

# Integrating Planning and Scheduling through Intensity Adaptation

Karen L. Myers<sup>1</sup>, Stephen F. Smith<sup>2</sup>,  
David W. Hildum<sup>2</sup>, Peter A. Jarvis<sup>1</sup>, Raymond de Lacaze<sup>1</sup>

<sup>1</sup> Artificial Intelligence Center  
SRI International  
Ravenswood Avenue  
Menlo Park, California 94025  
{myers,jarvis,delacaze}@ai.sri.com

<sup>2</sup> The Robotics Institute  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
{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, different 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 towards 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 dynamics of the operating environment complicate matters further, requiring efficient response to continual unexpected changes to system objectives and resource availability.

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 98). 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 88) – and the Advisable Planner (Myers 96). The Advisable Planner provides an advice-taking layer on top of CHIP that enables a user to guide and direct the plan generation toward solutions that match his or her individual preferences.

**Scheduling** ACS, a constraint-based scheduler, provides the base scheduling capability. ACS is an air operations scheduler constructed using OZONE (Smith, Lassila and Becker 96), 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 00).

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

Applications that require integrated planning and scheduling will have individual characteristics that dictate the relative importance of each capability. 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 92)), 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., anti-aircraft capability, electricity production, communications) modeled as hierarchical networks that ground out at the level of specific targets. We provide several strategies for attacking networks, ranging from attacking all components in a network, to attacking a coherent subset, or an isolated node (Lee 98).

Resources are assigned to support prosecution of individual targets. For a given target type, several aircraft/munitions combinations might be viable, although the numbers required to achieve the desired effect may vary with each choice. Capacities of different resources are positioned at various locations. The set of resources assigned to fly against a target can vary in type and may originate from multiple locations. This interchangeability of resources allows higher-level abstraction of capacities into intensity ‘dimensions’ (see below).

The style of planning within 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 any available solution. While most AI planning systems seek to minimize plan size, bigger plans tend to be better in this domain. 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 potential 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 models incorporate this *apportionment perspective* into their design. In particular, initial plans seek to capitalize on all available resources; as resource problems arise, strategies are adopted that decrease expected 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 domain model, actions with resource requirements do not appear until the lowest levels of a deep hierarchy. Approaches in which complete layers of a hierarchical plan are forwarded to a scheduler for resource allocation (e.g., (Wilkins and Myers 98)) 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 planned 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 performs incremental resource allocation for the actions introduced by the subplan, relative to resource assignments made for previous subplans. In the event that the scheduler is unable to produce a satisfactory resource assignment, the planner will modify a completed subplan 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.

Our 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. The incremental approach has the added benefit that it can be used for dynamically extending plans to include additional objectives as plan execution unfolds.

### Intensity Models of Resource Demand

In order to make informed decisions about its choices, a planner requires some model of the resource impact of its decisions. Although the specific actions that require resources don't appear until the lowest levels of our hierarchical models, high-level decisions have a great impact on resource requirements. For example, the 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 detail. For this reason, our approach to linking planning and scheduling 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 intensity represents a qualitative assessment of the operator's expected resource usage relative to alternatives for the same task. Our 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. When using 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 the individual resources 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. Our multidimensional intensity model for this domain groups 5 types of aircraft and 7 types of munitions into 4 distinct resource dimensions.

This model further improves on the single-dimensional qualitative approach by employing a situation-dependent characterization of intensity for an operator. In particular, operator intensities are defined by a heuristic function that estimates resource demand based on the number 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 should be readily assessable by the knowledge engineer who develops the planning operators (especially since they need only be approximate). 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 other domains.

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 be unable to 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, and hence strategy selection can be adapted to prefer approaches with minimal 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' use of resources 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 then selecting the most highly rated operator for application. Through appropriate definitions for the target intensity, operator intensities, and the operator rating strategy, the planner can adapt its strategy to match changing resource availability.

