

The 2014 International Planning Competition: Progress and Trends

Mauro Vallati, Lukáš Chrpa, Marek Grześ, Thomas L. McCluskey, Mark Roberts, Scott Sanner

Abstract

We review the 2014 International Planning Competition (IPC-2014), the eighth in a series of competitions starting in 1998. IPC-2014 was held in three separate parts to assess state-of-the-art in three prominent areas of planning research: the deterministic (classical) part (IPCD), the learning part (IPCL), and the probabilistic part (IPPC). Each part evaluated planning systems in ways that pushed the edge of existing planner performance by introducing new challenges, novel tasks, or both. The competition surpassed again the number of competitors than its predecessor, highlighting the competition's central role in shaping the landscape of ongoing developments in evaluating planning systems.

Introduction

Automated planning studies the problem of reasoning about actions to achieve goals or to maximise a reward. Actions are usually expressed in terms of preconditions and effects. Preconditions indicate the requirements that must hold to apply the action, while effects are the consequence (including the cost) of applying the action to the state of the world. Automated planning has been applied to diverse, real-world application areas such as space exploration, manufacturing, machine tool calibration, and road traffic management.

In order to foster the development and comparison of planning approaches, to assess the state-of-the-art in planning, and to coordinate new challenging benchmarks the International Planning Competition (IPC) has been organised every two or three years since 1998.

This article summarises the eighth IPC held in 2014 (IPC-2014), which focused on aspects that have a significant impact on planning and, in general, on the AI community: both (i) identifying the emerging planning trends, and (ii) describing the newly introduced domains, which have been designed for investigating the applicability of planning techniques to a range of real-world applications. More information about the competition, including complete results, source code of planning systems, and domain models, can be found on the IPC portal for all years and editions of the competition¹. A summary of the history behind these parts can be found at in the article about the prior competition, IPC-2011, (Coles et al. 2012).

Similar to IPC-2011, the competition was held in three distinct parts. The **deterministic part (IPCD)** has been running since IPC-1998 and is focused on fully observable environments where actions are atomic with deterministic effects, and planning is episodic. The **learning part (IPCL)** relaxes the episodic assumption to allow planners to learn from prior experience with the domain model or search process. The **probabilistic part (IPPC)** introduces episodic, cost-optimal problems with stochastic transitions, and (optionally) partial observability.

Deterministic Part (IPCD-2014)

IPCD is the longest running part of the competition and its evaluation tracks have evolved as the research community focused on new challenges. IPCD-2014 introduced three innovations: the Agile Track, a protocol for problem selection, and a unified planner submission system called DES.

The Agile Track evaluates how quickly planners solve challenging problems. The cutoff CPU time limit is 5 minutes compared to the usual limit of 30 minutes. This has been done because, in many planning applications, it might not be possible to wait a significant amount of time for having a plan to use; satisficing plans are required as soon as possible and, eventually, they can be optimised.

IPCD-2014 introduced a reproducible, unbiased and general protocol for efficient selection of testing problem instances². Clearly, such a selection is an extremely critical step of every competition, since it is strongly related to performance and outcomes of planning engines. This approach relies on: (i) generation of a large set of benchmarks; (ii) execution of anonymised planning systems on generated

¹<http://ipc.icaps-conference.org/>

²<http://helios.hud.ac.uk/scommv/IPC-14/selection.html>

instances; (iii) selection of suitable benchmarks. In particular, the third step aims to remove trivial and too complex planning problems.

Also, a new technique for submitting planners has been designed and exploited: the DES system. Such a system allowed participants to configure, compile, test and submit their planning engines directly on the competition premises. This minimised the number of last-minute bugs, that are usually mainly due to compiling issues, and let teams understand the actual performance of planners on the competition machines, which can significantly change on different hardware / software configurations (Howe and Dahlman 1993).

IPCD-2014 followed the evaluation formats of IPCD-2008 and IPCD-2011. The modelling language was the same as aforementioned editions of the competitions, but new core features required to be supported by planners were introduced: negative preconditions and conditional effects. This was done in order to foster features support and promoting planners that can be more appropriate to exploit in real-world planning applications. As in previous editions of the deterministic track, the VAL (Howey, Long, and Fox 2004) tool was used for validating the plans provided by planners. In all but the Agile track the evaluation metrics of IPC-2008 and IPC-2011, which favour quality and coverage over problem-solving speed, were maintained. Briefly, each planner receives a score between 0 and 1 for each planning task. The score is the ratio between the quality of the solution found, if any (0 if no solution is provided), and the quality of the best solution found by any competitor. The score is summed across all problems for a given planner, the winner is the planner with the highest score. Scores are not aggregated across tracks. In the Agile track, a similar metric is used, but scores are given by evaluating runtime. In this case, each planners receives a score between 0 and 1 for each planning task. The score is the ratio between the CPU time needed for solving the problem, 0 if no solution is provided, and the CPU time needed by the fastest competitor. In order to evaluate progress made in the planning area, results included a comparison to the winner of the last competition.

