

57-63

137233

93-18665

Adaptive Planning For Applications With Dynamic Objectives

Khosrow Hadavi
Wen-Ling Hsu
Siemens Corporate Research

Michael Pinedo
Columbia University

Abstract

Planning is commonly viewed as a task to devise a course of action or a *plan* that conforms as much as possible to a set of goals before acting. The plan will then be used to guide the activities. Most classic planning systems assume a static environment for the planning agents. In a static environment, states remain unchanged between actions, and the outcomes of actions are assumed to be deterministic. In reality, however, most applications are dynamic and stochastic in nature. External events, not caused by controlled actions, may occur; outcomes of actions may differ from expectations; new constraints may be introduced; and a new set of goals may evolve in response to the changes. Recently, we have proposed a multi-modal framework for adaptive planning in a *dynamic* environment with multiple objectives having the following characteristics:

- some of the objectives of the planning process may be conflicting
- some objectives may be ill-defined or difficult to measure quantitatively
- the objectives may change over time

The task domain of production planning and scheduling is a typical example of such an environment. The scheduling objectives typically include the following: meeting due dates; reducing lead times; reducing work-in-process and finished goods inventories; maximizing resource utilization and the throughput of the system; and minimizing the sensitivity of the schedule to random events. These objectives are sometimes in conflict with each other. In our previous work, we developed a real-time distributed scheduling system¹ that observes its environment from different perspectives. These perspectives stem from the different objectives, and the system can react to events as they occur while monitoring the various objectives. This multi-perspective monitoring helps our system achieve better control of the environment. During our study, we discovered that although these global objectives may not change over time, the relevance of each objective is actually a function of time and the state of the system. For example, given a set of N objectives O_1, O_2, \dots, O_n , at time t_1 , objective O_2 may be significantly more important than O_1 , whereas at another instance

¹For a detailed description of the system, please read the attached paper titled "An Architecture for Real Time Distributed Scheduling" to appear in "Applications of AI in Manufacturing," published by AAAI Press, edited by Dana S. Nau.

of time t_2 , objective O_1 may become most important. Furthermore, each heuristic implies a set of reactive strategies that move the system toward some objectives but *away* from other objectives (due to the conflicting nature of these objectives).

In our current research, we devise a qualitative control layer to be integrated into a real-time multi-agent reactive planner. The reactive planning system consists of distributed planning agents attending to various perspectives of the task environment. Each perspective corresponds to an objective. The set of objectives considered are sometimes in conflict with each other. Each agent receives information about events as they occur, and a set of actions based on heuristics can be taken by the agents. Within the qualitative control scheme, we use a set of *qualitative feature vectors* to describe the effects of applying actions. A *qualitative transition vector* is used to denote the qualitative distance between the current state and the target state. Given a target state and a set of heuristics, we have an algorithm to test the *reachability* of the target state. We will then apply on-line learning at the qualitative control level to achieve adaptive planning. Our goal is to design a mechanism to refine the heuristics used by the reactive planner every time an action is taken toward achieving the objectives, using feedback from the results of the actions. When the outcome is compared with expectations, our prior objectives may be modified and a new set of objectives (or a new assessment of the relative importance of the different objectives) can be introduced. Because we are able to obtain better estimates of the time-varying objectives, the reactive strategies can be improved and better prediction can be achieved.

END