

The 2008 Scheduling and Planning Applications Workshop (SPARK'08)

*Luis Castillo, Gabriella Cortellessa,
and Neil Yorke-Smith*

- SPARK'08 was the first edition of a workshop series designed to provide a stable, long-term forum where researchers could discuss the applications of planning and scheduling techniques to real problems. Animated discussion characterized the workshop, which was collocated with the 18th International Conference on Automated Planning and Scheduling (ICAPS-08) held in Sydney, Australia, in September 2008.

A perennial question in the AI planning and scheduling community is why planning and scheduling, a very applicable research field, finds so little use. What keeps the fine advances in this field made over recent years hidden? The international Scheduling and Planning Applications Workshop (SPARK) was established to help address this issue.

Building on precursory events, SPARK'08 was the first workshop designed to provide a stable, long-term forum where researchers could discuss the applications of planning and scheduling (P&S) techniques to real problems. Like its immediate predecessor (the ICAPS'07 Workshop on Moving Planning and Scheduling Systems), the 2008 SPARK workshop was collocated with the International Conference on Automated Planning and Scheduling (ICAPS), a premier forum for research in AI planning and scheduling, and the International Conference on Principles and Practice of Constraint Programming (CP).

A handful of outstanding application-oriented papers are presented each year at the ICAPS conference. Time and again, in invited talks and in open microphone discussion sessions such as ICAPS's Festivus (where conference participants air their grievances in an open and entertaining way), researchers have lamented the small number of applications papers accepted at conferences such as ICAPS, CP, and the AAAI Conference on Artificial Intelligence. They call out the challenges in taking P&S research advances into the real world. The few real world P&S domains that practitioners report (such as Mars rover planning and satellite scheduling) are discussed and cited until they become banal. Some lament theoretical work evaluated on artificial International Planning Competition (IPC) benchmarks. Others high-

light the gulf between accepted modeling languages such as the PDDL family and real problems. Still others point out the nontechnical factors that hinder technology adoption.

In this context, participants to SPARK'08 were chosen based on their work on problem domains and instances under study for, or closely inspired by, a real industrial or commercial deployment of P&S techniques. Some 30 researchers and practitioners gathered in Sydney, Australia, in September 2008 for the one-day workshop.

The challenges and discussions held at each of the preceding SPARK workshops generally set the baseline for successive workshops, so that the workshop series fosters an evolving and cumulative perspective of applications and their challenges. This effort is intended to converge to an evolving set of challenges that can aid practitioners and benefit the direction of the research community.

The diverse domains where P&S research can be applied were shown in the presented papers in 2008. Applications included grid computing, agile software development management, natural language generation, power station control, personal time management, route planning for damaged aircraft, clinical process management, and a range of robot, space, and satellite problems.

Alexander Koller and Ron Petrick from the University of Edinburgh were recognized as authors of the best paper for their novel use of P&S in service of natural language generation. They introduced two challenge domains: a sentence-generation domain and the GIVE (generating instructions in virtual environment) domain. Koller and Petrick had investigated the efficiency of leading ready-made planners FF and SGPLAN. They reported that while modern planners are able to solve many moderately sized problem instances in these domains quickly, the overall planning time is dominated by the grounding step that these planners perform, rather than search, which has a pronounced effect on domains that require relatively short plans but have large universes. Koller and Petrick posited that domains such as GIVE—which, when expressed in PDDL, has

an entirely different flavor to commonly used IPC domains—exhibit shortcoming for current AI planners.

Animated discussion characterized the workshop. The burden was put on the academic community to make P&S technology more appealing and functional for potential industry users. It was said that models are often far too complex to be understood by external users; only those who build a system can develop the models. At the same time, Gérard Verfaillie of ONERA noted that uncertainty had been considered in very few of the presented papers (although another workshop at ICAPS-08 was dedicated to uncertainty). He questioned whether it was because P&S could not deal with uncertainty or because of other reasons. The former provides a challenge for the P&S research community.

It was also said during the workshop that the concept of automation or autonomy is perceived as somewhat dangerous by users, who ask, “Can I really trust the system?” Explanation and controllability are important. David Smith of NASA Ames Research Center commented that this trust could be a cultural rather than a technological or even perceptual matter. Riccardo Rasconi of ISTC-CNR related experience with the European Space Agency, where P&S technology infusion came as mixed-initiative decision-making capabilities. This gave users sufficient time (and a control over the technology) to learn to appreciate what AI systems can do for them, Rasconi argued. Perhaps, he continued, the mixed-initiative approach will lead, in due course, to a slow but constant increase in trust and a change of culture, resulting in more effective distribution of competences between the human user and the software—an experience that was found in the deployment of MAPGEN on the NASA Mars rovers.

The workshop discussion then moved to benchmarks and the need for real-world data. Jeremy Frank of NASA Ames explained the difficulties in sharing data, among them proprietary formats, sensitive information, political approval, and funding for the necessary work in preparing the data. Similar comments were made by several participants from commercial sec-

tors of industry. The domain and data set for online planning, by J. Benton, Minh Do, and Wheeler Ruml, was held up as a noteworthy recent effort. Participants agreed with the vision of a library of domains and data sets evolving alongside the workshop. The data of some of the papers presented this year has been made publicly available at the workshop website.¹

In sum, the open challenges highlighted by SPARK'08 include a lack of users' trust in innovative technology, difficulties in modeling all the aspects of real world domains, and limited communication between P&S researchers and industrial practitioners.

The proceedings from SPARK'08 are available online. Selected papers from this year's workshop will appear in a forthcoming special issue of *Computational Intelligence Journal*. The 2009 SPARK workshop will be held in Thessaloniki, Greece, in September 2009.

Acknowledgements

The organizers acknowledge the Program Committee and reviewers, and thank especially Riccardo Rasconi (ISTC-CNR) for his help.

Note

1. decsai.ugr.es/~lcv/SPARK

Luis Castillo is an associate professor in the Department of Computer Science and Artificial Intelligence at the University of Granada, Granada, Spain and cofounder and scientific consultant of IActive Intelligent Solutions. He holds a Ph.D. in computer science from the University of Granada.

Gabriella Cortellessa is a research scientist at the Institute for Cognitive Science and Technologies, National Research Council of Italy (ISTC-CNR), in Rome, Italy. Her research spans mixed-initiative systems, intelligent interfaces for planning and scheduling systems, experimental methods for evaluating intelligent system features, and user-oriented explanations. She holds a Ph.D. in cognitive psychology from the University of Rome “La Sapienza.”

Neil Yorke-Smith is a computer scientist at SRI International's Artificial Intelligence Center. His research aims to help people make decisions in complex situations, with interests that include temporal reasoning, planning and scheduling, preferences, advisable agents, intelligent interfaces, and constraint programming, and their real-world applications. He holds a Ph.D. from Imperial College London.