

# Coordinating Activity in Knowledge-Intensive Dynamic Systems

Stephen F. Smith  
The Robotics Institute  
Carnegie Mellon University  
Pittsburgh PA 15213  
sfs@cs.cmu.edu

**Abstract**—In this paper we consider the problem of coordinating activity in knowledge-intensive dynamic systems (KIDS) -- large-scale, multi-agent systems in which agents with significant individual capabilities work together to accomplish complex, knowledge-intensive tasks. Given the time pressure and resource limitations under which a KIDS typically operates, the establishment of plans and schedules can significantly improve organizational performance. However, there are several complicating factors: (1) there is diversity and novelty in the structure of processes that must be executed over time, requiring tight coupling of action selection with resource allocation, (2) processes are unpredictable in their outcomes and require continual dynamic adjustment and revision, (3) the collective capabilities of the agents of a KIDS are its primary asset and, to minimize future resource limitations, task allocation should consider the side effects of acquired expertise, and (4) KIDS are large-scale enterprises, requiring an ability to effectively distribute decision-making. Our previous research has developed constraint-based search techniques for continuous, dynamic scheduling that have been successfully applied to complex, large-scale transportation and manufacturing domains. We outline current work aimed at extending these models to address the above issues and provide an effective basis for managing KIDS.

## 1. INTRODUCTION

Knowledge-intensive dynamic systems (KIDS) present complex scheduling and coordination problems. By KIDS, we refer generally to large-scale, multi-actor systems in which agents with significant individual capabilities work together to accomplish complex, knowledge-intensive tasks. Examples of KIDS include large news organizations such as CNN and AP, the research and development division of a large company, and the World Bank. KIDS plan and execute production processes which: (1) principally involve the collection, manipulation and management of knowledge products, (2) exhibit diversity and perpetual novelty in process structure, and (3) are unpredictable in their execution and require continual adjustment and revision.

The establishment of a schedule is crucial to effective management and control of KIDS; it is the means by which global coordination of executing agents (actors) is achieved and maintained. Thus, as evolving requirements, changing priorities, and new resource availability constraints continually force changes to previously planned tasks and resource assignments, attention must be paid to the cost of continual solution change and to maintaining continuity in currently executing processes. Like many other practical domains, KIDS are further complicated by the presence of multiple performance objectives, the need to optimize under complex (and often idiosyncratic) constraints and the need for flexible human decision-making involvement. KIDS are also systems whose components build expertise as processes are performed, and hence constituent resource capabilities evolve over time. Finally, KIDS are large-scale enterprises.

The scheduling requirements inherent in the coordination of KIDS are at odds with classical scheduling models and most current scheduling tools, which are conceived and designed as batch-oriented, black box solution generators. Most basically, this type of approach advocates a mode of operation wherein the scheduling problem is (re)solved from scratch each time there is a change to any problem constraint or requirement. This presents significant problems for use in dynamic and continuous scheduling environments, where such changes are the order of the day. First, there is generally no persistence of decisions over time. Small changes in inputs can lead to large changes in outputs, and it is quite difficult to control and manage solution change. Second, a resolve-from-scratch paradigm presents obvious computational barriers to integrating with execution; in large-scale domains it becomes problematic for the scheduler to keep pace with execution events, and its use becomes limited to off-line analysis. In general, classical scheduling models and tools are not well-suited for dynamic continuous scheduling environments like KIDS.

Recent research in incremental, constraint-based search has focused on the development of scheduling techniques that overcome these limitations and provide a practical basis for dynamic, continuous scheduling. These search procedures permit flexible, controlled management of schedules in response to execution dynamics over time, under quite general assumptions regarding the representation and incorporation of various domain constraints. This approach has been successfully applied to a number of complex, dynamic transportation and manufacturing scheduling

domains [10,11,14]. Our current research aims at investigating its applicability to KIDS.

In this paper, we outline some of the issues involved in applying this approach to scheduling to the problem of coordinating activity in KIDS. First we summarize the capabilities of constraint-based scheduling models that make them attractive in this context. Next we consider extensions to the current technology base that are necessary to enable effective, scalable management and control of KIDS.