IPCD-2014 received a record number of submissions: 67 planners took part in the deterministic track alone. In the following sections we will discuss trends of the planning area that emerged from the competition, and the new domain models introduced.

Results and Trends

IPCD-2014 held five tracks: agile (15 participants), multicore (9), optimal (17), satisficing (20) and temporal satisficing (6). Two tracks, the temporal optimal and preferences tracks, were cancelled due to too few competitors. We remark that few planners can deal with preferences or temporal models, which limits applicability of these systems in many real-world applications. We also note the number of entrants of the 2014 temporal satisficing track was lower than the corresponding 2011 track. It should be remarked that multicore solvers are not yet as well engineered as classical planners, a fact also observed IPCD-2011 (Coles et al. 2012), though this may be due to the multicore track being recently added.

In the sequential satisficing track, IBaCoP2 (Cenamor, de la Rosa and Fernández), which is a portfolio-based approach combining a set of state-of-the-art planning engines, was declared as the winner and Mercury (Katz and Hoffmann), which uses best first search with partial delete-relaxation heuristics, was declared as the runner-up. In the sequential optimal track, SymBA*-2 (Torralba, Alcázar, Borrajo, Kissmann and Edelkamp), which uses bidirectional blind search with perimeter abstraction heuristics, was declared as the winner and cGamer (Torralba, Alcázar, Kissmann and Edelkamp), an extension of the Gamer planner (Kissmann and Edelkamp 2011) which won the IPC-2008, which uses bidirectional symbolic search, was declared as the runner-up. In the sequential agile track, YAHSP3 (Vidal), which performs a search embedding delete-relaxed heuristics, was declared as the winner and Madagascar-pC (Rintanen), which translates planning problems into SAT, was declared as the runner-up. In the sequential multi-core track, ArvandHerd (Valenzano, Nakhost, Müller, Schaeffer and Sturtevant), which is a portfolio-based approach combining random-walk and best-first search based planning, was

declared as the winner and IBaCoP (Cenamor, de la Rosa and Fernández), a variant of IBaCoP2, was declared as the runner-up. In the sequential temporal track, YAHSP3-MT (Vidal), a variant of YAHSP3, was declared as the winner and Temporal Fast Downward (Eyerich, Keller, Aldinger and Dornhege), an extension of Fast Downward (Helmert 2006), a well known heuristic search based planner, was declared as the runner-up.

As confirmed by the record number of participants and the good overall performance, many high-performance sequential planners have been developed. We believe this to be the result of the availability of well-documented and supported planning platforms, such as FF (Hoffmann and Nebel 2001; Hoffmann 2003), Fast-Downward (Helmert 2006) and the very recent LAPKT³, that makes easier to test and develop new techniques. On the other hand, this leads to a large number of planning systems which are similar, and therefore share most of the weaknesses and strengths; for instance, 29 systems out of 67 are built on top of Fast-Downward (Helmert 2006).

In order to encourage the development of innovative planning techniques, and for avoiding the proliferation of too similar planning systems, two special jury awards were given for the most innovative planners. Mercury (Katz and Hoffmann) is based on the red-black heuristic, which relaxes only some state variables in order to balance between taking advantage of delete-relaxation and mitigating its drawbacks (Katz, Hoffmann, and Domshlak 2013). RPT (Alcázar, Fernández, Borrajo and Veloso) exploits Rapidly exploring random trees (RRTs), which is a well known technique in path planning, in order to decompose planning tasks into much smaller subtasks that connect randomly generated (non-spurious) states (Alcázar, Veloso, and Borrajo 2011).

The newly introduced Agile track evaluated planners on how quickly they found their first solution (solution quality was not considered). The speed of planners was not assessed in previous competitions, and the rationale behind introducing Agile track was to encourage development of “fast planners” since many application areas need solutions quickly. We found that planners usually performed very well in 2 or 3 domains, however, in other domains their performance was poor. On the other hand, most of the problems were solved by at least one of the planners.

Portfolio-based planning relies on combining different planning techniques for solving an instance; 29 out of 67 systems can be broadly classified as a portfolio planner. In IPCD-2011 portfolios were mainly static i.e., they were configured once on a set of training instances and exploited on testing problems. In contrast, IPCD-2014 saw a significant number of portfolio-based entrants that exploit different approaches for either online or offline configuration. Portfolio-based planners performed very well in the satisficing track and the multicore track participants were all portfolio-based. On the other hand, they tended to underperform in the optimal and agile tracks, where combining techniques usually increases coverage at the expense of runtime performance.

