

Plan-based Control of Robotic Agents

21.10.2001-26.10.2001

organized by

M. Beetz (Munich University of Technology), J. Hertzberg (FhG AIS, Bonn), M. Ghallab (LAAS - CNRS, Toulouse), M. Pollack (University of Michigan, Ann Arbor)

The Dagstuhl seminar Plan-based Control of Robotic Agents was held on October 21-26, 2001 and organized by Michael Beetz, Joachim Hertzberg, Malik Ghallab, and Martha Pollack.

In the Seminar, we brought together a team of 34 leading researchers who are researching different aspects of plan-based control such as flexible and reliable execution of plans, execution time plan formation and revision, and automatic learning of robot plans.¹ They discussed issues in plan-based control and worked towards a comprehensive framework for plan-based control of robotic agents. In the representations and discussions we were particularly addressing the following issues, both from the methodological and the application side.

- **Flexible and reliable execution of plan-based control.** What is the right representation and expressivity of plans for autonomous robot control? Should the plan language support flexible and reliable execution? What kind of plan management operations other than plan formation should the plan representation facilitate? How can the plan representation be grounded into the sensory and actuation capabilities of the robot?
- **Realization of runtime plan management and the integration of plan management into the overall control.** How can we accomplish the feasibility of plan management operations? Should plan management be incorporated at a plan layer of a hybrid robot control architecture? Are there other means for integrating plan management into the overall control?
- **Automatic learning of plan schemata and plan revision knowledge.** How can the robot automatically chunk its continuous behavior and learn complex behavior structures? How can the robot automatically

¹Despite having a great attendance of top researchers (even from the US), 12-14 top researchers from the US had to cancel their participation on short notice due to the terror attacks on September 11.

adapt to specific environments and tasks? How can the robot learn or acquire the plan schemata and planning knowledge needed for execution time plan management?

- **Formal models of plan-based control, plan formation, and plan revision.** How can we obtain realistic formal models of plan-based control? Can formal languages be used to define an abstract description layer for robot controller design? Do formalizations help clarify robot/human and robot/robot interaction, high-level task-oriented command, and telemanipulation interface languages?
- **Applications of plan-based robot controllers.** What are the challenge applications for plan-based control of robotic agents? Are there empirical results on the use of plan-based versus behavior-based robot controllers? Should we try to collect benchmarks for plan-based robot control?

The seminar covered the topics listed above by a program of long talks, technical talks and panel discussions. Details of the program are still available under <http://www.dagstuhl.de/DATA/Reports/01431/>. Selected contributions to the Seminar are currently under review and will be published as a book in the Springer Publisher's Lecture Notes in Artificial Intelligence series.

In the Seminar we have collaborated with PLANET's (the EU Network of Excellence in AI Planning and Scheduling), in particular the PLANET technical coordination unit (TCU) on Robot Action Planning. With this collaboration we could exploit synergies between the two activities. PLANET has provided us with an excellent infra structure and support for European researchers in the field. On the other hand, the Seminar has provided valuable input to PLANET's road map for research in robot planning. Nationally, the seminar has cooperated with the DFG *Schwerpunktprogramm Kooperierende Teams mobiler Roboter in dynamischen Umgebungen* (Cooperating Teams of Robots in Dynamic Environments). It was an explicit aim of the seminar to bring participants of these two activities together and, moreover, to deepen the contact with researchers from the U.S. working in the field.

The seminar has fueled several initiatives for research project proposals and multi lateral cooperations between participants of the seminar. In one of them we intend to make a project proposal that comprises several European research groups to develop a common testbed for plan-based control of autonomous robots. This intended project should then be followed up in a broader context and in the form of a basic research initiative within the next ESPRIT framework.

Other results of this seminar include a better, deeper understanding of issues and methodologies in plan-based control of robotic agents as well as insights into relevant real-world applications of this technology. The organizers also intend to write an article about the state of the field based on the Seminar and submit it to a magazine.

Many participants have expressed that they found the seminar highly interesting and stimulating. The typical approach of Dagstuhl seminars to bring together for ample exchange researchers from different communities has once more proven its charm, where mobile robotics, artificial intelligence and agent technology were the communities most prominently present. A follow-up seminar under the same title will be held in 2003, giving those who couldn't make it in 2001 the opportunity to meet those who have expressed their firm intention to come back next time.

