

Experimenting with Lifted Plan-Space Planning as Scheduling: Aries in the 2023 IPC

Arthur Bit-Monnot

LAAS-CNRS, Université de Toulouse, INSA, CNRS, Toulouse, France
abitmonnot@laas.fr

Abstract

In this paper we give a high level overview of the ARIES planner at the time of its participation in the hierarchical track of the 2023 *International Planning Competition (IPC)*.

ARIES is an experimental solver whose aim is to jointly evolve a combinatorial solver for scheduling-like problems and an encoding of task planning that exploits it.

Overview

Focus The focus of ARIES is on optimization planning problems with rich temporal and concurrency requirements, driven by applications in robotics and logistics. As such, the problems targeted by the 2023 IPC on hierarchical planning are certainly out of its comfort zone. Nevertheless, with ARIES participation in the IPC we hope to highlight the relative weaknesses and hopefully strengths of its approach.

As HDDL is not its primary target, at the time of its submission, very little effort has been dedicated to improving the performance of ARIES on a corpus of HDDL problems. This participation is also the occasion of setting a milestone and reference point to be used as a baseline for future improvements.

Search Space In terms of search space, ARIES pertains to the family of lifted plan-space planners like IxTeT (Ghallab and Laruelle 1994) and FAPE (Bit-Monnot et al. 2020). In practice, it means that the solver does not construct snapshots of the entire state at various epochs but instead reasons on the interactions and dependencies between individual actions (through threats and causal links, Ghallab, Nau, and Traverso 2004). In addition, being lifted implies that the problem is never grounded: instead each action of the plan is parameterized with variables that must obey a set of constraints, notably reflecting the causal requirements of the action. The solver must eventually impose a value for these parameters that ensures that the plan is valid and achieves the intended objectives.

Search Strategy Unlike IxTeT and FAPE, ARIES delegates the responsibility of search to an independent solver. We adopt the approach of generating a CSP of bounded size that is submitted to a combinatorial solver, as common in SAT-based planners such as Lilotane (Schreiber 2021).

ARIES uses in own specialized solver: a hybrid CP-SAT combinatorial solver for optional scheduling.

At a very high level, the planner adopts the following procedure:

1. Parse HDDL problem files and translates them into chronicles, where each chronicle consists of a collection of timed conditions, effects and subtasks linked by shared variables and constraints.
2. From the objective tasks of the problem, build the decomposition tree with a maximum depth (initially 1).
3. Encode the decomposition tree into a Constraint Satisfaction Problem (CSP)
4. Solve the CSP with the internal combinatorial solver
5. If the CSP was proven unsatisfiable, repeat from step 2 with an increased maximum depth.

Encoding HDDL

The planning model used internally by ARIES is based on *chronicles* (Ghallab, Nau, and Traverso 2004), a formalism to compactly represent requirements and effects of a process in time. In planning, chronicles find their first usage in IxTeT (Ghallab and Laruelle 1994) and have been then extended for hierarchical planning in FAPE (Bit-Monnot et al. 2020).

Translating HDDL to chronicles is fairly straightforward as they are for the most part strictly more expressive. In practice, it mostly consists in assigning temporal semantics to a non-temporal language. An example of a translation of an HDDL method into a chronicle is given in Figure 1.

Given a problem representation as chronicles, we encode it as a CSP using the formulae of Godet and Bit-Monnot (2022). In spirit, this encoding corresponds to a lifted plan-space representation. For each action/method node of the lifted decomposition tree, it introduces decision variables representing *(i)* the presence of the action/method in the solution plan, *(ii)* its parameters and *(iii)* its start and end times. These decision variables are linked by a set of constraints, notably:

- *Decomposition constraints* that force the refinement of all tasks with exactly one method or action.
- *Support constraints* that require each condition to be either absent or supported by an effect.

```

(:method m8_send_soil_data
 :parameters (
  ?x - rover
  ?from - waypoint
  ?l - lander
  ?w1 - waypoint
  ?w2 - waypoint)
 :task (send_soil_data ?x ?from)
 :precondition (and (at_lander ?l ?w2)
                   (visible ?w1 ?w2))
 :ordered-subtasks (and
  (t1 (do_navigate1 ?x ?w1))
  (t2 (comm_soil_data ?x ?l ?from ?w1 ?w2))
 ))

```

(a) HDDDL methods for the *rovers* domain.

