most AI planning systems seek to minimize plan size, bigger plans tend to be better in this domain since the inclusion of additional activities can increase the likelihood of achieving stated objectives. For example, eliminating more of an enemy’s missile sites tends to improve the quality of a plan for neutralizing enemy attack capability. Note that maximizing plan size is not equivalent to maximizing resource usage: the planner and scheduler must still decide how to allocate available resources economically to support chosen activities.

Air operations commanders generally apportion a set of resources for a given set of high-level objectives; human planners are expected to develop solutions that maximize the likelihood of objective attainment while staying within the resource allotment. Our planning model incorporates this *apportionment perspective* into its design. In particular, initial plans seek to capitalize on all available resources; as resource problems arise, strategies are adopted that decrease resource usage.

Technical Approach

Our integration method builds on an incremental model of planning and scheduling that assesses resource feasibility at the level of *subplans* for the overall set of objectives, using a model of *intensity* to approximate resource demand, and *adaptation* in response to scheduler feedback.

Incremental Planning and Scheduling

Within our hierarchical domain model, high-level operator choices can have a significant impact on resource requirements. However, actions with specific resource requirements do not appear until the lowest levels of a deep hierarchy. For example, the high-level decision of whether to employ a passive or more proactive approach to defending assets will greatly influence resource requirements, although the actual missions that require resources are planned at much lower levels of abstraction.

Approaches in which complete layers of a hierarchical plan are forwarded to a scheduler for resource allocation (e.g., (Wilkins and Myers 1998)) do not provide much value in this case, since most of the plan would have to be completed before any scheduler feedback could be obtained. Instead, we developed a hybrid top-down/incremental model for planning and scheduling. The approach involves planning in standard HTN fashion down to a specified level of detail (the *decomposition layer*), and then splitting into subplans that are elaborated separately. The decomposition layer, defined implicitly in terms of specific goals, separates the higher-level strategic decisions that define overall plan structure from the planning of (mostly independent) lower-level objectives.

After completion of each subplan, the scheduler incrementally allocates resources to the new actions introduced by the subplan, taking into account the resource assignments already made for previous subplans. In the event that the scheduler is unable to produce a satisfactory resource assignment, the planner will modify one or more completed subplans to reduce resource demand, and then forward the revisions to the scheduler for appropriate adjustments to the current schedule. Once all outstanding resource problems have been resolved, the planner continues with generation of remaining subplans until completion of a full plan and schedule. With this incremental approach, the integrated plan and schedule is built in piecewise, incremental fashion, with adjustments made in response to detected resource problems.

This incremental approach would be ineffective for domains in which extensive strategic dependencies link objectives. However, in our models for the air operations domain, most dependencies occur at the level of resource allocation, thus enabling the separation of the planning for individual objectives.

Intensity Models of Resource Demand

To make informed decisions about its choices, a planner requires some model of the resource impact of its decisions. Previous work on incorporating resource feasibility reasoning into hierarchical planning (e.g., (Drabble and Tate 1994)) has assumed the ability to determine *a priori* minimum and maximum resource requirements for individual operators at all levels of abstraction, and has used this information as decision-making guidance.

Two problems arise with approaches of this type. First, computing bounds on resource usage can be prohibitively expensive in complex domains, given the need to consider all possible goal expansions and resource allocation options. Second, for the air operations domain, the bounds obtained are likely to be weak and uninformative. This latter problem stems from two factors: the heterogeneity of the resources that might be assigned to a given subplan, and the fact that resources are physically distributed and must travel variable amounts to perform different tasks. Depending on the type of resource assigned, different numbers of resources (or different amounts of resource capacity) will be required to accomplish a particular task. The location and operating characteristics of assigned resources will dictate the overall length of time that resources must be allocated. Since in both cases, the potential variance across resource types is quite high, simple minimum (or maximum) bounds will provide overly optimistic (or pessimistic) estimates of resource demand.

