

# Preface

Automated systems for planning and scheduling the activities of individuals, and of coordinated groups, must scale into increasingly complex, dynamic, stochastic, and even adversarial application domains where plan/schedule failure can be costly or catastrophic. Substantial research continues to be directed towards techniques for robust/probabilistic planning and scheduling, monitoring and repairing plans and schedules, continual distributed planning, etc. This work typically assumes that agents' objectives stay relatively stable over time, but plans and schedules need to adjust to deviations from predicted execution paths to account for violated assumptions and to exploit emergent opportunities.

There is also growing interest in applications where agents' objectives (tasks, intentions, preferences...) can evolve over time. Even when a plan/schedule is progressing exactly as expected, the agents might reprioritize their tasks or alter how they intend to pursue their objectives. Such plans and schedules need to be repaired to reflect the agents' revised intentions.

For example, a planning/scheduling management process could be associated with a human user who, over time, adopts new tasks, reprioritizes existing goals as dependencies among them are discovered, reconsiders the effectiveness of alternative ways of pursuing an objective as more is learned about the outcomes of previously planned/scheduled activities, and so on. The process should ensure that these changes are addressed as well as possible in the revised plan/schedule. If the user is acting in an unpredictable world, this compounds the challenges faced by the planning/scheduling process as it must contend with adapting to the environmental trajectory as well. Further, if the user is interacting with other users, then the planning/scheduling process must account for and coordinate the ways in which the evolving plans/schedules of others affect it, and how the changes it makes to its plan/schedule affect them. And all of this asynchronous plan and schedule management must be done quickly enough to satisfy the demands of goals and tasks that must be achieved in a timely manner.

This symposium brings together researchers who are tackling various aspects of the distributed plan and schedule management problem to foster a broader understanding of this problem, establish a fruitful dialogue between researchers focusing on insular pieces of the problem, and instigate an exchange of ideas, techniques, and best practices to accelerate progress. Issues to be discussed at the symposium include:

How is dealing with instability in preferences and objectives similar to, and different from, dealing with environmental unpredictability during plan/schedule execution?

How can efficient techniques for robust conditional/probabilistic planning, for plan monitoring and repair, for continual plan refinement and replanning, etc., be gainfully adopted for distributed plan/schedule management?

How can distributed plans/schedules be configured to avoid, correct for, exploit, and/or tolerate asynchronously-arising local (environment or objective) deviations, and how can patterns of local deviations be distributively diagnosed and responded to?

How must distributed planning and scheduling algorithms change to adapt to open world situations in which the timing of events or activities are externally imposed and must be worked around?

What are the implications of bringing together distributed heterogeneous agents, whose planning and scheduling concerns are at very different temporal grain sizes or planning horizons, whose knowledge and beliefs could be incomplete or inconsistent with each other, or whose preferences and objectives might not be entirely aligned?

How should agents make meta-level decisions about how frequently and thoroughly changes within, and across, their plans/schedules should be carried out?

What role can learning play in improving distributed plan/schedule management efficiency and/or effectiveness?

To what extent should distributed plan/schedule management provide feedback about changes to the environment and to other agents' plans/schedules to inform decisions about how to prioritize tasks and form intentions, and in what form should the feedback be presented?

What principled/formal models of the stochastic environment, the agents' evolving individual/joint preferences and intentions, the mechanisms by which they can decide on plans/schedules, etc. are germane to defining and solving the distributed plan/schedule management problem?

What application experiences exist in designing, developing, and deploying distributed plan/schedule management systems, and what can we learn from them?

We would like to thank the authors of the papers in this volume for their contributions and their participation in this symposium. We would also like to thank the auxiliary reviewers for providing their valuable feedback. Finally, we appreciate the efforts of Alan Schultz and of the staff at AAI for their assistance.

*Ed Durfee, Dave Musliner, cochairs  
Marie desJardins, Bob Neches,  
Steve Smith, Regis Vincent,  
and Tom Wagner*