08461 Abstracts Collection
Planning in Multiagent Systems
— Dagstuhl Seminar —

Jürgen Dix[1], Edmund Durfee[2] and Cees Witteveen[3]

[1] TU Clausthal, Germany
dix@tu-clausthal.de
[2] University of Michigan, USA
durfee@umich.edu
[3] Delft University of Technology, The Netherlands
C.Witteveen@tudelft.nl

Abstract. From the 9th of November to the 14th of November 2008 the Dagstuhl Seminar 08461 "Planning in Multiagent Systems" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Multi-agent systems, AI-planning, coordination, robustness, temporal planning

08461 Executive Summary – Planning in Multiagent Systems

Planning in Multiagent Systems, or Multiagent Planning (MAP for short), considers the planning problem in the context of multiagent systems. It extends traditional AI planning to domains where multiple agents are involved in a plan and need to act together.

Research in multiagent planning is promising for real-world problems: on one hand, AI planning techniques provide powerful tools for solving problems in single agent settings; on the other hand, multiagent systems, which have made significant progress over the past few years, are recognized as a key technology for tackling complex problems in realistic application domains.

The motivation for this seminar is thus to bring together researchers working on these different fields in AI planning and multiagent systems to discuss the central topics mentioned above, to identify potential opportunities for coordination, and to develop benchmarks for future research in multiagent planning.

Dagstuhl Seminar Proceedings 08461
Planning in Multiagent Systems
http://drops.dagstuhl.de/opus/volltexte/2009/1874

Keywords: Multi-agent systems, AI-planning, coordination, robustness, temporal planning

Joint work of: Dix, Jürgen; Durfee, Edmund H.; Witteveen, Cees

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2009/1873

The Multi-Agent Contest Competition

Tristan Behrens (TU Clausthal, Germany)

The Agent Contest is an international competition that has been created in 2005 in order to allow the comparison of agent-based approaches to systems programming. Benchmarks are provided by letting agent teams solve a cooperative task in a dynamically changing environment, while they have to compete against other teams.

The first Agent Contest was held in 2005 in association with the CLIMA workshop. The scenario was a grid-like world, in which agents had to gather and store resources facing incomplete-information. The first contest was decentralized: the participants had to implement the agents as well as the environment and did not compete with other teams. In 2006 this has been changed by introducing the MASSim platform, which has been the fundament of the contest since then. We kept the scenario and let agent-teams compete for gold. In 2007 we moved to the ProMAS workshop and kept the scenario. In 2008 we have changed the scenario to the cows and cowboys scenario that was designed in order to put stress on the cooperation and coordination aspects of agent programming.

Keywords: Multi-agent systems programming, benchmarks, competition

Joint work of: Behrens, Tristan; Dastani, Mehdi; Dix, Jürgen; Novak, Peter

Multi-Agent Programming (MAP Without Planning, Unless You Want)

Rafael Bordini (University of Durham, UK)

In this talk, I gave an overview of the area of programming languages for Multi-Agent Systems and of Jason (a multi-agent platform based on a logic-based language for programming BDI agents). I mentioned a number of research projects aimed at extending Jason or combining it with other approaches, as well as a project aimed at formal verification of multi-agent systems using model-checking techniques. I also discussed some work that appeared recently in the Agents literature proposing the combination of a BDI programming language with a planner, then made some concluding remarks.

Keywords: Multi-Agent Programming, Jason, AgentSpeak, Model Checking

Where can Planning be used in the modeling of rational agents?

Nils Bulling (TU Clausthal, Germany)

Alternating-time temporal logic (ATL) is a temporal logic that can be used to model and to reason about multi-agent systems. The logic incorporates some basic game theoretic notions and has already been extended by additional concepts by several researchers, including ourselves.

In this talk, I will give an overview of previous work about extensions of ATL which focus on rational behavior in multi-agent systems. Then, we try to point out where planning might be used in combination with these logics, especially if reasoning within agents is considered.