## 2. CONSTRAINT-BASED SCHEDULING MODELS

From the standpoint of managing a knowledge intensive dynamic system, constraint based search models offer the following key capabilities:

- *Dynamic, incremental solution change and (re) optimization* – As suggested above, effective management of operations in dynamic domains requires controlled change; the ability to preserve previous commitments and minimize the ripple effect on currently executing processes while re-optimizing to address new requirements and changing priorities is fundamental to efficient global performance, and these incremental scheduling capabilities are the hallmark of constraint-based technologies (e.g., [8, 10, 11, 17, 18]).
- *Flexible, user-driven problem solving and option generation* – Constraint-based scheduling technologies provide a framework for interactive adjustment of constraints and objective criteria to explore tradeoffs and generate options [2, 7, 6, 13]. The need for such capabilities is crucial in situations where all constraints cannot be satisfied (e.g., there are more tasks that could be undertaken in a given time frame than current capabilities will allow). Given the time pressure and over-subscribed nature of resources in the case of KIDS, the ability to rapidly generate and evaluate various options (e.g., relaxed deadline, alternative process, etc.) is a second central component of efficient global performance.
- *Coordinated management of resources under complex and changing constraints* – Constraint-based scheduling technologies provide a rich framework for specifying and enforcing the myriad of idiosyncratic constraints encountered in KIDS scheduling environments, and for generating schedules that optimize resource usage under realistic assumptions [1,3]. Constraint-based scheduling technologies similarly accommodate the introduction of new constraints and incremental changes to existing constraints, allowing robust management of ill-structured production processes.
- *Collaborative coordination of multiple scheduling agents* – The above capabilities provide an infrastructure for collaborative negotiation and conflict resolution among hierarchically organized scheduling agents [12]. As scheduling decisions and authority are

pushed down to process level managers and executing agents, or as unsupportable assignments are pushed upward from executing agents to more global levels, the incremental constraint relaxation and solution change capabilities can be used to efficiently explore alternative options, propose schedule changes and achieve agreement.

## 3. RESEARCH ISSUES

Given the above capabilities, constraint-based scheduling procedures appear well-suited for managing activity in large-scale, dynamic systems. However, their application to the scheduling and management of KIDS presents some significant new research challenges. First, the models used by work to date in incremental, constraint-based scheduling make some rather restrictive assumptions from a KIDS perspective. Research has focused primarily on problems of synchronizing resource usage among large sets of well-defined processes (in contrast to the more customized and idiosyncratic processes encountered in KIDS). Research has also generally ignored issues relating to the task content, quality and outcomes, restricting the focus of rescheduling actions to problems involving task timing and resource unavailability. While such actions are certainly relevant to managing KIDS, capabilities for reasoning about and responding to information related to task outcomes are central. Another broad challenge in the context of managing KIDS is of course scalability. Though incremental scheduling is inherently well-equipped for application in large-scale domains, the ability to distribute overall scheduling and control responsibility remains a fundamental issue.

In the subsections below, we consider these research issues in more detail.

### *Dynamic Management of Changing Process Structure*

One immediate consequence of a decision to incorporate information relating to the outcomes of specific executing tasks is the need for process representations that relate specific outcomes to process structure and alternatives. Most previous work in constraint-based scheduling has focused on systems requiring synchronization of shared resources among large numbers of well-defined production processes (e.g., sets of airlift missions to accomplish various cargo movement requirements; sets of process routings to produce various requested manufactured parts). Models of individual processes have tended to require fairly simple process structures (e.g., linear sequences of tasks).

The ability to respond to outcome-based task status updates (e.g., task succeeded, task failed, task was 75% successful), first of all requires richer descriptions of process structure (e.g., incorporating conditional and looping control structures, process and sub-process alternatives). Modern planning formalisms such as hierarchical task networks [16] provide the sorts of general-purpose representational frameworks that are necessary here; however these representations are not well-reconciled with the resource models that are required for making resource allocation and

scheduling decisions. In [13], a hierarchical representation of activity types is linked to resource alternatives through the specification of *capabilities*, and this representation is used to dynamically manipulate process structures (i.e., plan fragments) as various resource assignments are considered (and re-considered). Some steps toward the development and use of an integrated representation have also been taken in [9].

#### *Reasoning About Task Quality and Cost*

Given a representational framework that allows incorporation of information relating to task content and outcomes in process descriptions, a second set of issues concerns the nature of the parameters (or metrics) used to characterize this information, and how this information might be used to drive an incremental scheduling process. Certainly binary indicators of task success and failure need to be accommodated as possible task status updates and this can serve as a basic model for directing process change. However, in the context of KIDS, more discriminating (e.g., partial) models are fundamental to effective process management and scheduling. Of particular relevance and interest are scheduling models that account for the *quality* of the results obtained as a function of the *cost* of task execution.