Contents

1	Plan Representation for Robotic Agents	6
2	XFRMLearn - A System for Learning Structured Reactive Navigation Plans	6
3	A Preliminary Report on Integrating State Abstraction and HTN Planning	6
4	Plan-Based Adaptation for Control of Robotic Agents	7
5	Reliable Multi Robot Coordination Using Minimal Communication and Neural Prediction	8
6	Coordinated Multi-robot Exploration	8
7	Mental Models for Robot Control	9
8	Planners that Learn: Why they Must; How they Might	9
9	The WITAS UAV Project: A Case Study for Planning Techniques	10
10	Reactive Natural Language	10
11	Approaches to Planning for a Mobile Robot: from Planned Activities to Robust Execution	11
12	Strength in Numbers: A Team of Robotic Agents for Surveillance	12
13	A Perspective of Plans in Robot Control	12
14	Reasoning About Robot Actions: A Model Checking Approach	13
15	Progressive planning for mobile robots	13
16	Lifelong Planning for Mobile Robots	14
17	Anchoring: a key concept for plan execution in embedded systems	14
18	Use of Cognitive Robotics Logic in a Double Helix Architecture for Autonomous Systems	15

19 Cooperative probabilistic state estimation for vision-based autonomous mobile robots	15
20 The DD&P Robot Control Architecture	16
21 Decision-Theoretic Control of Planetary Rovers	16

1 Plan Representation for Robotic Agents

Michael Beetz

Most robotic agents cannot fully exploit plans as resources for better problem-solving performance because of imminent limitations of their plan representations. In this talk I propose plan representations that are, for a given job, representationally and inferentially adequate and inferentially and acquisitionally efficient. I state what these properties mean in the context of robotic agents and describe how plan representations can be designed to satisfy them. The proposed plan representations have been successfully employed in several longterm experiments on autonomous robots.

2 XFRMLearn - A System for Learning Structured Reactive Navigation Plans

Thorsten Belker

Autonomous robots, such as robot office couriers, need navigation routines that support flexible task execution and effective action planning. In the talk I present two generations of XFRMLearn, a system that learns structured reactive navigation plans. I first describe Model- and Diagnosis-based Transformational Learning (MDTL) as a framework for learning such plans. I then identify three problems of this approach and discuss how Model- and Projection-based Transformational Learning (MPPTL -not yet fully implemented) might overcome these limitations.

3 A Preliminary Report on Integrating State Abstraction and HTN Planning

Susanne Biundo

Our approach imposes a hierarchical structure on both planning operators and state descriptions. We adopt from classical HTN planning the description of operators by abstract and primitive tasks and so-called decomposition methods. The methods are used to stepwise refine abstract tasks until a non-linear primitive plan is obtained. Tasks – abstract as well as primitive ones – are carrying pre- and postconditions. This allows to include tasks into a partially developed plan on every level of abstraction using state-based planning techniques. In order to guarantee correctness of these insertion operations on abstract plans, so-called decomposition axioms are introduced. They describe a hierarchical relation between preconditions and effects of tasks on different abstraction levels. The approach relies on a formal, logic-based semantics. This allows to define the notion of legal decompositions and with that forms the basis for a consequent and well-founded integration of operator-based and state-based planning techniques.

4 Plan-Based Adaptation for Control of Robotic Agents

Michael Bowling

The goal of this talk is to make an argument that online adaptation is a crucial element of robot systems. Examples from RoboCup are presented where failure to adapt had drastic consequences on performance. In fact all domains with uncertainty (e.g, environmental, partial observability, agents with unknown behavior) require online adaptation. After motivating the need for adaptation I discuss some approaches to achieve practical plan-based adaptive systems. I explore stochastic games and WoLF learning (Bowling and Veloso, 2001) as a framework for concurrent learning problems. I examine an opponent modelling solution that learns models of opponent behavior and then uses these models to plan solutions (Riley and Veloso, 2002). Finally, I present a very simple and practical technique to add adaptation to even reactive systems. This technique is based on a simple trial-and-error mechanism to account for uncertainty in execution. These three techniques offer very different approaches to the problem, from complex to simple. Since adaptation is a critical part of robot planning, these and other approaches need further exploration in real systems.