Given these problems, our approach to linking planning and scheduling instead builds on a heuristic characterization of expected resource usage by a planning operator, which we refer to as an operator's *intensity*. Our work to date has explored two models for intensity, which vary both the dimensionality (*single* vs. *multi*) and the precision (*qualitative* vs. *quantitative*).

Single-dimensional Qualitative Intensity Model. In this model, an operator's intensity represents a qualitative assessment of the operator's expected resource usage relative to alternatives for the same task. The air operations domain, for example, contains multiple operators for neutralizing an enemy's communication capability, ranging from taking out a single site, to destroying some select subset of communication devices, to eliminating all communication nodes. For an intensity scale of [0 10], the first operator might be ranked a 2, the second a 5, and the third a 10 to reflect their relative levels of expected resource consumption.

Multidimensional Quantitative Intensity Model. This model captures expected resource usage at a finer level of granularity. Resources are grouped into functional

categories intended to capture similarities in resource applicability. These groupings provide an aggregation over individual resource classes, thus simplifying the resource models inherent to the scheduler; however, the aggregation has greater detail than the single-dimensional intensity model and so would be expected to provide improved predictive value for resource usage estimation.

Within our air operations domain, for example, aircraft and munitions can be grouped according to the different types of missions in which they can be used (which is a function of target type). Our multidimensional intensity model for this domain groups 5 types of aircraft and 7 types of munitions into 4 resource dimensions. Because aircraft and munitions can be used for different types of missions, these dimensions are not mutually exclusive. This connectivity introduces additional complexity into the multidimensional intensity adaptation process, since decisions related to one dimension can impact results for others.

The multidimensional quantitative model also improves on the single-dimensional qualitative approach by employing a situation-dependent characterization of operator intensity. In particular, operator intensities are defined by a heuristic function that estimates resource demand based on the number and type of targets that an operator is expected to introduce.

The single-dimensional model has the virtue of requiring little effort to define the qualitative rankings within the underlying planning models: such rankings could be readily determinable by the knowledge engineer who develops the planning operators. In contrast, the multidimensional quantitative model requires the identification and modeling of resource abstractions. Such abstractions fall out naturally in the air operations domain but may be more problematic to define in others.

The weakness of the single-dimensional approach lies in its lack of granularity. Consider a situation with relatively low overall resource demand but where the class of resources required for a key type of action has been almost exhausted. The single-dimensional approach would not adjust strategy selection to adapt to the shortage because of the overall abundance of resources. In contrast, the multidimensional model can represent a lack of capacity for specialized groups of resources, thus enabling an adjustment in strategy selection to prefer approaches that minimize demand for the oversubscribed resource.

Intensity-based Adaptation

The incorporation of intensity information to guide planning occurs at the level of subplans. For a given subplan, the planner calculates a *target* intensity, denoted by I^T . This value represents the expected ‘ideal’ level of

resource usage for a particular subplan, relative to availability and expected demand for remaining subplans. When faced with a choice among multiple applicable operators O_i for a subgoal, the intensity I^{O_i} for each is computed. Each operator is assigned a rating $Rating(O_i)$ based on how closely its intensity matches the subplan’s target intensity, with the planner selecting the most highly rated operator for application. The specific definitions for the target intensity, operator intensity, and operator rating used in our work are presented below.

Adjustment of the target intensity across subplans enables the planner to adapt its strategy to match changing resource availability. The planner is provided with updates on resource availability after every interaction with the scheduler. Suppose that upon successful allocation of resources to a subplan, the scheduler’s assessment of remaining resource availability indicates a shortage (excess) of remaining resources relative to the subplans yet to be generated and scheduled. By reducing (increasing) the target intensity for the next subplan to reflect this shortage (excess), the planner will be biased toward selecting operators with lower (higher) intensity values that will decrease (increase) resource consumption levels. In this way, the planner dynamically adjusts its decision-making in response to scheduler feedback.