From the standpoint of maximizing task quality, we can identify at least two potentially useful types of derivative interpretations of the source of task quality: (1) task quality is an exogenous property (e.g., one process for obtaining information is inherently more reliable than another, or the quality of an information collection task increases the longer it executes), and (2) task quality is a function of the skills (expertise) or quantity of assigned resources (e.g., an experienced reporter will get better information than a novice). The former has been the classic focus of research in deliberation scheduling [5]. However, this work has tended to focus on allocating independent sets of tasks, subject to individual task deadlines. In the case of KIDS, alternatively, inter-dependencies between tasks (processes) are more typical, and task quality can be influenced (positively or negatively) by task interactions. Scheduling models that consider the impact on quality of the skill level or quantity of resources assigned have received very little attention as well, since the context of deliberation scheduling is generally a single processor (the problem solving agent). Here one general issue is producing (and maintaining) a schedule that achieves the "best fit" of available resources to input tasks.

#### *Hedging Against the Possibility of Task Failure.*

At one extreme end of the task quality spectrum is the possibility of complete task failure. In circumstances of certain high priority requirements, it may make sense to attempt to maximize the likelihood of success by introducing some amount of process redundancy in the coverage (e.g., allocating several, possibly different processes to a particular information gathering objective). Likewise, it might be useful to budget extra time to increase the likelihood of a desired outcome. There has been

relatively little work in the area of building such forms of robustness into production schedules.

#### *Improving the Collective Capability of the KIDS*

Given that task quality is generally a function of the expertise of resources that are assigned to perform the task, one organizational objective of a KIDS is to improve its collective skill set over time. On one hand, this objective can lead to performing tasks better and more efficiently in the future; on the other hand it can also minimize the potential deleterious effects of future resource losses (e.g., as resources leave the organization). Generally speaking, improving the skill set or level of a given resource is a function of gaining experience and practice at executing various tasks. Accordingly, task allocation and scheduling procedures that pro-actively consider the side effects of acquired expertise and seek to enrich the capabilities of individual resources wherever performance constraints allow are clearly desirable.

#### *Distribution of Scheduling and Control Responsibility*

Although techniques for incremental management of schedules provide some measure of scalability in their own right, a second key aspect of scalable coordination of KIDS is distributed decision-making. The scale and time-stress associated with effectively keeping pace with execution in the context of KIDS make distribution of scheduling and control responsibility a practical necessity. In [12], a framework and methodology for decomposition and distribution of scheduling activity is proposed in the context of large-scale, work-flow management processes. This framework presumes a hierarchical description of processes and resources, mapped to various levels of decision-making and authority within the organization. Some reliance on such a framework appears necessary to ensure effective global coordination of complex inter-dependent processes. The downside, however, is the potential for computational bottlenecks, if models at successive levels are too tightly linked.

At the other end of the spectrum, recent research in agent-based approaches has given rise to so-called "self-scheduling" systems [4,15]. Such systems operate in a bottom-up, totally decentralized fashion. They do not compute schedules in advance but instead make allocation decisions as needed via local interaction between individual executing agents. Self-scheduling systems provide scalability, but it is generally difficult to predict global performance characteristics.

Given the tradeoffs between globally focused and self-scheduling approaches to decentralization, the key to effective scalable management of KIDS seems to lie in finding middle ground that capitalizes on their respective strengths and minimizes their respective weaknesses. For example, self-scheduling systems have proven most effective in fairly structured situations (e.g., where a single performance objective is sufficient) or in resource rich execution environments. Both of these types of circumstances are most likely to be present at lower levels

in an organizational hierarchy, where managed processes are narrow in scope and decision-time frames are short. By the same token, these types of circumstances are fairly unlikely at more global decision-making levels, and here explicit advance scheduling is more likely to pay performance dividends. One key here will be the ability to develop summary performance profiles of localized methods that enable establishment of reasonable operating constraints at higher, circumscribing decision-making levels.

#### 4. SUMMARY

In this paper, we have advocated the use of incremental, constraint-based scheduling techniques as an effective basis for coordinating activity in knowledge-intensive dynamic systems. Constraint-based scheduling models are well matched to the requirements of KIDS: they promote controlled continuous change as execution circumstances evolve and warrant it, they operate under quite general representational assumptions, they integrate well with user-decision making, and they directly support negotiation and conflict resolution in distributed settings. At the same time, the characteristics of KIDS pose new research challenges for constraint-based scheduling research. Extensions are required to allow richer models of process structure, to support reasoning about task quality as a function of execution cost and resources assigned, to make allocation decisions that promote enhancement of organizational expertise and to ensure scalability of continuous scheduling processes. Our current research is focused on this set of issues.