For example, scheduler feedback could indicate a shortage/excess of remaining resources, relative to the subplans yet to be generated and scheduled. Such a shortage/excess would be reflected in the setting of the next target intensity at a lower/higher level; the planner would then be biased towards selecting operators with lower/higher intensity values to reduce/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 schedulability problems. The experiments reported on in this paper adopt a *chronological backoff* strategy: when the scheduler encounters a problem with a subplan, the planner

successively reduces the intensity of that subplan by some factor until a successful 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 reaches a point where 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 the target intensity, operator intensity, and operator rating scheme for the multidimensional case, followed by a discussion of 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:

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

Provided that resource usage remains below allotment levels, the value of  $I_j^T$  will exceed  $TopIntensity$ . Values below  $TopIntensity$  indicate that planning choices should seek to 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 the equation:

$$I^{O_i} = \begin{pmatrix} I_1^{O_i} \\ \vdots \\ I_m^{O_i} \end{pmatrix} \quad \text{where} \quad I_j^{O_i} = \frac{\frac{ExpectedDemand(O_i, I_j)}{\pi}}{\frac{1}{\pi} \leftarrow Capacity(I_j)} \leftarrow TopIntensity$$

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** Figure 1 presents our scheme for ranking operators according to their proximity to the target intensity values. 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. 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 over/underutilization to different degrees. With this rating scheme, the preferred operator will be that with the lowest rating.

$$Rating(O_i) = \sum_j \frac{Penalty(d_j)}{d_j \cdot D^{O_i}}$$

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

Figure 1: Operator Rating

**Single-Dimension Case** For the single-dimensional case, the target intensity  $I^T$  reduces to the following:

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

The operator intensity is simply the qualitative annotation defined for the operator, while the rating is the difference between the target and operator intensities.

## Experimental Evaluation

We conducted a series of experiments to evaluate the effectiveness of the intensity adaptation methods. 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 an acceptable schedule, the plan generation and scheduling process repeats.

Our test problem yields plans with 8 subplans and 50-724 actions, depending upon the aggressiveness of the planning strategies applied. Experiments involved running the test problem with different resource profiles, as shown in Figure 2. The BIG profile contains a large amount of resources relative to the maximal plan, while the 100% profile provides just sufficient resources to enable the maximum plan to be constructed; the profiles then decay gradually until there are insufficient resources to support the minimal plan.

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 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 98). As discussed above, air operations plans can generally be




Figure 2: Experiment Resource Profiles

made more effective by adding additional actions to them. For this reason, we use plan size as a rough indicator of plan quality.

Our experiment compares the single- and multidimensional approaches (with  $P^+ = P^- = I$ ) to the iterative waterfall. For performance, we consider three main factors: generation time, plan size, and number of planner/scheduler interactions. Figure 3 shows the results.

The topmost graph displays generation times in log scale 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 also outperforms the single dimensional approach. The middle 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 resources become more limited.

These results show an impressive speed-up by the intensity adaptation methods over the waterfall baseline. The tests used a scaled-down version of the domain in which goals that did not involve intensity decisions were limited to a single applicable operator. As such, the waterfall backtracking was limited to the same choices as the intensity adaptation methods. An additional experiment was run where non-intensity goals had 2 applicable operators. Runtimes for the intensity methods were virtually identical to those in Figure 3, since the intensity method backtracks at the level of intensity values rather than operators (hence 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 through to reach one that impacts resource usage significantly.

The waterfall approach produces larger plans than the intensity-based methods for the 100% through 50% profiles; as resource availability decreases though, it produces smaller (i.e., less aggressive) plans. In comparing the corresponding runtimes, it is clear that the small




Figure 3: Comparison of Waterfall, Single-dimensional and Multidimensional Methods

increase in plan size is obtained at the cost of an increase of several orders of magnitude in planning/scheduling time. While there is some variation between the single 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.

Additional experiments, documented in (Myers, Smith et al 2001), provide further evidence of the effectiveness of the intensity adaptation method for integrating strategic planning and scheduling. One set considers variants to the ground interdiction problem described here; a second set employs an air superiority problem with related but different structure. These experiments produced similar results to those reported in this paper, thus providing evidence for the generality of the method. We also performed sensitivity analyses for the 2 key parameters to the intensity adaptation method: the resource over/underutilization ratio in the operator ranking scheme ( $P^+ / P^-$ ) and the distribution of resources among intensity dimensions. These experiments show that the adaptive

nature of the intensity method makes it robust in the face of reasonable variations in these parameters.