Within this adaptive framework, different control strategies can be defined for selecting the subplan to be revised in response to scheduling problems. The experiments in this paper adopt a *chronological backoff* strategy: when the scheduler encounters a problem with a subplan, the planner reduces the target intensity for that subplan in accord with a *target intensity reduction policy* and then generates an alternative plan. This process continues until either a resource feasible subplan is found, or there is no more room for intensity reduction. In the latter case, the algorithm removes the unsuccessful subplan from the plan; if the target intensity of the previous subplan can be reduced, then planning and scheduling are tried at that lower level; otherwise, the planner continues to remove subplans until it encounters a subplan that is not yet at the minimal intensity value. From that point, it tries to plan with the lower target intensity and then restarts the generation process in the forward direction.

Below, we provide the basic definitions for target intensity, operator intensity, operator rating scheme and target intensity reduction policy for the multidimensional case, followed by their definitions for the simpler single-dimensional case.

Target Intensity I^T . The target intensity for a given intensity dimension is defined in terms of the ratio of the resources available per remaining subplan to the resources allotted originally to each subplan (assuming uniform apportionment to each); this ratio is then normalized

relative to the interval of intensity values in use (namely, $[0, TopIntensity]$). More formally, let $Capacity(I_j)$ be the overall capacity for resources in dimension j and let R_j^i be the remaining capacity for dimension j after the first i of n subplans have been created and scheduled. The following equation defines the target intensity I^T for the $i+1^{st}$ subplan:

$$I^T = \begin{bmatrix} I_1^T \\ \vdots \\ I_m^T \end{bmatrix} \text{ where } I_j^T = \frac{\frac{1}{n-i} \times R_j^i}{\frac{1}{n} \times Capacity(I_j)} \times TopIntensity$$

Provided that resource usage remains below allotment levels, the value I_j^T will exceed $TopIntensity$. Values below $TopIntensity$ indicate that planning choices should decrease demand for resources within that dimension below the original allotment level.

Operator Intensity I^{O_i} . The intensity I^{O_i} of a planning operator O_i is defined by:

$$I^{O_i} = \begin{bmatrix} I_1^{O_i} \\ \vdots \\ I_m^{O_i} \end{bmatrix}$$

$$I_j^{O_i} = \frac{ExpDemand(O_i, I_j)}{\frac{1}{n} \times Capacity(I_j)} \times TopIntensity$$

The intensity for each dimension is defined to be the ratio of the expected resource demands introduced by the operator to the original allotment of resources for that subplan and dimension (assuming uniform allotment). For the air operations domain, the resource demands of an operator are measured in terms of the expected munitions and aircraft required to prosecute the targets associated with the operator. These estimates are calculated by summing the expected number of targets of a given type multiplied by a capacity estimate for the type.

Operator Ranking. Our scheme for ranking operators according to their proximity to the target intensity values is defined by the following equations. The ranking method builds on the *intensity difference vector* $D^{O_i} = I^T - I^{O_i}$, which gives the difference between the target intensity and operator intensity vectors.

$$Rating(O^i) = \sum_{d_j \in D^{O_i}} Penalty(d_j)$$

$$Penalty(d) = \begin{cases} P^+ \times d & \text{for } d \geq 0 \\ P^- \times ABS(d) & \text{for } d < 0 \end{cases}$$

The operator rating, denoted by $Rating(O^i)$, is defined to be the sum of the magnitudes in the intensity difference vector, adjusted by a *penalty factor*. In cases where the difference value d_j is positive (i.e., the operator requires fewer resources than indicated by the target intensity), the penalty is defined by P^+ ; in cases where $d_j < 0$ (i.e., the operator is expected to use more resources than indicated by the target intensity), the penalty is defined by P^- . Through appropriate settings of the ratio of these penalty factors, different strategies can be defined that penalize resource overutilization/underutilization to different degrees. With this rating scheme, the preferred operator will be that with the lowest rating.