In terms of progress made by planning engines, LAMA-11 (Richter and Westphal 2010; Richter, Westphal, and Helmert 2011) the winner of the 2011 competition, would have placed 12th in the IPCD-2014 quality rankings, which reveals a strong trend of progress in the state of the art, larger than the trend we observed between IPCD-2011 and IPCD-2008 (Coles et al. 2012). To evaluate progress on runtime, we considered LPG-td (Gerevini, Saetti, and Serina 2003; 2010) and Metric-FF (Hoffmann and Nebel 2001; Hoffmann 2003); they would have ranked respectively the 13th and the 17th position (out of 17) of the Agile track. Interestingly, it is worth noting that the winner of the sequential temporal track, YAHSP3-MT is not able to address concurrency, and it won due to its very good performance on a small number of domains.

Domain Models Introduced

Since one of the aims of the competition is to evaluate the domain-independent performance of planners, and given the fact that planning techniques can be finely

³Available at: <http://lapkt.org>

“tuned” on existing and known domains, it is of crucial importance that new challenging domains and problems are introduced at every edition. Moreover, competition is a good occasion for evaluating how state-of-the-art planners perform in potential application domains. In 2014, we introduced 9 new domains, out of 23 that have been used among all the tracks. Some of them have been specifically designed for testing planners’ ability in handling the required PDDL features, and had not directly lead to applications. On the other hand, several domains deal with relevant real-world problems. It is the case of traffic control, which is modelled in three domain models: RTAM (Shah et al. 2013), in which planners have to deal with accident management i.e., planning actions to be taken by police, ambulances, etc., and, CityCar and MapAnalyser (M. Vallati and L. Chrupa), where planners are required to shape a map, in terms of road connections between junctions, for maximising the flow of cars between different areas of the city. Other newly introduced domains are: CaveDiving (N. Robinson, C. Muise, and C. Gretton), where a group of divers has to be organised in order to visit a number of underwater caves, Maintenance (J. Rintanen) that schedules maintenance of aircrafts in airports, ChildSnack (T. de la Rosa and R. Fuentetaja) that creates sandwiches and serves them to child, Tetris (M. Vallati) that models the well known Tetris game, Hiking (T.L. McCluskey) that organises hiking trips, and, finally, the Genome Edit Distances domain (Haslum 2011) that encodes the problem of finding a minimum cost sequence of operations that transforms one genome (signed permutation of genes) into another.

It is worth noting that – for various reasons – a number of new domains were not suitable to be used in the competition. Domains that were submitted but not used in the competition included: Airport (W. Westerman and T. L. McCluskey), where automated planning is exploited for performing ground traffic control operations in airports; Nrp (J. Piip and J. Ernits) that models the nurse scheduling problem in a hospital; Pizza (T. de la Rosa and R. Fuentetaja) that cuts pizzas into slices and then serves them to customers; Crisp (Koller and Hoffmann 2010) that describes the problem of sentence generation in natural languages; and Calibration (Parkinson et al. 2014) that optimises the calibration process of machine tools.

Learning Part (IPCL-2014)

IPCL-2014 built upon the two previous learning parts to include a Quality Subtrack that employed the plan quality evaluation from the deterministic track. This track is designed to evaluate learning versus non-learning planners in the context of problems similar to the deterministic track. An Integrated Execution Subtrack was proposed but did not have enough competitors to run.

The Quality subtrack featured three awards: overall, best learner, and basic solver. The overall award included any approach fitting within the competition framework (CPU, memory, disk space, etc.), so there were few, if any, restrictions for this award. The hope in keeping the format similar to the IPCD is to facilitate direct comparison to IPCD results. The best learner award identifies the planner that most improved the difference in quality (i.e., the learning delta) when learning was applied over when it was not. In numerous conversations leading up to the competition, there was no clear consensus on how to assess the learning delta and some felt the pareto-optimal metric from IPCL-2011 confounded the two metrics and was too easily manipulated (i.e., a competitor can artificially inflate the learning delta by failing during the no-knowledge runs). Thus, metrics were evaluated separately with an elimination round to ensure competitors could not artificially inflate the learning delta.