Keywords: Logics, Multi-agent systems, Rationality, Planning

Joint work of: Bulling, Nils; Dix, Jürgen; Jamroga, Wojciech

A MultiAgent Systems Perspective on the MultiAgent Planning

Edmund H. Durfee (Univ. of Michigan - Ann Arbor, USA)

Developers of multiagent systems often confront problems in multiagent planning as a means, as well as sometimes an end, for coordinating agents. In this talk, I take a few arguably extreme positions in order to encourage discussion and introspection. Launching from the conjecture in the MAS community that "all agent systems are multiagent systems" I extend this to say that "all multiagent systems are multiagent planning systems." By the latter, I mean that solving the problem of how agents should coordinate their activities requires solution of a multiagent planning problem. Sometimes this problem is solved offline, and even by an agent external to the system in question; for example, a system developer might design a protocol (which arguably is a partially-instantiated multiagent plan), or an organizational design (which arguably is a multiagent plan at a level of abstraction where "roles" replace operators). At other times, agents explicitly model their intended problem-solving activities, and identify potential future interactions to seek or avoid.

After describing how multiagent planning is at least implicitly required as a means for coordination in multiagent systems, I then turn to how multiagent systems can also be used for producing multiagent plans as an end in itself. I discuss motivations for doing so, including exploiting parallel computation (e.g., plan merging), utilizing distributed expertise, avoiding expensive centralization, and handling issues such as privacy, autonomy, authority, etc. I conclude with various open questions for the field, for defining objectives, constraints, and characteristics for multiagent planning.

Challenges panel introduction

Robert P. Goldman (SIFT - Minneapolis, USA)

Introduction to the panel on Challenges for Multi-Agent Planning, containing framing material concerning multi-agent planning and a discussion of three specific challenges. The framing material contains a rudimentary taxonomy of multiagent planning problems (together with some odd cases in the space), motivations for doing multi-agent planning, and techniques and challenges (less clearly separated than should be). The three specific challenges addressed were planning communication actions (in connection with the author's work on the multiagent CIRCA system), incentivizing self-interested agents to act as part of a team (in connection with the author's work on the DARPA COORDINATORS program), and finally a discussion of the possibility for shared representations and problem sets for multi-agent planning.

Distributed Task Allocation in Social Networks

Tomas B. Klos (Delft University of Technology, The Netherlands)

This paper proposes a new variant of the task allocation problem, where the agents are connected in a social network and tasks arrive at the agents distributed over the network.

We show that the complexity of this problem remains NP-hard. Moreover, it is not approximable within some factor. We develop an algorithm based on the contract-net protocol. Our algorithm is completely distributed, and it assumes that agents have only local knowledge about tasks and resources.

We conduct a set of experiments to evaluate the performance and scalability of the proposed algorithm in terms of solution quality and computation time. Three different types of networks, namely small-world, random and scale-free networks, are used to represent various social relationships among agents in realistic applications. The results demonstrate that our algorithm works well and that it scales well to large-scale applications.

Keywords: Task allocation, social networks, multiple agents, distributed protocol

Joint work of: Weerdt, Mathijs de; Zhang, Yingqian; Klos, Tomas

Full Paper:
http://ifaamas.org/Proceedings/aamas07/html/AAMAS07_0389_76a1ec696d0d8aafbab6fd454c3ec490.xml

See also: Proceedings AAMAS 2007

Planning for Interactions with Multiple Autonomous Agents

Dana S. Nau (University of Maryland - College Park, USA)

There are several kinds of planning algorithms for environments in which there is a single agent whose actions may have multiple outcomes. I'll discuss ways to use these planning algorithms to solve multi-agent planning problems. The basic idea is to develop models of the agents, use these models to translate the multi-agent problem into one or more single-agent planning problems, and solve those problems using a single-agent planner.

I'll discuss two kinds of agent models: capability models, and probabilistic forecasting models. I'll explain how these enable us to use algorithms for planning with nondeterminism and planning over MDPs, respectively. I'll also discuss whether to do the planning online or offline, and some similarities and differences to non-zero-sum games. I'll include experimental results from several multi-agent planning domains.

Keywords: Planning with nondeterminism, Planning on MDPs, Multi-agent systems, Agent models, Game theory

Coordinating Plans Through Distributed Constraint Optimization

Brammert Ottens (EPFL - Lausanne, Switzerland)

In this paper we show how the coordination of agent plans can be performed using Distributed Constraint Optimisation (DCOP) techniques. In particular, we show how a Truck Task Coordination problem can be modelled as a DCOP. We introduce a complete asynchronous DCOP algorithm, Asynchronous Open DPOP (ASODPOP), based on the DPOP algorithm that exhibits fast convergence to the optimal solution compared with both ADOPT and Distributed Stochastic Search (DSA). Fast convergence is useful when agents are time bounded and are thus unable to wait for an optimal solution.