Target Intensity Reduction Policy. In situations where the scheduler is unsuccessful in an attempt to allocate resources for a given subplan, it provides feedback to the planner in the form of the list of problematic resources whose limited capacity have contributed to the failure. The intensity reduction policy used to adjust the target intensity for that subplan incorporates this information. In particular, each intensity dimension that includes resources from the problematic set is decreased by an amount Δ . For the experiments presented in this paper, $\Delta = .25 \times TopIntensity$.

Single-Dimensional Case. For the single-dimensional case, the target intensity I^T reduces to

$$I^T = \frac{\frac{1}{n-i} \times R^i}{\frac{1}{n} \times Capacity} \times TopIntensity$$

The operator intensity is simply the qualitative annotation defined for the operator, while the rating is the difference between the target and operator intensities. The target intensity reduction policy consists of decreasing the current target intensity by Δ .

Experimental Evaluation

We conducted a series of experiments to evaluate the effectiveness of our intensity adaptation methods. Our test problem yields plans with eight subplans and 50 to 724 actions, depending upon the aggressiveness of the planning strategies applied. Experiments involved running the test problem with different resource profiles, as shown in Figure 1. The 100% profile provides just sufficient resources for the maximum plan; the profiles then decay gradually until there are insufficient resources to support the minimal plan. Additionally, the experiments employ a profile labeled BIG that contains a large amount of resources relative to the maximal plan.

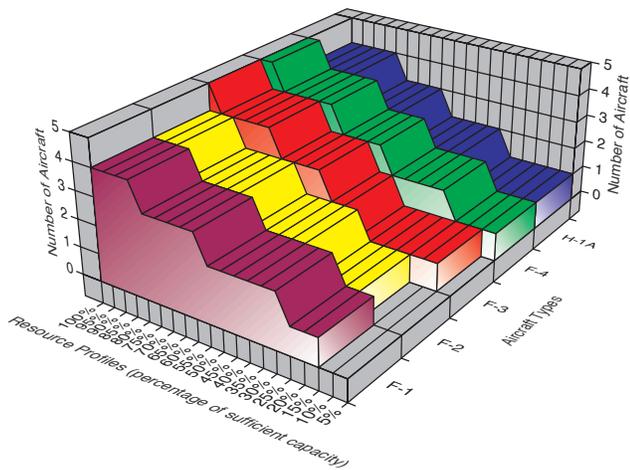


Figure 1: Experiment Resource Profiles

Generation time constitutes one important criterion for evaluating planner/scheduler behavior. Some measure of plan quality must also be considered. Otherwise, the best strategy is simply to generate the smallest plan that satisfies stated objectives: because it contains fewer activities, it will require fewer resources and so should be easier to schedule. Plan quality can be difficult to assess as it involves multiple dimensions and can be highly subjective (Gil 1998). As discussed above, air operations plans can generally be made more effective by adding more actions to them. For this reason, we use plan size as a rough indicator of plan quality.

For a baseline, we adopted a loosely coupled iterative waterfall integration of the planner and scheduler in which the planner generates complete plans and then passes them to the scheduler for resource allocation and time-on-target assignments. If the scheduler fails to produce a feasible schedule, the process repeats with the planner performing chronological backtracking to generate alternative plans. To draw fair comparisons with the intensity-based approaches, the waterfall method considers operators in decreasing order of intensity. This strategy generally yields a plan that is close to the largest supportable for the available resources but is not necessarily optimal (i.e., chronological backtracking stops at the first solution, even though undoing an earlier operator choice might enable more aggressive subsequent choices).

Our evaluation consists of two experiments. Experiment A compares the single-dimensional and multidimensional approaches (with $P^+ = P^- = 1$) to the iterative waterfall. Experiment B assesses the sensitivity of the multidimensional method to the *penalty factors* P^+ and P^- . For each, we consider three performance factors: generation time, plan size, and number of planner/scheduler interactions.