The basic solver award was introduced in IPCL-2014. While any definition of a basic solver can be perceived as somewhat arbitrary, the intent was to make a succinct definition available before the competition and take one small step toward encouraging fundamental research in basic solvers without limiting any approach. A basic solver is defined as any single (meta) algorithm that does not leverage more than one general purpose solver at its core. This definition includes as a basic solver any meta-algorithmic approach (e.g., Iterated Local Search, Iterated WA*,

managing calls to one SAT solver, randomised restarting A*, etc.) provided the parametrised variants use the same core solver. To make this concrete we will use some examples: A core solver cannot itself be an ensemble of heterogeneous solvers. A fallback strategy that applies a different core solver is considered more than one core solver and is excluded as a basic solver. Simultaneously searching with different heuristics (with possibly distinct open lists) is no different in spirit than making iterated calls of WA* with different weights and is included as a basic solver. A randomised restart algorithm that adjusts its restart strategy is included; iterated solvers that select (or adjust) parameters for a single base algorithm is similar in spirit to other iterated meta-algorithmic approaches. A planner that re-encodes the task for a single core solver (while possibly learning to select among distinct encodings) is considered a basic solver. Competitors certified their planners as a basic solver subject to a final decision by the organisers.

Competitors were judged on 6 domains chosen from previous competitions: elevators, floortile, nomystery, parking, spanner, and transport. No new domains were submitted for this track, so the domains from the 2008 and 2011 deterministic and learning tracks were used. In order to use domains from these competitions, the language for the learning part of IPC-2014 was a restricted subset of PDDL 3.1. All solutions were validated with VAL (Howey, Long, and Fox 2004). A challenge in the learning competition is to generate problem instances that extend beyond non-learning approaches, lie within reach of the anticipated performance of learning approaches, extend beyond the most capable learning planners, and draw the testing set from a similar distribution from the training set while avoiding results affected by overfitting. To mitigate the bias in selecting domains and in selecting the instance distributions, the domains were chosen to balance the approaches that did well on them in previous competitions and instances were selected randomly from a range.

At the start of a six-week learning stage, competitors were provided generators for these domains, a representative set of validation problems, and guideline ranges for the evaluation distributions. Errors in the domains and generators were corrected at this point. After the learning stage was complete, competitors were provided runs from selected validation problems to ensure that their planner was performing as expected. Problems found in those runs were corrected before collecting the final results. For the final evaluation, 5 problems from each domain were randomly generated from the distributions, resulting in 30 problem instances. Competitors did not know these evaluation problems until after the results were released. The planners were run on the EC2 cloud compute platform with the support of a generous grant from Amazon Web Services; each compute platform had a compute equivalent of 2 cores and 3.75 GB memory and ran Ubuntu 12.04 LTS. To account for variations in the actual computing resources on the cloud platform, each planning system was run 30 times with and without domain knowledge on each problem instance. There are three categories of awards, each with a first, second, and third place.

Of the 14 planners from 8 teams who expressed an interest in competing for the Quality track, 11 planners from 7 teams competed in the final evaluation. 5 planners were registered as basic solvers that used only one algorithm to solve problems.

Results and Trends

The overall best quality award compares planners on the quality of the best plan they produced for each problem. The awards for best overall quality (out of a possible score of 30) go to: First Place - MIPlan (Quality: 21.88) from team members Nuñez, Borrajo, López; Second Place - Fast Downward Cedalion (Quality: 19.98) from team members Seipp, Sievers, Hutter; and, Third Place - Fast Downward SMAC (Quality: 17.45) from team members Seipp, Sievers, and Hutter.

The best learner award compares planners on the learning delta between their overall improvement on plan quality when knowledge was applied over when it was not. To ensure that the baseline performance without knowledge was fair, any *problem* solved by seven or more (i.e., half or more) planners was removed from this evaluation, resulting in 24 problem instances. No planners were eliminated

from consideration, but the ranking did change slightly due to this adjustment. The awards for best learner go to: First Place - Fast Downward Cedalion (Adjusted Quality Delta: 10.40), Second Place - Eroller (Adjusted Quality Delta: 9.97) from team members de la Rosa and Fuentetaja; and, Third Place - Fast Downward SMAC (Adjusted Quality Delta: 9.18). The basic solver award compares planners the quality of the best plan while using only a single core algorithm. The best core solver awards go to: First Place - Fast Downward SMAC (Quality: 17.45); Second Place - LLama (Quality: 14.30) from team members Virseda, Alcázar; and, Third Place - Eroller (Quality: 12.51).

Assessing the benefit of learning for classical planning remains a challenge. The elimination round seems to have partially resolved one issue in assessing a learning delta. However, there remains no consistent measurement of the computational effort invested during the learning stage, suggesting that teams with access to greater computational resources may be at an advantage. Further, it remains possible to win the overall award with a strong base system and a modest performance increase with learning. It seems worthwhile to continue discussions on how to reward learning in future competitions in spite of the challenges associated with defining such a metric.

Learning planners have advanced considerably since IPCL-2008 and IPCL-2011. The problem distributions extended the range of the previous competitions, and the most capable planners solved many of these harder problems with learning where they failed without learning. However, the lack of new domains for IPCL-2014 was unfortunate, and it will be illuminating to evaluate these systems on the new problems from IPCD-2014. It may also be worth considering in future competitions how to leverage new domains in both IPCD and IPCL so that the results are based on the same problems and can be more directly comparable.

