

Third International Conference on Artificial Intelligence Planning Systems

Brian Drabble

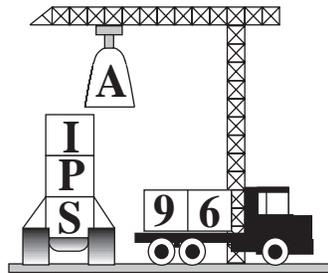
■ The Third International Conference on Artificial Intelligence Planning Systems (AIPS-96) was held in Edinburgh, Scotland, from 29 to 31 May 1996. The main gathering of researchers in AI and planning and scheduling, the conference promoted the practical applications of planning technologies. Details of the conference papers and sessions are provided as well as information on the Defense Advanced Research Projects Agency-Rome Laboratory Planning Initiative.

The Third International Conference on Artificial Intelligence Planning Systems (AIPS-96) was held in Edinburgh, Scotland, from 29 to 31 May 1996. Previous conferences were held at the University of Maryland in June 1992 (AIPS-92), organized by Jim Hendler and Drew McDermott, and the University of Chicago in June 1994 (AIPS-94), organized by Kristian Hammond.

The generation of plans and related fields, such as scheduling, resource allocation, and reasoning about action, have a long research tradition in AI. The International Conference on Artificial Intelligence Planning Systems is the main gathering of researchers in AI planning and scheduling. The key aim of the conference was to bring together the AI planning community and those involved in a number of related areas. In particular, it was hoped that by devoting two of the sessions to papers from this area, organizers could attract key researchers in AI scheduling and stimulate discussions and explore common interests and problems. Given the enthusiasm of the discussions that took place during the breaks, we achieved this aim. Another objective of the conference was

to promote the practical applications of planning technologies. Thus, a special session was devoted to describing the technological achievements of the Defense Advanced Research Projects Agency (DARPA)-Rome Laboratory Planning Initiative (ARPI). Also, two panels were convened to identify the areas that researchers should be looking into to make major impact in the real world and to discuss the problems of eliciting knowledge for planning and scheduling problems.

Over 100 papers were submitted to the conference, but only 32 were accepted. Topics covered in the proceedings included practical algorithms for achieving efficiency in plan and schedule generation, formal



results on completeness and complexity of planning domains, and the formal specification of planning and scheduling knowledge.

Invited Speakers

The speakers at the conference presented some of their latest work and ideas in intelligent planning: Alan Bundy of Edinburgh University gave a talk entitled "Proof Planning"; Dan Weld of Washington University pre-

sented his work entitled "Planning-Based Control of Software Agents"; and Brian Williams of the National Aeronautics and Space Administration (NASA) Ames Research Center gave a talk entitled "Model-Based Autonomous Systems in the New Millennium."

Each of the three talks showed how the planning techniques and systems that are currently being developed are finding positive uses in actual applications. For example, the focus of Brian Williams's talk was the intelligent planning system being developed for the *Deep Space One* spacecraft on its journey to the outer reaches of the solar system; he convincingly showed that "a little bit of planning can go a long way"!

Panels

The panels at the conference aimed to show that planning systems and techniques were reaching the point at which they could be exploited. However, as yet, this take-up has been slow. It was the aim of the two panels to try and identify why this was the case and whether the areas in which researchers were concentrating were the ones that would provide the greatest returns.

The first panel was "Planning: What Stones Remain Unturned?" Chaired by Jim Hendler of the University of Maryland, the panel included members from the U.S. academic and the United Kingdom and U.S. research and development (R&D) communities. Panel members agreed on the maturity of the techniques but disagreed about the types of problem the techniques should be addressing. One approach was to try and influence mass-market products, for example, project management tools and work-flow engines, but another approach was to concentrate on large-scale systems, for example, military command and control systems.

