A Survey of the Seventh International Planning Competition

Amanda Coles   Andrew Coles
Department of Informatics, King's College London (United Kingdom)

Angel García Olaya   Sergio Jiménez   Carlos Linares López
Computer Science Deparment, Universidad Carlos III de Madrid (Spain)

Scott Sanner
NICTA & the Australian National University, Canberra (Australia)

Sungwook Yoon
Palo Alto Research Center (United States)

March 15, 2012

Abstract

In this article we review the 2011 International Planning Competition. We give an overview of the history of the competition, discussing how it has developed since its first edition in 1998. The 2011 competition was run in three main separate tracks: the Deterministic (Classical) Track; the Learning Track; and the Uncertainty Track. Each track proposed its own distinct set of new challenges and the participants rose to these admirably, the results of each track showing promising progress in each area. The competition attracted a record number of participants this year, showing its continued and strong position as a major central pillar of the international planning research community.

1 Introduction

Automated Planning is the process of finding an ordered sequence of actions that, starting from a given initial state, allows the transition to a state where a series of objectives are achieved. Actions are usually expressed in terms of preconditions and effects; i.e. the requirements a state must meet for the action to be applied, and the changes subsequently made. Domain-independent planning relies on general problem solving techniques to find an (approximately) optimal sequence of actions and has been the focus of numerous International Planning Competitions (IPCs) over the years.

The first IPC was organised by Drew McDermott in 1998. For the following 10 years it was a biennial event and remains a keystone in the world-wide planning research community: the most recent, seventh, IPC took place in 2011. The major important contribution of the first competition was to establish a common standard language for defining planning problems — the Planning Domain Definition Language (PDDL) [McDermott, 1998] — which has been developed and extended throughout the competition series. Today, the extended PDDL is still widely used, and is key in allowing fair benchmarking of planners. Participation has increased dramatically over the years and a growing number of tracks have formed, representing the broadening community — see Figure 1 for details. The three main tracks now operating are the Deterministic, Learning and Uncertainty Tracks.

The IPC has two main goals: to produce new benchmarks; and to gather and disseminate data about the current state-of-the-art. Entering a planner represents significant work, and the contribution of all participants in pushing planner development, along with the data gathered, are the major prized value of the competition. The impact of the IPC on the planning and scheduling community is broader than just determining a winner: benchmarking test sets are used for evaluating new ideas, and the defined state-of-the-art, the most recent winner, is a useful benchmark. Typically, entrants in the competition come from academia, though some industrial colleagues have been involved, and industrial sponsorship secured. The independent assessment of available systems is useful to potential users of planners outside the research community.

The competition is run by the organisers over a period of several months, with participants submitting their planning systems electronically. The results of each edition of the competition are presented in a special session of the International Conference on Automated Planning and Scheduling, ICAPS[1]. The IPC council, chaired by Lee McCluskey, oversee the competition series (and the knowledge engineering competition series ICKEPS) and are seeking chairs for the next competition, expected to take place in 2013. More information about the competition can be found on the IPC[2] website.

2 Deterministic Track

The deterministic part of the competition is the longest-running track. Its focus is on the ability of planners to solve problems across a wide range of unseen domains: a challenging test of the ability of planners to succeed as domain-independent systems. Several sub-tracks of the competition have developed over the years, with all tracks at the centre of Figure 1 being considered sub-tracks of the deterministic competition. The 2011 competition saw the introduction of a new track for multi-core planners. Furthermore, another key contribution was to release all the software used to run the competition[3], thus reducing workload for future potential organisers.

The 2011 competition followed the successful 2008 competition, and was run in a very similar way. For 2011 we decided to keep the language the same, without introducing extensions, as planners still need to 'catch up' with the currently available features. We also made use of the plan validator VAL [Howey et al., 2004]. We maintained the evaluation metrics introduced in IPC-2008, favouring quality and coverage

[1] Videos of the 2011 presentations are at http://videolectures.net/icaps2011_freiburg/
[2] http://icaps-conference.org/index.php/Main/Competitions
[3] available at http://www.plg.inf.uc3m.es/ipc2011-deterministic/FrontPage/Software