Learning alone does not account for the best overall performance, although learning was important for ranking well in the overall award. For example, IBaCoP performed very well in IPCD-2014 and its companion learning planner, LIBaCoP, similarly placed well in the ranking before learning, but it was surpassed in the overall ranking by systems that had larger learning deltas. The best overall and best learner rankings differ and the rankings changed substantially (sometimes for the worse) when learning was applied for all but the winning overall planner, MIPlan. For the 24 problems used to assess learning delta, less than half (5 of 11) of the systems improved the quality score by an average of more than 5 problems; 3 of those were basic solvers.

Basic solvers deserve greater attention. The best three basic solvers scored overall rankings of 3rd, 7th, and 8th while scoring 2nd, 3rd, and 4th in the best learner ranking. This ranking difference underscores a performance gap between ensembles and basic solvers, reveals that basic solvers still remain viable for improved learning, and strongly suggests it is worthwhile to highlight basic solvers in future competitions. Basic solvers, which usually excel on specific problem instances, are often the atomic building blocks of ensembles, which usually push the performance envelope across problem sets. Both are necessary to advance planning and search; it is prudent to consider how we can encourage advancements in both. More broadly, separately awarding basic solvers (rather than penalising or eliminating ensemble systems) may also prove useful in other competitions facing concern over how to balance performance differences between ensemble systems and basic solvers.

The Probabilistic Part (IPPC-2014)

IPPC-2014 continued with the competition format initiated in IPPC-2011 that provided tracks for both discrete MDP and POMDP problem specifications described in the Relational Dynamic influence Diagram Language (RDDL) language (Santer 2010). Continuous tracks of IPPC-2014 were also planned but not run due to a lack of competitors.

While probabilistic tracks dating back to the IPPC-2004 used a probabilistic extension of PDDL known as PPDDL (Younes et al. 2005), RDDL was introduced in IPPC-2011 as an alternative specification language to PPDDL to provide

joint modelling of stochasticity, concurrency, and complex reward and transition structure not possible in (lifted) PPDDL. The use of RDDDL in the IPPC-2011 marked the first time that complex stochastic domains such as Traffic Control and Elevator Control could be modelled. A slight extension of RDDDL for the IPPC-2014 allowed the further modelling of real-world domains taken from the Ecology literature. As in the IPPC-2011, IPPC-2014 provided translations of RDDDL domains and problem instances to various alternative formats (e.g., factored MDPs and POMDPs and grounded PPDDL) to facilitate participation by a wide range of competitors.

IPPC-2014 used the purely reward-based evaluation metric introduced in IPPC-2011 — for each of ten problem instances and for each of eight problem domains (80 problem instances altogether), 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. In contrast to the IPPC-2011, which only used a single time limit for completion of the entire competition (leading to highly varying amounts of time allocated to different problem instances), IPPC-2014 enforced a strict per-instance time limit of 18 minutes to complete all 30 trials. This timing change was made to ensure fair comparison of planner performance on a per-instance basis.

Participation in IPPC-2014 offered a rematch of the top two competitors from the respective MDP and POMDP tracks of IPPC-2011 using variants of Monte Carlo tree search (MCTS) and online value iteration. In addition, two new competitors joined the MDP track IPPC-2014. Overall results showed that the three year head-start of the previous competitors paid off in this competition with the top two places in each track unchanged from IPPC-2011; nonetheless, results on four of the domains common between IPPC-2011 and IPPC-2014 show that planners competing in both have demonstrated significant performance improvements since 2011. Below, we elaborate on general planning trends that can be inferred from these results along with a discussion of new domains that were introduced in IPPC-2014 to provide novel technical challenges and real-world applications.

Results and Trends

With the change from PPDDL used in IPPCs 2004, 2006 and 2008 to RDDDL used in IPPCs 2011 and 2014, the field of competitors and winning methods have shifted substantially. Given the derivation of PPDDL from its deterministic PDDL subset, high-performance planners for PPDDL in IPPCs 2004-2008 often relied on deterministic replanning methods that used an underlying PDDL planner — the first and quite successful of these approaches being FF-Replan (Yoon, Fern, and Givan 2007), which replanned on unexpected outcomes in a determinised translation of PPDDL to PDDL. With the change to RDDDL in IPPC-2011 that made heavy use of exogenous stochasticity, concurrency, and general reward (cost-optimal) objectives in a finite horizon setting, the winning approaches in IPPC-2011 shifted to variants of Monte Carlo tree search and in some cases, online value iteration.