Keywords: DCOP, Logistics, Planning, Coordination

Full Paper:
http://liawww.epfl.ch/People/ottens/papers/Ottens2008d.pdf

See also: Proceedings of the Multi Agent Planning Workshop - ICAPS 2008, September 2008, Sydney Australia

ASODPOP: Making Open DPOP Asynchronous

Brammert Ottens (EPFL - Lausanne, Switzerland)

In this paper we show how ODPOP can be adapted to an asynchronous environment where agents might have to decide their values before the algorithm has ended, giving us Asynchronous ODPOP (ASODPOP). We have compared the algorithm with both ADOPT and distributed local search (DSA). Compared to ADOPT we show that our approach sends fewer messages, converges to a reasonable solution faster, and uses an equal amount of NCCCs. We also show that this convergence is much faster than local search, whilst the solution that local search converges to is far from optimal.

Keywords: DCOP, Logistics, Planning, Coordination

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1871

Using options with set exercise prices to reduce bidder exposure in sequential auctions

Valentin Robu (CWI - Amsterdam, The Netherlands)

The exposure problem appears whenever an agent with complementary valuations bids to acquire a bundle of items sold sequentially, in separate auctions. In this talk, we review a possible solution that can help solve this problem, which involves selling options for the items, instead of the items themselves. We provide a brief overview of the state of the art in this field and discuss, based on our recent results, under which conditions using option mechanisms would be desirable for both buyers and sellers, by comparison to direct auctioning of items. We conclude with a brief discussion of further research directions in this field, as well as the relation to other techniques proposed to address the problem, such as leveled commitment mechanisms.

Keywords: Options, sequential auctions, multi-agent systems, exposure problem, bidding strategies, mechanism design, leveled commitment

Joint work of: Mous, Lonneke; Robu, Valentin; La Poutre, Han

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1872

Coordination of Distributed Planning and Scheduling Agents

Stephen Smith (Carnegie Mellon University - Pittsburgh, USA)

The practical constraints of many application environments require distributed management of executing plans and schedules.

Such factors as geographical separation of executing agents, limitations on communication bandwidth, constraints relating to chain of command and the high tempo of execution dynamics may all preclude any single agent from obtaining a complete global view of the problem, and hence necessitate collaborative localized planning and scheduling decisions. For the past 3 years, as part of the DARPA Coordinators program, my group at CMU has been pursing the development of scalable frameworks for collaborative distributed schedule management in uncertain execution environments. This work has led to the development of the cMatrix agent architecture, which implements a philosophy of rapid localized response to unexpected circumstances, subsequent propagation of consequences to other agents with inter-dependent decisions, and then, as time permits, coordinated negotiation with these agents to improve on local solutions through joint change. In early evaluation tests with a simulator playing the uncertain environment, the cMatrix system was shown to produce execution results within .02% of those produced by an expected optimal (but non-scalable) centralized MDP solver on small problem instances, and to effectively scale to problems involving up to 100 agents and 10,000 tasks. In August 2008, the system competed in a field test exercise, where the objective was to direct a team of human agents in carrying out a coordinated response to a mock natural disaster. In this talk, I will discuss the challenges presented by this class of problem, summarize our technical approach and the results obtained to date, and outline our current research directions.

Keywords: Multi-agent planning and scheduling

Joint work of: Smith, Stephen

Full Paper:
http://www.ri.cmu.edu/publication_view.html?pub_id=5707

See also: Smith, S.F., A. Gallagher, T. Zimmerman, L. Barbulescu and Z. Rubinstein, Distributed Management of Flexible Times Schedules, Proceedings 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 07), Honolulu Hawaii, May 2007.

Multi-Agent Systems for the Real World

Pedro Szekely (USC/ISI - Marina del Rey, USA)

Creating multi-agent applications for human users in the real world requires approaches that can cope with a wide variety of unpredictable events. Users may make errors or improvise, communication bandwidth and latency is highly variable, activities may fail, and the models that systems use may be unexpectedly inaccurate or plain wrong. We present a formal framework to understand the different types of issues that systems must address, and argue that the particular algorithms used to make decisions are far less important than allowing users to provide strategic guidance providing agents with good information upon which

to base those decisions. We show how unrealistic assumptions about information quality foreordain poor real-world performance for several popular multi-agent approaches. We propose our PCM approach as a first step towards simultaneously addressing the full range of interacting issues raised by trying to make multi-agent systems effective in the real-world.