5 Reliable Multi Robot Coordination Using Minimal Communication and Neural Prediction

Sebastian Buck

In many multi robot applications, such as robot soccer, robot rescue, and exploration, a reliable coordination of robots is required. Robot teams in these applications should therefore be equipped with coordination mechanisms that work robustly despite communication capabilities being corrupted. In this talk I propose a coordination mechanism in which each robot first computes a global task assignment for the team that minimizes the cost of achieving all tasks, and then executes the task assigned to itself. In this coordination mechanism a robot can infer the intentions of its team mates given their belief states. Lack of information caused by communication failures causes an increase of uncertainty with respect to the belief states of team mates. The cost of task achievement is estimated by a sophisticated temporal projection module that exploits learned dynamical models of the robots. We will show in experiments, both on real and simulated robots, that our coordination mechanism produces well coordinated behavior and that the coherence of task assignments gracefully degrades with communication failures.

6 Coordinated Multi-robot Exploration

W. Burgard

In this paper we consider the problem of collaborative multi-robot exploration. This problem is enormously complex, because the search spaces grows exponentially in the number of robots. As in single-robot exploration the goal of multi-robot exploration is to minimize the overall exploration time. The key problem to be solved therefore is to choose appropriate target points for the individual robots so that they simultaneously explore different regions of their environment. We present a probabilistic approach for the coordination of teams of mobile robots which, in contrast to previous approaches, simultaneously takes into account the costs of reaching a target point and the utility of target points. We present different experiments illustrating the capabilities of our approach and demonstrating a significant reduction of the time needed to complete the exploration task.

7 Mental Models for Robot Control

Hans-Dieter Burkhard

Control of autonomous robots in dynamic environments is interesting from a cognitive point of view as well as under application view points. It is now widely accepted that emergent behavior provided by stimulus-response approaches as well as goal-directed behavior are necessary to meet the different purposes of intelligent autonomous robots. The paper discusses the use of mental states for control and compares the different approaches. The double pass architecture is proposed for combining the different requirements.

8 Planners that Learn: Why they Must; How they Might

Gerald DeJong

We argue that learning is central to planning in real-world applications such as Robotics. Consider the historical progression from classical planners to reactive systems to reinforcement learning. The model of the world for a classical planner is concise, it handles extreme nonlinearly effortlessly, and plans are automatically constructed. Unfortunately, these plans seldom work as anticipated in the real world. The model, relying on logic and theorem-proving-like inference, is brittle. Reactive systems are much less brittle, but a human implementor performs plan construction. Reinforcement Learning (direct or indirect) can be seen as acquiring its own model of the world (implicitly or explicitly). Like classical planning, the result is produced automatically without human intervention. Like reactive systems, goal-achievement behavior in the real-world is robust. Unfortunately, also like reactive systems, the solution is very narrow, typically tied to a specific problem. How can we simultaneously achieve the breadth of classical planning, the robustness of reactive systems & reinforcement learning, and the automaticity shared by classical planning & reinforcement learning? The answer must lie in learning. Only through learning (including sampling adequately from the underlying problem population of interest) can we justify that our system is likely to respect the relevant subtleties of a real world. To learn such complex concepts using only a tractable number of training examples demands that some

additional source of evidence be provided to the learner beyond training set. Conventional induction will not suffice. We suggest that a prior human-expert-supplied domain theory is needed. This suggests that approaches of the sort illustrated by Explanation-Based Learning may be successful.

9 The WITAS UAV Project: A Case Study for Planning Techniques

Patrick Doherty

The purpose of this talk is to provide a broad overview of the WITAS Unmanned Aerial Vehicle Project. The WITAS UAV project is an ambitious, long-term basic research project with the goal of developing technologies and functionalities necessary for the successful deployment of a fully autonomous UAV operating over diverse geographical terrain containing road and traffic networks. The project is multi-disciplinary in nature, requiring many different research competences, and covering a broad spectrum of basic research issues, many of which relate to current topics in artificial intelligence.

A number of topics considered are knowledge representation issues, active vision systems and their integration with deliberative/reactive architectures, helicopter modeling and control, ground operator dialogue systems, actual physical platforms, and a number of simulation techniques.

10 Reactive Natural Language

R. James Firby

Although natural language interpretation is not typically conceived as a robot planning problem, even very simple natural language exchanges often require the robot to interpret ambiguous references in the context of the immediate, real, physical situation. This talk describes a reactive natural language system built using a robot planner to get direct access to internal robot state and generate the physical actions needed to eliminate ambiguity.