The dominance of planners using Monte Carlo tree search continued in IPPC-2014 with revised entries from the same winning teams and runners-up as in 2011. For the MDP track, the winner was PROST-2014 (Keller and Geisser) using a modified UCT approach (Keller and Helmert 2013); the runner-up was G-Pack (Kolobov, Mausam and Weld) using an iterative deepening version of labeled RTDP (Bonet and Geffner 2003) with sampled Bellman backups (Kolobov, Mausam, and Weld 2012). For the POMDP track, the winner was POMDPX-NUS (Ye, Wu, Zhang, Hsu and Lee), which used a combination of heuristic search on a sampled sparse belief tree (Somani et al. 2013) along with UCT for POMDPs known as POMCP (Silver and Veness 2010); the runner-up was KAIST-AIPR-Lab (Han, Nam, Lee, and Kim) also using POMCP in combination with symbolic heuristic search value iteration (Sim et al. 2008). All of these planners showed improvement on their previous versions on the four domains common to both the 2011 and 2014 competitions; in some domains, the 2014 versions of planners halved the average cost of their 2011 versions indicating

substantial progress in the MDP and POMDP tracks since IPC 2011. An innovation award was not offered in IPPC-2014 as done for other tracks since the number of competitors is still small and diverse and does not (yet) suffer from a proliferation of planners based on a common codebase.

To compare planners from IPPC-2008 to IPPC-2014, a variant of the Triangle Tireworld domain (Little and Thiébaux 2007) used in IPPC-2008 was translated to RDDL for use in IPPC-2014. An analysis of results from both competitions suggests that IPPC-2008 replanners could solve larger problems, albeit somewhat suboptimally, while the IPPC-2014 Monte Carlo planners were closer to optimal on the problems they could solve, but could not scale to the largest problems. This inherently reflects the strengths and weaknesses of the two approaches. Triangle Tireworld was intentionally designed to prove difficult for replanners since it required careful reasoning about the probabilities of different paths, hence the potential suboptimality of determinising approaches that discard probabilities. On the other hand, Monte Carlo tree search approaches reason directly about expected costs, but in doing so must reason about a tree of contingency paths with varying probabilities in contrast to probability-free deterministic replanning. This observation suggests that a merger of Monte Carlo tree search and replanning techniques might provide more optimal planners for longer horizons; notably techniques such as RFF (Teichteil-Königsbuch, Kuter, and Infantes 2010) that won the fully observable uncertainty track of IPPC-2008 do offer some potential guarantees for replanning by incorporating Monte Carlo simulations, but so far RFF has only been evaluated in the goal-oriented setting of IPC 2008.

One final trend observed in 2014 was that an unprecedented half of the planners entered in IPPC-2014 used symbolic (decision diagram) representations. While these decision diagram based planners did not take the top places, one of them – KAIST-AIPR-Lab – did achieve runner-up in the POMDP track. It remains to be determined whether any particular type of structure for MDP or POMDP domains can be identified where symbolic planners can definitively outperform non-symbolic planners.

Domain Models Introduced

A principle guiding the switch from PPDDL to RDDL in IPPC-2011 was to enable the representation of rich planning domains with concurrency and exogenous stochastic events such as Traffic Control and Elevator Control that could not be represented in PPDDL. IPPC-2014 sought to continue this trend by introducing three new contributed domains of interest to the broader scientific community. The first two new domains were from Ecology and respectively represent the problems of Wildfire Management (Karafyllidis and Thanailakis 1997) (i.e., how to allocate firefighting resources to protect assets) contributed by Zhenyu Yu and Invasive Species Management (Muneepeerakul et al. 2007) (i.e., how to allocate personnel resources to limit the spread of the Tamarisk invasive plants in stream systems) as modelled in (Dietterich, Taleghan, and Crowley 2013). A slight extension of RDDL for the IPPC-2014 allowed transition distribution parameters to be composed from exponentials and other elementary functions, which enabled the modeling of these two domains that used empirically derived transition models. A third domain of Academic Advising contributed by Libby Ferland represented an academic advisor's task of recommending courses for students to help them graduate as quickly as possible.

Concluding Remarks

All tracks in IPC-2014 have shown significant advances in planning technology in the last three years and, in some cases, the techniques used by winning planners in particular tracks indicate a fundamental shift in the approach that appears to be most effective. Hence the IPC-2014 has served its purpose of both encouraging and tracking progress in the planning community by comparing state-of-the-art planning technologies in a controlled evaluation setting on domains of common interest to the community. As the AI and planning communities move forward to

consider future IPCs and other similarly motivated competitions in AI in general, we remark on two imperatives that we believe future organisers should bear in mind.