Experiment A: Intensity Adaptation Evaluation

Figure 2 shows the results for Experiment A. The upper-left graph displays generation time for the three methods. As can be seen, the waterfall method requires substantially more time when resources become constrained, while the intensity-based methods perform much better. The multidimensional approach outperforms the single-dimensional approach, with the advantage increasing as resource availability drops. The upper-right graph displays the number of interactions between the planner and scheduler required to find a solution. As with generation time, these results show that the multidimensional method outperforms the single-dimensional method, and that they both are far superior to the waterfall method as resource availability decreases.

Experiment A used a scaled-down version of our air operations domain in which goals that do not involve intensity decisions are limited to a single applicable operator. This restriction was introduced to ensure that the waterfall backtracking was limited to precisely the same choices as the intensity adaptation methods, in essence providing the best possible comparative analysis conditions for the waterfall model. An additional experiment was run where non-intensity goals had two applicable operators. Runtimes for the intensity methods were virtually identical to those in Figure 2 since the intensity method backtracks at the level of intensity values rather than operators (hence, it is not impacted by the additional operators). In contrast, the waterfall method was unable to find a solution below the 100% resource profile after 239 trials and almost 30 hours of runtime. The waterfall method fails so badly in this larger problem because many planning decisions must be backtracked over to reach one that impacts resource usage significantly.

The waterfall approach produces slightly larger plans than the intensity-based methods for the 100% through 50% profiles; as resource availability decreases further though, it produces smaller (i.e., less aggressive) plans. In comparing runtimes, it is clear that the small increases in plan size come at the cost of an increase of several orders of magnitude in planning/scheduling time. While there is some variation between the single-dimensional and multidimensional methods, the difference is relatively small. Overall, these results show that the performance benefits realized by the multidimensional approach do not adversely impact solution quality.

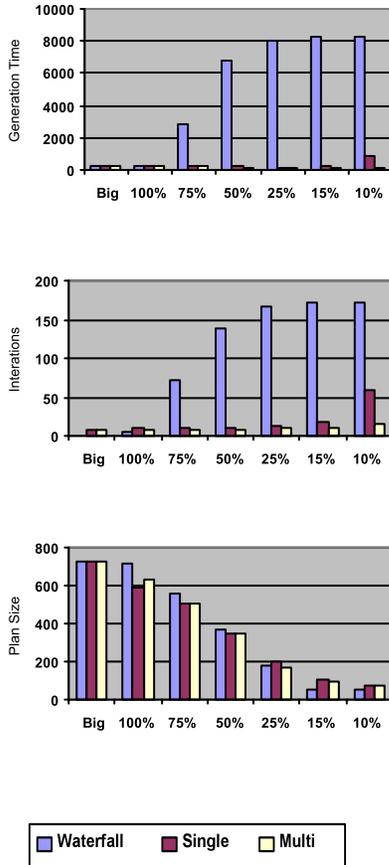


Figure 2: Comparison of the Waterfall, Single-dimensional, and Multidimensional Methods

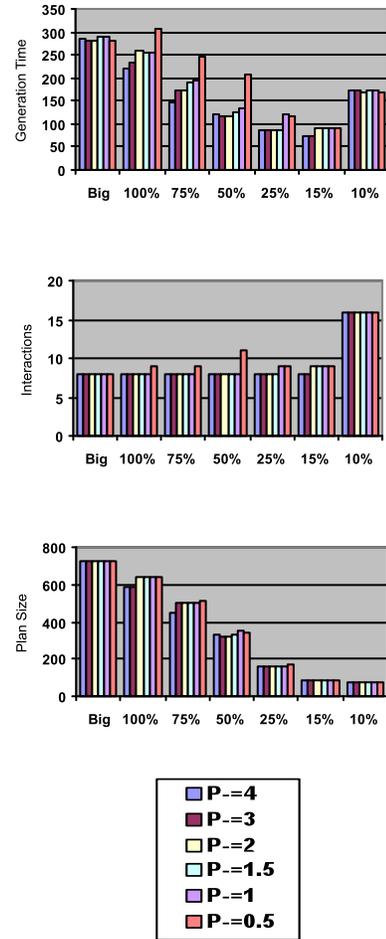


Figure 3: Sensitivity of the Multidimensional Approach to the Ratio of P^- to P^+