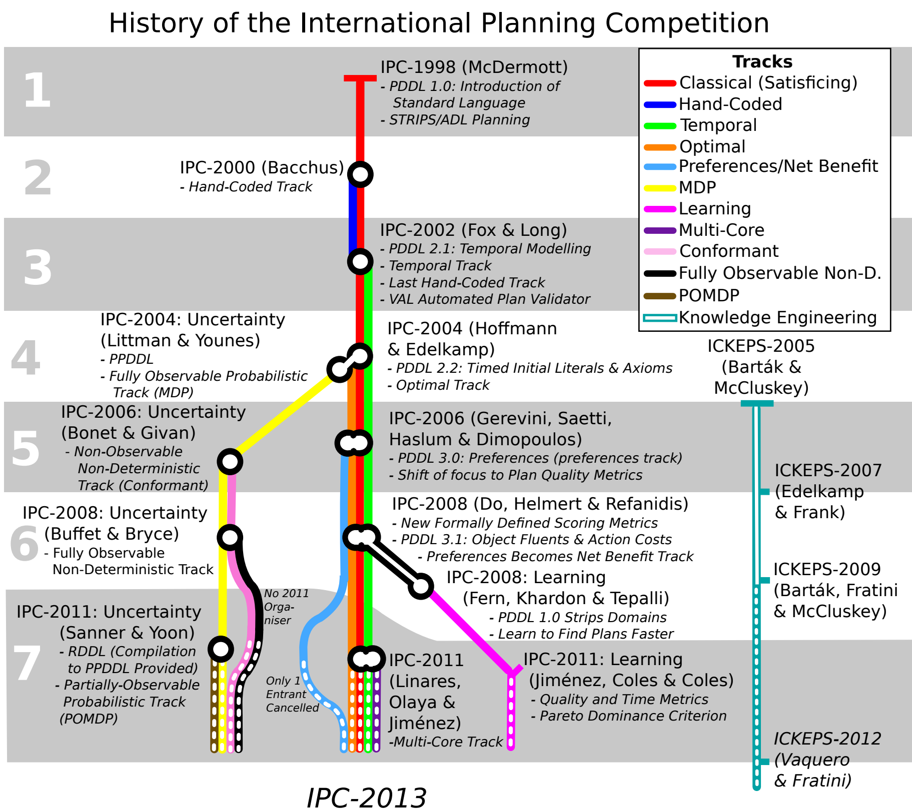

Figure 1: The History of the International Planning Competition

over problem-solving speed. Briefly, each planner is allowed 30 minutes on each planning task, and receives a score between 0 and 1. The score is the ratio between the quality of the solution found, if any (if not, it is given zero), and the quality of the best solution found by any entrant. The score is summed across all problems for a given planner: the winner and runner up for each track being those with the highest scores. Scores are not aggregated amongst tracks. We included in the results a comparison to the winner of the last competition to ensure progress is being made.

The 2011 competition was extremely popular: a record number of 55 entrants took part in the deterministic track alone, almost eight and three times more than the first and sixth competitions respectively, showing significant growth in community involvement. A summary of each of the sub-tracks follows.

2.1 Satisficing Track

LAMA won the satisficing track for the second year running, in its new incarnation LAMA-2011 (Richter, Westphal, Helmert & Röger). LAMA follows in a long history of successful planners using forward-chaining search — including previous winners HSP

(Bonet & Geffner) in 1998, FF (Hoffmann) in 2000 and FAST DOWNWARD (Helmert & Richter) in 2004 — with further guidance obtained from landmarks (facts that must be true in any solution plan). Interestingly the only non-forward-search planner to win this track was LPG (Gerevini & Serina) in 2002, using stochastic local search. A number of other interesting techniques have been seen throughout the years, including the use of pattern databases, and planning as satisfiability. 9 out of 27 of the planners in 2011 outperformed the 2008 winner LAMA-2008 (Richter & Westphal), showing good progress in the state-of-the-art.

2.2 Multi-Core Track

With the advent of parallel computers at affordable prices we wanted to ask the question: can planners using multiple cores at the same time perform better than using the single core allowed in the classical track? The winner of the multi-core track was ARVANDHERD (Nakhost, Mueller, Schaeffer, Sturtevant & Valenzano); but it did not outperform the classical-track winner, LAMA-2011. This is not so concerning, however — the history of the IPC shows that classical planners are highly engineered in terms of data structures, and are difficult to beat in the first editions of new tracks.