First, we reinforce a remark from the IPC 2011 (Coles et al. 2012) that the competition languages and domains have profoundly influenced the direction of planning research dating back to the first IPC in 1998. Hence we believe it is imperative that future competition domains are chosen so as to maintain relevance not only to trends in the planning community but also to potential end users of planning technology. In many cases, modelling domains of interest to the broader research community may simply require elicitation from domain experts and formalisation within the language constraints of existing tracks. In other cases this may require using more expressivity already available in existing languages or even novel language extensions to include, e.g., object fluents (supported by both PDDL 3.1 and RDDDL but not used in existing competitions), constraints and timelines (supported by ANML (Smith, Frank, and Cushing 2008)), hierarchical decomposition (Erol, Hendler, and Nau 1994; Shivashankar et al. 2013), or more holistic views of the planning system (Ghallab, Nau, and Traverso 2014; Pollack and Horty 1999).

Second, we believe the adage of “what you measure is what improves” is critical to keep in mind for future competitions. Future organisers should thus consider whether the current evaluation metrics encourage planning research that finds wider adoption in the AI community and beyond. For instance, organisers might consider different or richer quality metrics that explicitly trade off time and objective quality or perhaps even a change of setting, e.g., to focus on real-time planning scenarios that limit online deliberation and require effective offline planning for a variety of scenarios that may be encountered at execution time. Furthermore, it would be useful if planner computational effort could be measured independently of its implementation in order to better compare planning paradigms, but no such measurement has been proposed to date.

In concluding, we believe that as long as future IPCs make a concerted effort to remain relevant — to model domains of interest for real-world applications and to measure what matters — future IPCs can continue to drive forward rapid progress in planning research and its practical application.

Acknowledgements

The organisers of IPCD-2014 would like to acknowledge the use of the University of Huddersfield Queensgate Grid in carrying out this work as well as Ibad Kureshi and John Brennan for their invaluable support with the cluster. The organisers of IPCL-2014 thanks Colorado State University and the Naval Research Laboratory for providing compute resources; Mark Roberts was funded partially by the Office of Strategic Defense. The IPCL and IPPC were supported by a generous grant from Amazon Web Services allowing the competitions to run on the Elastic Compute Cloud.

References

- Alcázar, V.; Veloso, M. M.; and Borrajo, D. 2011. Adapting a rapidly-exploring random tree for automated planning. In *Proceedings of SOCS*.
- Bonet, B., and Geffner, H. 2003. Labeled rtdp: Improving the convergence of real-time dynamic programming. In *Proceedings of the 13th International Conference on Automated Planning and Scheduling (ICAPS-03)*, 12–21. AAAI Press.
- Coles, A. J.; Coles, A.; Olaya, A. G.; Celorrio, S. J.; López, C. L.; Sanner, S.; and Yoon, S. 2012. A survey of the seventh international planning competition. *AI Magazine* 33(1).
- Dietterich, T. G.; Taleghan, M. A.; and Crowley, M. 2013. Pac optimal planning for invasive species management: Improved exploration for reinforcement learning from simulator-defined mdps. In *Proceedings of the 27th National Conference on Artificial Intelligence (AAAI-13)*.