The system parses each language utterance into a symbolic description that captures its basic isolated structure. The result is treated as sensory data to be interpreted by the robot planner. A key step in this interpretation is grounding object references by matching them to internal models that can be "anchored" to real physical objects. Finding appropriate real-world objects may require physically looking in the world, checking object properties, examining the speaker for physical reference cues, and querying the speaker for clarification. Interpreting query responses may require further elaboration, additional queries and simple emergent dialogs.

11 Approaches to Planning for a Mobile Robot: from Planned Activities to Robust Execution

Malik Ghallab

This talk is concerned with practical task planning for an autonomous mobile robot and with the robust achievement of planned activities. The first part introduces a functional representation for temporal planning which relies on a CSP based-approach and on a hierarchy of domain attributes for synthesizing a plan. Because of that hierarchy, the task planner can be integrated to a motion planner that takes into account the geometry and kinematics of the robot. Motion planning based on a probabilistic road-maps is explained and illustrated. The second part of the talk considers the robust execution of a plan, that is how to decompose a task such as "goto(position)" into sensory-motor actions. Our approach is to specify a collection of Hierarchical Tasks Networks, called modalities, whose primitives are sensory-motor functions. Each modality is a possible way of combining some of these functions to achieve the desired task. The relationship between supervision states and the appropriate modality for pursuing a task is learned through experience as a Markov Decision Process (MDP) which provides a general policy for the task. This MDP is independent of the environment, it characterizes the robot abilities for that task.

12 Strength in Numbers: A Team of Robotic Agents for Surveillance

Maria Gini

This talk presents the hardware and software components of a robotic team designed for security and surveillance applications. The team consists of two types of robotic agents. The first type is a larger, heavy-duty robotic platform, called the “ranger.” Rangers are used to transport, deploy, and supervise a number of small, mobile sensor platforms called “scouts,” the second type of robotic agent. In an example scenario, the scouts are deployed into an office/lab environment, navigate towards dark areas, and position themselves to detect moving objects using their cameras. A ranger communicates with each of the scouts and determines whether there are objects of potential interest within the observed area. The paper includes experimental results for individual scout and ranger-scout activities.

13 A Perspective of Plans in Robot Control

Joachm Hertzberg

Plan-based robot control has yet to prove its empiric value. We argue that plans are not in the first place a means to achieve better robot performance, but that they can serve as a means to communicate with other agents and as a tool for engineering the robot control software. An increase in robot performance through the use of plans is to be expected as a side effect of these two primary purposes. We then present guidelines for designing hybrid robot control architectures by specifying a framework for such an architecture. Its features include asynchronous concurrency among modules, usage of the plans-as-advice metaphor, and exploitation of information from the robot dynamics for the world model update. We introduce the DD&P architecture as a concrete example of a robot control architecture implementing the guidelines. DD&P is currently under development.

14 Reasoning About Robot Actions: A Model Checking Approach

Frodoald Kabanza

Robots already play an important role in industrial manufacturing for which the environment is static and well-known, but extending their use in real world applications where the operating conditions are changing and unpredictable, raises many many challenges. Reacting to unanticipated events, interacting and coordinating with other agents (living or artificial), and acquiring information about the world are all still difficult problems. These should be the direct product of the robot's capabilities to perceive, act, and process information intelligently, taking into account its state, that of the environment, and the goals to be achieved. An interesting question is the role, if any, of reasoning about the robot's actions in that intelligent processing, for what type of behaviours and contexts? This paper discusses the use of model-checking as a framework for investigating answers to some important aspects of this question. The approach proposed is to study behaviours that allow abstract, but informative models, such that a computer program can reason with them efficiently. For such behaviours, we could then use model-checking as a means for implementing off-line and online verification and planning of robot actions.

15 Progressive planning for mobile robots

Lars Karlsson

I have presented a possibilistic/probabilistic conditional planner called PTLplan, and how this planner can be integrated with a behavior-based fuzzy control system called the Thinking Cap in order to execute the generated plans. Being inspired by Bacchus and Kabanza's TLplan, PTLplan is a progressive planner that uses strategic knowledge encoded in a temporal logic to reduce its search space. Actions' effects and sensing can be context dependent and uncertain, and the resulting plans may contain conditional branches. When these plans are executed by the control system, they are transformed into B-plans which essentially are combinations of fuzzy behaviors to be executed in different contexts.