2.3 Temporal Track

Since the introduction of PDDL 2.1 in 2003, only a subset of the temporal planners available have been able to reason with the full temporal semantics of the language. As such, for the 2011 temporal track, we included a special class of temporal problems that include required concurrency [Cushing et al., 2007]. That is, no solution to the problem exists if the planner is not able to run two actions in parallel at the same time. The most successful planners in this track were the winner DAEYAHSP (Dréo, Schoenauer, Savéant & Vidal) and runner up ex-aequo YAHSP2-MT (Vidal) which performed best on the standard temporal problems, and runner up ex-aequo POPF2 (Coles, Coles, Fox & Long), which was the only planner to solve problems in all domains with required concurrency.

2.4 Optimal Track

As planning technology develops, writing planners that find optimal, as opposed to simply satisfying, solutions to problems becomes more feasible. FAST DOWNWARD STONE SOUP 1 (Helmert, Hoffmann, Karpas, Keyder, Nissim, Richter, Röger, Seipp & Westphal) won this year's competition outperforming the new version of the 2008 winner, GAMER (Edelkamp & Kissmann). FAST DOWNWARD STONE SOUP is portfolio based, in contrast to the symbolic search using BDDs of GAMER. The major shift towards forward search and away from planning as satisfiability in the two most recent competitions can be attributed to a change in the definition of optimality: the last two competitions have required a lowest-cost plan; whereas previous editions required a solution with the minimum number of actions. The former is much less amenable to a planning as satisfiability approach.

3 Learning Track

Efficient domain-independent search is a major challenge for AI. Using a single solver for many different problems significantly reduces human effort; the trade-off being that domain-specific systems, whilst time consuming to write, are generally much more efficient. Creating a system that can automatically learn to solve problems more efficiently is a promising approach for combining the advantages of both types of systems. This is the inspiration for research in learning for planning, a topic widely explored since the 1970s. The first IPC learning track in 2008 [Fern et al., 2011], was an important milestone for research in learning in planning, providing a platform for fair comparison. The track comprises two phases: a *learning* phase where the planners, given training problems, learn domain-specific knowledge; and an *evaluation* phase, where the planners exploit this knowledge in solving a set of unseen problems.

We took much inspiration from the 2008 learning track in organising its 2011 successor. However, in light of lessons learnt we make several changes to the running of the competition. A somewhat controversial outcome of the first learning track was that best-performing planners on the *evaluation* phase were not those that improved the most upon learning, indeed the winner showed little improvement, and several planners performed worse after learning. OBTUSEWEDGE (Yoon, Fern & Givan), awarded best learner in 2008, was one of the few planners to improve. A major innovation in 2011 was to use Pareto dominance as the metric for determining competition winners: a planner must both perform better than other planners *and* must have improved more by learning in order to be considered 'better' than its competitor. We further extended the scope for learning by allowing a longer learning period and providing problem generators, to allow an unrestricted number of available training problems.

A total of eight systems participated, broadly falling in to two categories: parameter tuners, learning to adjust the parameters of planners (or portfolios) for best performance; and knowledge learners, planners learning heuristics or policies for the given domain. The competition made use of many previous planning benchmarks, generating larger challenging instances, and introduced two new domains challenging for commonly used delete relaxation heuristics. These were the Spanner domain, in which delete relaxation planners tend to head towards dead ends, challenging planners to learn to avoid them; and the Barman domain, in which delete relaxation misses relevant knowledge about the state of limited resources.

The results of the 2011 competition painted a much more positive picture of learning in planning than those of its predecessor. Out of eight participants, six improved performance with learning in seven of the nine domains. Further, four of the competitors outperformed the deterministic track winner, LAMA-2011 (Richter, Westphal, Helmert & Röger), demonstrating that learning can improve upon the state-of-the-art. The winner PBP2 (Gerevini, Saetti & Vallati), uses statistical learning to define the time-slots dedicated to each planner in its portfolio. The runner up, FD-AUTOTUNE (Fawcett, Helmert, Hoos, Karpas, Röger & Seipp), learns the best set of parameters for the popular planner FAST-DOWNWARD (Helmert). The most successful group of planners were parameter tuning systems, the results reveal a major open challenge: making planners that learn knowledge from the domain (e.g. macro-action, heuristic or policy learners) competitive with the state-of the art.