- Erol, K.; Hendler, J. A.; and Nau, D. S. 1994. HTN planning: Complexity and expressivity. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI)*, 1123–1128.
- Gerevini, A.; Saetti, A.; and Serina, I. 2003. Planning through stochastic local search and temporal action graphs. *Journal of Artificial Intelligence Research* 20:239 – 290.
- Gerevini, A.; Saetti, A.; and Serina, I. 2010. An empirical analysis of some heuristic features for planning through local search and action graphs. *Fundamenta Informaticae* 105:1 – 31.
- Ghallab, M.; Nau, D.; and Traverso, P. 2014. The actor’s view of automated planning and acting: A position paper. *Artificial Intelligence* 208:1–17.
- Haslum, P. 2011. Computing genome edit distances using domain-independent planning. In *ICAPS’11 Scheduling and Planning Applications Workshop (SPARK)*.
- Helmert, M. 2006. The Fast Downward planning system. *Journal of Artificial Intelligence Research* 26:191 – 246.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- Hoffmann, J. 2003. The metric-ff planning system: Translating “ignoring delete lists” to numeric state variables. *Journal Artificial Intelligence Research (JAIR)* 20:291–341.
- Howe, A., and Dahlman, E. 1993. A critical assessment of benchmark comparison in planning. *Journal of Artificial Intelligence Research* 1:1–15.
- Howey, R.; Long, D.; and Fox, M. 2004. Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *Proceedings of the Sixteenth International Conference on Tools with Artificial Intelligence*, 294 – 301.
- Karafyllidis, I., and Thanailakis, A. 1997. A model for predicting forest fire spreading using gridular automata. volume 99, 87–97.
- Katz, M.; Hoffmann, J.; and Domshlak, C. 2013. Red-black relaxed plan heuristics. In *Proceedings of the 27th National Conference of the American Association for Artificial Intelligence (AAAI13)*, 489–495.
- Keller, T., and Helmert, M. 2013. Trial-based Heuristic Tree Search for Finite Horizon MDPs. In *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS-13)*, 135–143. AAAI Press.
- Kissmann, P., and Edelkamp, S. 2011. Improving cost-optimal domain-independent symbolic planning. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, (AAAI 2011)*.
- Koller, A., and Hoffmann, J. 2010. Waking up a sleeping rabbit: On natural-language sentence generation with ff. In *Proceedings of the 20th International Conference on Automated Planning and Scheduling (ICAPS10)*, 238–241.
- Kolobov, A.; Mausam; and Weld, D. 2012. Lrtdp vs. uct for online probabilistic planning. In *Proceedings of the 26th National Conference on Artificial intelligence (AAAI-12)*.
- Little, I., and Thiébaux, S. 2007. Probabilistic planning vs. replanning. In *ICAPS Workshop on IPC: Past, Present and Future*.
- Muneepeerakul, R.; Weitz, J. S.; Levin, S.; Rinaldo, A.; and Rodriguez-Iturbe, I. 2007. A neutral metapopulation model of biodiversity in river networks. *Journal of Theoretical Biology* 245(2):35–63.
- Parkinson, S.; Longstaff, A. P.; Crampton, A.; and Gregory, P. 2014. Automated planning for multi-objective machine tool calibration: Optimising makespan and measurement uncertainty. In *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS14)*.
- Pollack, M. E., and Horty, J. F. 1999. There’s more to life than making plans: Plan management in dynamic, multiagent environments. *AI Magazine* 20(4):71.
- Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *Journal of Artificial Intelligence Research* 39:127 – 177.

- Richter, S.; Westphal, M.; and Helmert, M. 2011. LAMA 2008 and 2011. In *Booklet of the 7th International Planning Competition*.
- Sanner, S. 2010. Relational dynamic influence diagram language (RDDL): Language description. http://users.cecs.anu.edu.au/~ssanner/IPPC_2011/RDDL.pdf.
- Shah, M. M. S.; Chrupa, L.; Kitchin, D. E.; McCluskey, T. L.; and Vallati, M. 2013. Exploring knowledge engineering strategies in designing and modelling a road traffic accident management domain. In *IJCAI 2013*.
- Shivashankar, V.; Alford, R.; Kuter, U.; and Nau, D. 2013. The GoDeL planning system: a more perfect union of domain-independent and hierarchical planning. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, 2380–2386. AAAI Press.
- Silver, D., and Veness, J. 2010. Monte-Carlo planning in large POMDPs. In *Proceedings of 24th Conference on Neural Information Processing Systems (NIPS-10)*, 2164–2172.
- 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)*, 1088–1093.
- Smith, D. E.; Frank, J.; and Cushing, W. 2008. The ANML language. In *Proceedings of Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. Despot: Online pomdp planning with regularization. In *Proceedings of 27th Conference on Neural Information Processing Systems (NIPS-13)*.
- Teichteil-Königsbuch, F.; Kuter, U.; and Infantes, G. 2010. Incremental plan aggregation for generating policies in MDPs. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10): Volume 1*, 1231–1238.
- 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)*, 352–359.
- 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.

Short Bio

Mauro Vallati is a research fellow at the PARK research group of the University of Huddersfield, United Kingdom. His main research interest is learning for planning. He was a co-organiser of the IPCD-2014.

Lukáš Chrupa is a research fellow at the PARK research group of the University of Huddersfield, United Kingdom. His research interests mainly consist of applying Machine Learning and Knowledge Engineering techniques in AI planning. He was a co-organiser of the IPCD-2014.

Marek Grześ is a lecturer at the University of Kent at Canterbury, UK. His research interests include artificial intelligence, decision-theoretic planning, probabilistic reasoning, and applications thereof. He was a co-organiser of the IPPC-2014.

Thomas L. McCluskey is a co-leader of the PARK research group of the University of Huddersfield, UK. His research interests include automated plan generation, domain modeling, information extraction, knowledge engineering, machine learning, data mining, requirements validation and formal specification. He co-founded and helps run the Knowledge Engineering for Planning and Scheduling competitions which are run biennially at ICAPS, and complement the IPC series. He was a co-organiser of the IPCD-2014.

Mark Roberts is a Postdoctoral Fellow for the National Research Council at the Naval Research Laboratory. He was a co-organizer of the IPCL-2014.

Scott Sanner is a Principal Researcher at NICTA and an Adjunct Research Fellow at the Australian National University. He was a co-organizer of both the IPPC-2011 and IPPC-2014.