16 Lifelong Planning for Mobile Robots

Sven Koenig

Mobile robots often have to replan as their knowledge of the world changes. Lifelong planning is a paradigm that allows them to replan much faster than with complete searches from scratch, yet finds optimal solutions. To demonstrate this paradigm, we apply it to Greedy Mapping, a simple sensor-based planning method that always moves the robot from its current cell to the closest cell that it has not observed yet, until the terrain is mapped. Greedy Mapping has a small mapping time, makes only action recommendations and can thus coexist with other components of a robot architecture that also make action recommendations, and is able to take advantage of prior knowledge of parts of the terrain (if available). We demonstrate how a robot can use our lifelong-planning version of A* to repeatedly determine a shortest path from its current cell to the closest cell that it has not observed yet. Our experimental results demonstrate the advantage of lifelong planning for Greedy Mapping over other search methods. Similar results had so far been established only for goal-directed navigation in unknown terrain.

17 Anchoring: a key concept for plan execution in embedded systems

Alessandro Saffiotti

Anchoring is the process of creating and maintaining the correspondence between symbols and percepts that refer to the same physical objects. Although this process must necessarily be present in any physically embedded system that includes a symbolic component (e.g., an autonomous robot), no systematic study of anchoring as a problem per se has been reported in the literature on intelligent systems. In this talk, I will advocate the need for a domain-independent definition of the anchoring problem, and will report some initial steps in this direction. I will illustrate these steps showing experiments performed on a real mobile robot.

18 Use of Cognitive Robotics Logic in a Double Helix Architecture for Autonomous Systems

Erik Sandewall

Cognitive robots that operate in situations with strict real-time requirements need to accommodate two different approaches to time in the computational system: the use of synchronous architectures and controlled, preferably fixed duration of the main computational cycle which is natural from the control perspective and, on the other hand, the virtual impossibility of controlling deliberation time which is only partly dealt with using anytime algorithms. This raises not only an architectural challenge, but also a question how the logic of time and action that is needed in the deliberative subsystem shall be able to refer to the system's real time and to the processes that occur in real clock time.

In this talk I describe an approach to this problem which is being developed in the dialogue system part of the WITAS project, based on the author's earlier work using Cognitive Robotics Logic. The goal of WITAS as a whole is to develop methods for, and an actual system for equipping a medium-sized unmanned helicopter with high-level cognitive capabilities, including computer vision for the purpose of dynamic scene understanding, multiple levels of control of the vehicle and its sensor systems, and speech-based multi-modal dialogue with an operator on the ground. Information about WITAS and other work within this project can be found at the project's web page, <http://www.ida.liu.se/ext/witas/>.

19 Cooperative probabilistic state estimation for vision-based autonomous mobile robots

Thorsten Schmitt

With the services that autonomous robots are to provide becoming more demanding, the states that the robots have to estimate become more complex. In this paper, we develop and analyze a probabilistic, vision-based state estimation method for individual, autonomous robots. This method enables a team of mobile robots to estimate their joint positions in a known environment and track the positions of autonomously moving objects. The state estimators of different robots cooperate to increase the accuracy and reliability of the estimation

process. This cooperation between the robots enables them to track temporarily occluded objects and to faster recover their position after they have lost track of it. The method is empirically validated based on experiments with a team of physical robots.

20 The DD&P Robot Control Architecture

Frank Schoenherr

The talk presented a new technique for extracting symbolic ground facts out of the sensor data stream in autonomous robots for use under hybrid control architectures, which comprise a behavior-based and a deliberative part. The sensor data are used in the form of time series curves of behavior activation values. Recurring patterns in individual behavior activation curves are aggregated to well-defined patterns, like edges and levels, called qualitative activations. Sets of qualitative activations for different behaviors occurring in the same interval of time are summed to activation gestalts. Sequences of activation gestalts are used for defining chronicles, the recognition of which establishes evidence for the validity of ground facts. The approach in general was described, and examples for a particular behavior-based robot control framework in simulation were presented and discussed.

21 Decision-Theoretic Control of Planetary Rovers

Shlomo Zilberstein

Planetary rovers are small unmanned vehicles equipped with cameras and a variety of sensors used for scientific experiments. They must operate under tight constraints over such resources as operation time, power, storage capacity, and communication bandwidth. Moreover, the limited computational resources of the rover limit the complexity of on-line planning and scheduling. We describe two decision-theoretic approaches to maximize the productivity of planetary rovers: one based on adaptive planning and the other on hierarchical reinforcement learning. Both approaches map the problem into a Markov decision problem and at-

tempt to solve a large part of the problem off-line, exploiting the structure of the plan and independence between plan components. We examine the advantages and limitations of these techniques and their scalability.