Time-Critical Scheduling in Stochastic Domains

Lloyd Greenwald **Thomas Dean**
Department of Computer Science
Brown University, Box 1910, Providence, RI 02912
lgg@cs.brown.edu tld@cs.brown.edu

In this work we look at extending the work of (Dean *et al.* 1993) to handle more complicated scheduling problems in which the sources of complexity stem not only from large state spaces but from large action spaces as well. In these problems it is no longer tractable to compute optimal policies for restricted state spaces via policy iteration. We, instead, borrow from operations research in applying bottleneck-centered scheduling heuristics (Adams *et al.* 1988). Additionally, our techniques draw from the work of (Drummond and Bresina 1990).

Consider the problem of scheduling planes and gates at a busy airport. A stochastic process describes the arrival of planes at the airport and is affected by uncontrollable events such as weather. Stochastic processes also govern the processing requirements for unloading and loading passengers at arrival and departure gates. Deadlines for theses operations are determined by prespecified desired arrival and departure times. Other sources of uncertainty include gate closings. The optimization problem is to assign planes to gates at each time step to minimize some global measure of tardiness. In any given state there is, in general, one action for every possible assignment of planes to gates.

The work of (Dean *et al.* 1993) introduces a general approach to planning and scheduling in stochastic domains in which a two-phase iterative procedure is employed. The first phase determines a restricted subset of the state space on which to focus (called the *envelope*) and the second phase constructs a policy for this envelope. Deliberation scheduling is employed to allocate on-line computation time between the anytime algorithms that make up each phase and across iterations of both phases. This approach directly addresses uncertainty by modeling the environment as a stochastic automaton and constructing policies to account for alternative trajectories reachable from a given start state. Restricting policy construction to a given envelope addresses the large state space issue.

Our work addresses the additional combinatorial explosion of large action spaces by focusing processing on *time windows* rather than envelopes of specific states, and by selectively exploring the space described by the window rather than exhaustively exploring the space via policy iteration. While planning domains such as robot navigation may adhere to a restricted neighborhood of states over time, states solved for prior time steps in scheduling domains with large action spaces do not remain relevant as the process progresses. Additionally, alternative actions from any given state lead to disjoint state spaces with little chance that the trajectories will merge on a common envelope of specific states. By partitioning the state space along the time dimension we capture the appropriate context in scheduling domains.

We generate policies by selectively exploring the state space described by the time window. For any given state we use dispatch scheduling rules such as earliest deadline first to select an action. We then employ Monte Carlo simulation on a stochastic model of the domain to determine the most probable reachable states. This process is repeated for a fixed amount of time to determine a partial policy. A second phase attempts to improve the expected value of the policy by detecting bottlenecks and constraining associated actions in further iterations of policy generation. By using greedy dispatch rules on unconstrained actions, we avoid exhaustively searching large action spaces. This procedure is augmented with default reflexes for low probability states not explicitly simulated. We employ deliberation scheduling to allocate on-line processing time across time windows and phases based on anticipated quality-time tradeoffs.

References

Adams, J.; Balas, E.; and Zawack, D. 1988. The shifting bottleneck procedure for job shop scheduling. *Management Science* 34(3):391–401.

Dean, Thomas; Kaelbling, Leslie; Kirman, Jak; and Nicholson, Ann 1993. Planning with deadlines in stochastic domains. In *Proceedings AAAI-93*. AAAI. 574–579.

Drummond, Mark and Bresina, John 1990. Anytime synthetic projection: Maximizing the probability of goal satisfaction. In *Proceedings AAAI-90*. AAAI. 138–144.