The second panel was "Knowledge Acquisition for Planning," chaired by Nigel Shadbolt of the University of Nottingham. The members of the panel were drawn from the United Kingdom and U.S. R&D communities. The outcomes of the panel were

that current tools for knowledge acquisition and modeling were adequate for most applications but that the developers of planning and scheduling systems should pay careful attention to acquiring and modeling certain features, such as constraints, that are common in realistic planning and scheduling domains but are not found in other task types, such as diagnosis.

Papers

The conference papers were divided into a number of sessions, each concentrating on a specific aspect of planning or scheduling. The sessions were as follows: (1) refinement planning, (2) representing and reasoning about uncertainty, (3) analysis of planning domains, (4) controlling the search in planning problems, (5) scheduling, (6) planning representations, (7) domain representations, and (8) plan execution and repair.

The papers presented were a mixture of basic research and developed ideas and techniques that were reported previously. The majority of the papers were in the basic research category and showed the current strength and interest in the areas of planning and scheduling. One common theme in a number of papers was the importance of constraints and the major part they played in identifying acceptable solutions. Constraint representation and reasoning is expected to be a major focus of research in planning and scheduling over the next few years.

Special Session with the DARPA-Rome Laboratory Planning Initiative

Continuing the theme of the practical benefits of AI planning and scheduling, the conference devoted the afternoon of the first day to a special session on the work of ARPI. The aim of the session was to highlight the achievements and successes of this \$66 million initiative in bringing AI planning and scheduling technology into use in the U.S. military. ARPI began in 1989 and is now in its third phase. The aim of the ARPI program

The aim of the ARPI program is to build relevant demonstrations of advanced concepts for planning and scheduling to respond to crises requiring military intervention.

is to build relevant demonstrations of advanced concepts for planning and scheduling to respond to crises requiring military intervention, for example, logistics support, noncombatant evacuation, disaster-relief missions, and air campaign planning. ARPI is also concerned with the use of these ideas in other applications and sectors, for example, oil-spill recovery, civil emergencies, and manufacturing. ARPI brings together university researchers, government agencies, companies, and operational military staff.

An overview of ARPI was provided by the program managers (Nort Fowler and Tom Garvey), who described the successes so far and the main technical foci of ARPI: generative planning, case-based planning, scheduling, planning under uncertainty, temporal reasoning and planning, decision-theoretic control of planning, simulation and plan evaluation, learning and mixed-initiative planning. The philosophy behind ARPI is to pick a challenging problem related to military need and bring the user in early. The program managers reported that the ARPI participants were working toward a June 1996 demonstration of air campaign planning and replanning that would last over several simulated days of operation. A successful demonstration subsequently took place.

The papers in this session covered planning and databases, transportation scheduling, planning ontologies, simulation, and mixed-initiative user interfaces. The papers presented in the session were refereed papers accepted for the conference, and they appear in the proceedings of the conference (Drabble 1996). In addition, a special volume of papers covering the technological achievements of each

project in each of the three phases of ARPI was produced to complement the session (Tate 1996).

The main message to come from the session was that ARPI was making progress and that it had made a major impact in certain areas of military operation. For example, the U.S. Department of Commerce reported in 1994 that the deployment of a single logistics support aid called the dynamic analysis and replanning tool (DART), during Operation Desert Shield paid back all U.S. government investment in AI and knowledge-based-system research over a 30-year period.

Summary

The overall message from AIPS-96 was that currently, there is a strong international research and R&D community in planning and scheduling. The techniques and systems described at the conference were being applied to ever-more complex problems, and companies and organizations were obtaining direct benefits from them. However, the number of organizations aware of programs such as ARPI was low, and for the field to progress further, there needs to be a concerted effort from the academic and R&D communities to make industry more aware of work they are doing and the benefits they could achieve.

References

- Drabble, B. 1996. Proceedings of the Third International Conference on Artificial Intelligence Planning Systems. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Tate, A. 1996. Advanced Planning Technology. Menlo Park, Calif.: American Association for Artificial Intelligence.