4 Uncertainty Track

The uncertainty part of the IPC was initiated in 2004 by Michael Littman and Håkan Younes with the introduction of PPDDL, the probabilistic extension of PDDL [Younes et al., 2005]. PPDDL extends PDDL with stochastic action effects, allowing a variety of Markov Decision Processes (MDPs) to be encoded in a relational PDDL-like manner. The 2006 competition (Givan & Bonet) added a track for Conformant planning (i.e., non-observable non-deterministic domains) and the 2008 competition (Bryce & Buffet) added a track for fully-observable non-deterministic (FOND) domains. In the 2011 competition, we dropped the Conformant and FOND tracks due to lack of interest, but added a partially observed MDP (POMDP) track. We also made a major change of language from PPDDL to RDDL [Sanner, 2010] (while providing automated translations from RDDL to ground PPDDL and factored MDPs and POMDPs), which allowed modeling a variety of new problems with stochasticity, concurrency, and complex reward and transition structure not jointly representable in lifted PPDDL. The 2011 competition saw five MDP and six POMDP planner entrants.

Previous competitions saw the emergence of FF-REPLAN [Yoon et al., 2007] — which replanned on unexpected outcomes in a determinised translation of PPDDL — as an influential and top-performing planner. With our language change from PPDDL to RDDL in 2011 and our variety of new problem domains, planners based largely on the UCT Monte Carlo tree search algorithm [Kocsis and Szepesvári, 2006] placed first in both the MDP and POMDP tracks in the 2011 competition. For the MDP track, the winner was PROST (Keller & Eyerich), which used UCT in combination with determinisation techniques to initialise heuristics; the runner up was GLUTTON (Kolobov, Dai, Mausam & Weld), which used an iterative deepening version of RTDP [Barto et al., 1995] with sampled Bellman backups. For the POMDP track, the winner was POMDPX NUS (Wu, Lee & Hsu), which used a Point-based Value Iteration (PBVI) technique [Kurniawati et al., 2008] for smaller problems, but a POMDP-variant of UCT [Silver and Veness, 2010] for larger problems; the runner up was KAIST AILAB (Kim, Lee & Kim), which used a symbolic variant of PBVI [Sim et al., 2008] with a number of enhancements.

Evaluation for the 2004, 2006, and 2008 competitions relied on analysis of one or more of the following metrics: (1) average action cost to reach the goal, (2) average number of time steps to reach the goal, (3) percent of runs ending in a goal state, and (4) average wall-clock planning time per problem instance. Because lack of planner attempts on some harder domains made it difficult to aggregate average performance results on these metrics, we introduced an alternate purely reward-based evaluation approach in 2011 — for every problem instance of every domain, a planner was assigned a normalised $[0, 1]$ score with the lower bound determined by the maximum average performance of a noop and random policy and the upper bound determined by the best competitor; any planner not competing or underperforming the lower bound was assigned a score of 0 and all normalised $[0, 1]$ instance scores were averaged to arrive at a single final score for each planner.

A recurring debate at each competition is whether problem domains have reflected the full spectrum of probabilistic planning (e.g., [Little and Thiébaux, 2007]). This issue partially motivated our change from PPDDL to RDDL in 2011 in order to model stochastic domains like multi-intersection traffic control and multi-elevator control that could not be modeled in lifted PPDDL. How the language and domain choice for the 2013 IPC shapes up remains to be seen; however, given the profound influence the uncertainty track of the IPC has had on the direction of planning under uncertainty research in the past seven years, we believe it is imperative that the competition domains in 2013 are chosen to ensure the greatest relevance to end applications of interest to the planning under uncertainty community.

References

[Barto et al., 1995] Barto, A. G., Bradtke, S. J., and Singh, S. P. (1995). Learning to act using real-time dynamic programming. *Artificial Intelligence*, 72:81–138.

[Cushing et al., 2007] Cushing, W., Kambhampati, S., Mausam, and Weld, D. (2007). When is temporal planning *really* temporal planning? In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 1852–1859, Hyderabad, India.

[Fern et al., 2011] Fern, A., Khardon, R., and Tadepalli, P. (2011). The first learning track of the international planning competition. *Machine Learning*, 84:81–107.

[Fox and Long, 2003] Fox, M. and Long, D. (2003). PDDL2.1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124.

[Gerevini and Long, 2005] Gerevini, A. and Long, D. (2005). Plan constraints and preferences in PDDL3. Technical report, Department of Electronics for Automation, University of Brescia, Italy.

[Gerevini et al., 2009] Gerevini, A. E., Haslum, P., Long, D., Saetti, A., and Dimopoulos, Y. (2009). Deterministic planning in the fifth international planning competition: PDDL3 and experimental evaluation of the planners. *Artificial Intelligence*, 173:619–668.