Experiment B: Sensitivity to the ratio of P^- to P^+

As noted above, the ratio of P^- to P^+ in the operator ranking scheme for the multidimensional approach can be adjusted to vary the penalty for overutilization/underutilization of resources relative to the established target intensity. To assess sensitivity to these values, we ran test cases with $P^+=1$ and P^- ranging from 0.5 to 4.

Figure 3 displays the results. For $P^-=4$ (and to some extent, $P^-=3$), there is a noticeable drop in plan size for the 100% through 50% profiles. For $P^-=5$, generation times and the number of planner/scheduler interactions are appreciably higher over that same range. Such results are to be expected: when resource overutilization is penalized relative to underutilization (i.e., $P^-/P^+ > 1$), the intensity adaptation process will be more cautious, resulting in a tendency toward smaller plans. In contrast, when resource

overutilization is favored relative to underutilization (i.e., $P^-/P^+ < 1$), the intensity adaptation process will be more aggressive in its strategy selection, possibly resulting in the need for more backtracking due to overly aggressive strategy choices.

We had expected to see more dramatic variation as P^- changed but the adaptive nature of the intensity method appears to compensate for overly aggressive or weak decisions induced by large/small penalty ratios. This robustness makes the intensity adaptation approach strongly insensitive to reasonable values for parameters P^- and P^+ .

Related Work

As mentioned earlier, much of the previous work in integrated planning/scheduling systems has been motivated by resource-driven applications. The early Hubble Space Telescope scheduling application of the HSTS system (Mussettola et al. 1992) provides a representative example, where a set of independent (or loosely-coupled) requests for telescope viewing time, each requiring a complex set of spacecraft actions for setup, observation, and cleanup, must be selected and sequenced for execution. Here, the overriding concern is efficient allocation of system resources, with planning decisions localized to implementation of individual tasks. The Remote Agent Planner/Scheduler (Jonsson et al. 2000) and the ASPEN mission planner (Rabideau et al. 1999) also fall into this category, as does IP3S (Sadeh et al. 1998), a system that integrates process planning and production scheduling in the manufacturing domain.

The REALPLAN system (Srivastava et al. 2001) places greater emphasis on strategic planning. Like our approach, REALPLAN partitions a problem into separate planning and scheduling components rather than solving the entire problem in a single integrated search space (see (Smith et al. 2000) for a survey of integrated search approaches). We similarly believe that such partitioning provides essential computational leverage. REALPLAN employs an iterative waterfall control model, with feedback of failure information in the most sophisticated variant. As shown in this paper, such an approach can be intractable in nontrivial domains.

The CIRCA-based planning and scheduling system described in (McVey et al. 1997) builds on an iterative waterfall model of interaction, but incorporates feedback from the scheduler to planner that is similar in spirit to our intensity adaptation approach. Based on a probabilistic state model, the planner generates *control plans* designed to prevent runtime transition to failure states. Planning relies on a specified probability threshold on states, with higher thresholds leading to consideration of fewer eventualities and simpler plans. When the scheduler is unable to meet the stated deadlines of all actions in a generated plan, it recommends a higher probability threshold to the planner for the next iteration. Similarly, when schedules underutilize resources, the scheduler suggests a lower probability threshold to enable the incorporation of additional activities.