#### ACKNOWLEDGEMENTS

The work reported in this paper has been supported in part by the Department of Defense Advanced Research Projects Agency and the U.S. Air Force Research Laboratory – Rome, under contracts F30602-97-2-0666 and F30602-00-2-0503, by the National Science Foundation under contract # 9900298 and by the CMU Robotics Institute. Thanks to David Hildum for helpful comments on an earlier draft.

#### REFERENCES

- [1] Baptiste, P., C. LePape and W. Nuijten, *Constraint-Based Scheduling: Applying Constraint Programming to Scheduling Problems*, Kluwer Academic Press, Boston, 2001
- [2] Becker, M. and S.F. Smith, "Mixed-Initiative Resource Management: The AMC Barrel Allocator", Proceedings 5th Int. Conference on Artificial Intelligence Planning and Scheduling (AIPS-2000), Breckenridge, CO, April, 2000.
- [3] Cheng, C. and S.F. Smith, "Applying Constraint Satisfaction Techniques to Job Shop Scheduling", *Annals of Operations Research*, Vol. 70, 1997, pp. 327-357.
- [4] Cicirello, V.A. and S.F. Smith, "Wasp Nests for Self-Configurable Factories", Agents '01, Proceedings 5th International Conference on Autonomous Agents, Montreal Canada, May 2001
- [5] Dean, T. and M. Boddy, "Deliberation Scheduling for Problem Solving in Time-Constrained Environments", *Artificial Intelligence*, 67(2): 245-285, 1994
- [6] Derthick, M. and S.F. Smith, "An Interactive 3D Visualization for Requirements Analysis" submitted for Publication.
- [7] Kramer, L.A. and S.F. Smith, Optimizing for Change: Mixed-Initiative Resource Allocation with the AMC Barrel Allocator, *Proc. 3<sup>rd</sup> International NASA Workshop on Planning and Scheduling for Space*, November 2002.
- [8] Kramer, L.A. and S.F. Smith, "Maximizing Flexibility: A Retraction Heuristic for Over-subscribed Scheduling Problems", *Proceedings 18<sup>th</sup> International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, Aug 2003.
- [9] Rubinstein, Z.B., "Efficient Scheduling of Evolving Nondeterministic Process Plans in Dynamic Environments", Ph.D. Thesis, Computer Science Department, University of Massachusetts, September 2002.
- [10] Smith, S.F., "OPIS: A Methodology and Architecture for Reactive Scheduling", in *Intelligent Scheduling*, (eds. M. Zweben and M. Fox) Morgan Kaufmann Pub., 1994.
- [11] S.F. Smith, O. Lassila and M. Becker, Configurable, Mixed-Initiative Systems for Planning and Scheduling, in *Advanced Planning Technology*, (ed. A. Tate), AAAI Press, Menlo Park, CA, 1996.
- [12] Smith, S.F., D.W. Hildum and M.A. Becker, "Workflow Management from a Scheduling Perspective", *Proc. 1999 AAAI Workshop on Agent-Based Systems in a Business Context*, Orlando, FL, July, 1999.
- [13] Smith, S.F., D.W. Hildum and D. Crimm, "Interactive Resource Management in the Comirem Planner", *Proceedings IJCAI-03 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco Mexico, August 2003.
- [14] Smith, S.F., M.A. Becker, and L.A. Kramer, "Continuous Management of Airlift and Tanker Resources: A Constraint-based Approach", *Journal of Mathematical and Computer Modeling*, Special Issue on Defense Transportation: Algorithms, Models and Applications for the 21st Century, In Press.
- [15] Smith, S.F., "Is Scheduling a Solved Problem?", *Proc. 1<sup>st</sup> Multi-Disciplinary International Conference on Scheduling Theory and Applications (MISTA)*, Nottingham, UK, August 2003.
- [16] Wilkins, D. *Practical Planning*, Morgan Kaufmann Publishers, San Mateo CA, 1988.
- [17] Zhou, Q. and S.F. Smith, "A Priority-Based Preemption Algorithm for Incremental Scheduling with Cumulative Resources", CMU Robotics Institute Technical Report CMU-RI-TR-02-19, June 2002
- [18] Zweben, M., B. Daun, E. Davis and M. Deale, "Scheduling and Rescheduling with Iterative Repair, in *Intelligent Scheduling*, (eds. Zweben & Fox), Morgan Kaufmann Publishers, 1994.