Keywords: Multi-Agent Systems, Planning and Scheduling, Real World, Coordination

Execution Monitoring of Human-Machine Teams

David E. Wilkins (SRI - Menlo Park, USA)

There is an increasing need for automated support for humans monitoring the activity of distributed teams of cooperating agents, both human and machine. We characterize the domain-independent challenges posed by this problem, and describe how properties of domains influence the challenges and their solutions. We will concentrate on dynamic, data-rich domains where humans are ultimately responsible for team behavior. Thus, the automated aid should interactively support effective and timely decision making by the human. We present a domain-independent categorization of the types of alerts a plan-based monitoring system might issue to a user, where each type generally requires different monitoring techniques. We describe a monitoring framework for integrating many domain-specific and task-specific monitoring techniques and then using the concept of "value of an alert" to avoid operator overload.

We use this framework to describe an execution monitoring approach we have used to implement Execution Assistants (EAs) in three different dynamic, data-rich, real-world domains to assist a human in monitoring team behavior. One domain (Army small unit operations) has hundreds of mobile, geographically distributed agents, a combination of humans, robots, and vehicles. Another domain (teams of unmanned ground and air vehicles) has a handful of cooperating autonomous robots. Our EAs alert the human controller when reported events threaten plan execution or physically threaten team members. Alerts were generated in a timely manner without inundating the user with too many alerts (less than 10% of alerts are unwanted, as judged by domain experts).

Keywords: Execution, monitoring, human-robot teams, plans

Full Paper:
http://www.ai.sri.com/~wilkins/papers/ex-mon-jair.pdf

See also: "Airlift mission monitoring and dynamic rescheduling", Engineering Applications of Artificial Intelligence Journal, March, 2008, volume 21. Interactive Execution Monitoring of Agent Teams, Journal of Artificial Intelligence Research, volume 18, pages 217-261, March 2003

Creating incentives to prevent execution failures

Yingqian Zhang (Delft University of Technology, The Netherlands)

When information or control in a multiagent planning system is private to the agents, they may misreport this information or refuse to execute an agreed outcome, in order to change the resulting end state of such a system to their benefit.

In some domains this may result in an execution failure. We show that in such settings VCG mechanisms lose truthfulness, and that the utility of truthful agents can become negative when using VCG payments (i.e., VCG is not strongly individually rational). To deal with this problem, we introduce an extended payment structure which takes into account the actual execution of the promised outcome. We show that this extended mechanism can guarantee a nonnegative utility and is (i) incentive compatible in a Nash equilibrium, and (ii) incentive compatible in dominant strategies if and only if all agents can be verified during execution.

Keywords: Mechanism design, multiagent planning

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1870

Mechanism Design for Multiagent Planning: Difficulties of applying VCG to MAP

Mathijs de Weerdt (Delft University of Technology, The Netherlands)

Multiagent planning methods are concerned with planning by and for a group of agents. If the agents are self-interested, they may be tempted to lie in order to obtain an outcome that is more rewarding for them. We therefore study the multiagent planning problem from a mechanism design perspective, showing how to incentivise agents to be truthful. We prove that the well-known truthful VCG mechanism is not always truthful in the context of optimal planning, and present a modification to fix this. Finally, we present some (domain-dependent) polytime planning algorithms using this fix that maintain truthfulness in spite of their non-optimality.

Keywords: Multiagent planning, mechanism design, VCG

Joint work of: de Weerdt, Mathijs; van der Krogt, Roman; Zhang, Yingqian

Full Paper:
http://www.st.ewi.tudelft.nl/%7Emathijs/publications/ecai08.pdf

See also: Roman P.J. van der Krogt and Mathijs M. de Weerdt and Yingqian Zhang. Of Mechanism Design and Multiagent Planning, in Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08), IOS Press, 2008, Malik Ghallab and Constantine D. Spyropoulos and Nikos Fakotakis and Nikos Avouris, 423–427.