Conclusions

The two intensity-based methods presented in this paper provide complementary methods for supporting effective planner/scheduler integration in domains that require significant strategic planning. The single-dimensional

qualitative approach provides a simple, easily implemented method that shows significant performance gains over waterfall-style methods. The multidimensional quantitative approach provides even better results but requires somewhat more modeling effort to operationalize.

This work represents one thrust of a larger effort to develop an integrated planning and scheduling system for management and control of large-scale enterprises (Myers and Smith 1999). Beyond the work on plan and schedule generation described here, we are developing intensity-based methods to support efficient plan and schedule repair in response to the addition or revision of objectives and changes to resource availability.

Acknowledgments. The work reported in this paper was supported by DARPA contracts F30602-97-2-0066 (CMU) and F30602-97-C-0067 (SRI) under the supervision of Air Force Research Laboratory – Rome.

References

- Becker, M., and Smith, S.F. 2000. Mixed-Initiative Resource Management: The AMC Barrel Allocator. In *Proceedings of the 5th International Conference on AI Planning and Scheduling*.
- Drabble, B., and Tate, A. 1994. The Use of Optimistic and Pessimistic Resource Profiles to Inform Search in an Activity Based Planner. In *Proceedings of the 2nd International Conference on AI Planning and Scheduling*.
- Gil, Y. 1998. On Evaluating Plans. Technical Report, University of Southern California Information Sciences Institute.
- Jonsson, A.; Morris, P.; Mussettola, N.; Rajan, K.; and Smith, B. 2000. Planning in Interplanetary Space: Theory and Practice. In *Proceedings of the 5th International Conference on AI Planning and Scheduling*.
- Lee, T. 1998. The Air Campaign Planning Knowledge Base. Technical Report, SRI International.
- McVey, C.B.; Atkins, E.M.; Durfee, E.H.; and Shin, K.G. 1997. Development of Iterative Real-Time Scheduler to Planner Feedback. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*.
- Muscettola, N.; Smith, S.F.; Cesta, A.; and D'Aloisi, D. 1992. Coordinating Space Telescope Operations in an Integrated Planning and Scheduling Framework. *IEEE Control Systems*, 12(1):1369 - 1376.
- Myers, K.L. 1999. CPEF: A Continuous Planning and Execution Framework. *AI Magazine* 20(4).
- Myers, K.L., and Smith, S.F. 1999. Issues in the Integration of Planning and Scheduling for Enterprise Control. In *Proceedings of the DARPA Symposium on Advances in Enterprise Control*.
- Rabideau, G.; Knight, R.; Chien, S.; Fukunaga, A.; and Govindjee, A. 1999. Iterative Repair Planning for Spacecraft Operations using the ASPEN System. In *Proceedings of the International Symposium of Artificial Intelligence, Robotics and Automation for Space*.
- Sadeh, N.; Hildum, D.W.; LaLiberty, T.J.; McAnulty, J.; Kjenstad D.; and Tseng, A. 1998. A Blackboard Architecture for

- Integrating Process Planning and Production Scheduling. *Concurrent Engineering: Research & Applications* 6(2).
- Smith, S.F.; Lassila, O.; and Becker, M. 1996. Configurable Systems for Mixed-Initiative Planning and Scheduling. In *Advanced Planning Technology* (ed. A. Tate), AAAI Press.
- Smith, D.E.; Frank, J.; and Jonsson, A.K. 2000. Bridging the Gap Between Planning and Scheduling. *Knowledge Engineering Review* 15(1):47-83.
- Srivastava, B.; Kambhampati, S.; and Minh, B.D. 2001. Planning the Project Management Way: Efficient Planning by Effective Integration of Causal and Resource Reasoning in RealPlan. *Artificial Intelligence* 131(1-2):73-134.
- Wilkins, D.E. 1988. *Practical Planning: Extending the Classical AI Planning Paradigm*. Morgan Kaufmann.
- Wilkins, D.E. and Myers, K.L. 1998. A Multiagent Planning Architecture. In *Proceedings of the 4th International Conference on AI Planning Systems*.