[Hoffmann and Edelkamp, 2005] Hoffmann, J. and Edelkamp, S. (2005). The deterministic part of IPC-4: An overview. *Journal of Artificial Intelligence Research*, 24:519–579.

[Howey et al., 2004] Howey, R., Long, D., and Fox, M. (2004). VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *The Sixteenth IEEE International Conference on Tools with Artificial Intelligence (ICTAI-2004)*, pages 294–301, Boca Raton, Florida, United States.

[Kocsis and Szepesvári, 2006] Kocsis, L. and Szepesvári, C. (2006). Bandit based Monte-Carlo planning. In *Proceedings of the 17th European Conference on Machine Learning (ECML-06)*, pages 282–293.

[Kurniawati et al., 2008] Kurniawati, H., Hsu, D., and Lee, W. S. (2008). SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Proceedings of Robotics: Science and Systems IV*, Zurich, Switzerland.

[Little and Thiébaux, 2007] Little, I. and Thiébaux, S. (2007). Probabilistic planning vs. replanning. In *ICAPS Workshop on IPC: Past, Present and Future*.

[Long and Fox, 2003] Long, D. and Fox, M. (2003). The 3rd international planning competition: Results and analysis. *Journal of Artificial Intelligence Research*, 20:1–59.

[McDermott, 1998] McDermott, D. (1998). PDDL the planning domain definition language. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.

[Sanner, 2010] Sanner, S. (2010). Relational dynamic influence diagram language (RDDL): Language description. http://users.cecs.anu.edu.au/~ssanner/IPPC_2011/RDDL.pdf.

[Silver and Veness, 2010] Silver, D. and Veness, J. (2010). Monte-Carlo planning in large POMDPs. In *Proceedings of 24th Conference on Neural Information Processing Systems (NIPS-10)*, pages 2164–2172.

[Sim et al., 2008] Sim, H. S., Kim, K.-E., Kim, J. H., Chang, D.-S., and Koo, M.-W. (2008). Symbolic heuristic search value iteration for factored POMDPs. In *Proceedings of the 23rd national conference on Artificial intelligence (AAAI-08)*, pages 1088–1093.

[Yoon et al., 2007] Yoon, S., Fern, A., and Givan, R. (2007). FF-replan: A baseline for probabilistic planning. In *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS-07)*, pages 352–359.

[Younes et al., 2005] Younes, H. L. S., Littman, M. L., Weissman, D., and Asmuth, J. (2005). The first probabilistic track of the international planning competition. *Journal of Artificial Intelligence Research (JAIR)*, 24:851–887.

5 Short Bio

Amanda Coles holds an EPSRC research fellowship, and is based in the Department of Informatics, King's College London and was a co-organizer of the IPC-2011 Learning Track. Her research interests include planning with preferences, and with time and resources. She has co-authored several planners, including MARVIN, a macro-learning planner that competed in IPC-2004, and more recently LPRPG-P, a planner handling preferences. She is an author or co-author of 30 publications on AI planning.

Andrew Coles is lecturer in the Department of Informatics, King's College London. His research interests include planning with rich domain models, and in learning and inference for planning. He has over 30 publications in the area of AI planning, and has co-written several expressive temporal planners, including CRIKEY3 and POPF; the latter becoming the basis of application-based industrial work. He was a co-organizer of the Learning Track of IPC-2011.

Angel García Olaya is an associate professor in the Planning and Learning Group (PLG) of Universidad Carlos III de Madrid, Spain. His research interests are planning with soft-goals and planning and execution for robotics and real-time environments. He was a co-organizer of the deterministic track of the IPC-2011.

Sergio Jiménez is a research assistant at the PLG of Universidad Carlos III de Madrid, Spain. His research interest is learning for planning. Sergio was co-organizer of the Learning Track of the IPC 2011.

Carlos Linares López is associate professor at the PLG of Universidad Carlos III de Madrid, Spain. His research interests are domain-dependent problem solving with heuristic search techniques and domain-independent automated planning. He was co-organizer of the deterministic track of the IPC-2011.

Scott Sanner is a Senior Researcher at NICTA and an Adjunct Research Fellow at the Australian National University. He was a co-organizer of the Uncertainty Track of the IPC-2011.

Sungwook Yoon is a Research Scientist at Palo Alto Research Center. He was a co-organizer of the Uncertainty Track of the IPC-2011.