## 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 (Muscettola et.al 92) 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 clean-up, 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. 00) and the DCAPS mission planner (Chien et.al.) also fall into this category, as does IP3S (Sadeh et.al 98), a system that integrates process planning and production scheduling in the manufacturing domain.

The REALPLAN system places greater emphasis on strategic planning (Srivastava, Kambhampati and Minh). Like our approach, REALPLAN partitions a problem into separate planning and scheduling subproblems rather than solving the entire problem in a single integrated search space (see (Smith, Frank, Jonsson 01) 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 planning and scheduling system described in (McVey, et.al. 97) builds on an iterative waterfall model of interaction, but incorporates a form of feedback from 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. Plans are generated based on a specified probability threshold on states, with higher thresholds resulting in consideration of fewer eventualities and hence 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 as feedback to the planner for the next iteration. Similarly, when schedules underutilize resources, the scheduler suggests a lower probability threshold for the planner to enable the incorporation of additional activities.

## Conclusions

The two intensity-based methods defined 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 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 99). Beyond the work on plan and schedule generation described here, we are also 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.

## References

- Becker, M. and S.F. Smith, "Mixed-Initiative Resource Management: The AMC Barrel Allocator", *Proc. 5<sup>th</sup> Intl. Conf. on AI Planning and Scheduling*, 2000.
- Chien, S., G. Rabideau, J. Willis and T. Mann, "Automated Planning and Scheduling of Shuttle Payload Operations", *Artificial Intelligence*, 114(1-2), 1999.
- Gil, Y. "On Evaluating Plans", Technical Report, USC/ISI. 1998.
- Jonsson, A., P. Morris, N. Muscettola, K. Rajan and B. Smith, "Planning in Interplanetary Space: Theory and practice", *Proc. 5<sup>th</sup> Intl. Conf. on AI Planning and Scheduling*, 2000.
- Lee, T. "The Air Campaign Planning Knowledge Base", SRI International Technical Report, 1998.
- McVey, C.B., E.M. Atkins, E.H. Durfee and K.G. Shin, "Development of Iterative Real-Time Scheduler to Planner Feedback", *Proc. 16<sup>th</sup> Intl. Joint Conf. On AI*, 1997.
- Muscettola, N., S.F. Smith, A. Cesta and D. D'Aloisi, "Coordinating Space Telescope Operations in an Integrated Planning and Scheduling Framework", *IEEE Control Systems*, 12(1), 1992.
- Myers, K. L., "Strategic Advice for Hierarchical Planners", *Principles of Knowledge Representation and Reasoning: Proc. of the 5<sup>th</sup> Intl. Conf.*, Morgan Kaufmann, 1996.
- Myers, K.L., "CPEF: A Continuous Planning and Execution Framework", *AI Magazine*, 20(4), 1999.
- Myers, K.L. and S.F. Smith, "Issues in the Integration of Planning and Scheduling for Enterprise Control", *Proc. DARPA Symposium on Advances in Enterprise Control*, 1999.
- Myers, K.L., S.F. Smith *et al*, "Cooperative Planning and Scheduling", in preparation, 2001.
- Sadeh, N., D.W. Hildum, T.J. LaLiberty, J. McAnulty, D. Kjenstad and A. Tseng, "A Blackboard Architecture for Integrating Process Planning and Production Scheduling", *Concurrent Engineering: Research & Applications*, 6(2), 1998.
- Smith, D.E., J. Frank, and A.K. Jonsson, "Bridging the Gap Between Planning and Scheduling", *Knowledge Engineering Review*, 15 (1) 2000.
- Smith, S.F., O. Lassila and M. Becker, "Configurable Systems for Mixed-Initiative Planning and Scheduling", in *Advanced Planning Technology*, (ed. A. Tate), AAAI Press, 1996.
- Srivastava, B., S. Kambhampati and B.D. Minh, "Planning the Project Management Way: Efficient Planning by Effective Integration of Causal and Resource Reasoning in RealPlan", *Artificial Intelligence*, to appear, 2001.
- Wilkins, D. *Practical Planning: Extending the Classical (AI) Planning Paradigm*, Morgan Kaufmann, 1998.
- Wilkins, D. and K. Myers, "A Multiagent Planning Architecture", *Proc. 4<sup>th</sup> Intl. Conf. on AI Planning Systems*, 1998.