m8-send-soil-data

variables: $x, from, l, w_1, w_2$ (parameters)
 $t_s, t_e, t_s^1, t_e^1, t_s^2, t_e^2$ (timepoints)
task: $[t_s, t_e]$ *send-soil-data*($x, from$)
conditions: $[t_{start}]$ *at-lander*(l, w_2) = true
 $[t_{start}]$ *visible*(w_1, w_2) = true
subtasks: $[t_s^1, t_e^1]$ *do-navigate1*(x, w_1)
 $[t_s^2, t_e^2]$ *comm-soil-data1*($x, l, from, w_1, w_2$)
constraints: $t_s = \min(t_s^1, t_s^2)$
 $t_e = \max(t_e^1, t_e^2)$
 $t_e^1 < t_s^2$ (total order)

(b) Equivalent chronicle representation.

Figure 1: HDDDL and chronicle representation of an HTN method of the *rovers* domains

- *Coherence constraints* that ensure that each state variable is never given more than one value at a time.

Compared to SAT-based encodings for HTN, the size of this encoding is mostly independent of the size of the plan and the number of predicates in the state. On the other hand, it grows mostly quadratically with the size of the lifted decomposition tree, because of the support and coherence constraints.

Matching PDDL semantics Despite the overall compatibility of chronicles with {HP}DDL, several peculiarities of PDDL must be explicitly handled.

PDDL introduces the concept of *mutex actions* to forbid interfering actions from being executed at the same time. This is required in PDDL due to the instantaneous nature of actions but introduces an awkwardness for a temporal model as the handling of a condition statement depends on the place where it was asserted. To handle such cases, we extend the constraints of Godet and Bit-Monnot (2022) with *mutex constraints* that forbid a condition’s interval to meet the start of an effect from another action chronicle. Note that mutexes induced by pairs of *effects* are already enforced by *coherence constraints*.

For historical reasons, the *delete-before-add* semantics of PDDL stipulates that if a single action assigns both *true* and *false* to a common boolean state-variable, then only the positive assignment should be considered. This violates the principle that a state variable should not be assigned two variables at the same time enforced by the *coherence constraints*. ARIES does not support this particular corner case of the PDDL semantics as handling it properly in a fully lifted representation would notably complexify the encoding and may induce non-negligible runtime penalties. As a result, ARIES will consider inapplicable any action that relies on the *delete-before-add* semantics, making it incomplete in the presence of such problems. This is for instance the case of the *woodworking* domain where ARIES is unable to find a solution.

Finally, support for HDDDL syntax remains partial. For instance, the multiple inheritance of types exploited in the *UMTranslog* domain of IPC 2020 is unsupported and ARIES

would refuse to parse the problem.

Combinatorial Solver

The encoding of Godet and Bit-Monnot (2022) is generic and could in theory exploit any solver capable of representing disjunctions and reified difference constraints, which is the case of most CP, SMT or MILP solvers. Our experience with existing solvers for this encoding is however lukewarm.

Previous experiences

First let us state that the highly combinatorial structure of the encoding does not appear to play well with the linear relaxations of MILP solvers that were at difficulty even for the simplest instances.

Experiments with the Z3 SMT solver (De Moura and Bjørner 2008), showed interesting results with reasonable performance on non-hierarchical domains (Bit-Monnot 2018). Most impressive in particular was the ability of the activity-based search of SAT/SMT solver to robustly converge to a solution or prove its absence without any planning-specific knowledge. On the other hand, SMT-based solvers remain hard to tune with limited support for optimization. While Z3 excelled at proving unsatisfiability, the runtime to a first solution remained very high even on simple problems.

Our own experiments with CPOptimizer (Laborie et al. 2018), a leading CP solver for scheduling, showed that it had complementary strength: fine-tuning and strong propagation allowed it to quickly converge to a first solution and incrementally improve it. A key feature for strong propagation was a native support of optional activities, that are ubiquitous in task planning. CPOptimizer may however struggle to prove unsatisfiability or escaping dead-ends in its search space in situations that would be easily handled by the conflict-directed clause learning of SAT/SMT solvers.

While our use of existing solvers clearly showed limitations, FAPE provided an alternate model. FAPE uses best-search in the space of lifted partial plans, with each partial plan encoded as a CSP that is iteratively refined and constrained as search progresses. This CSP was handled with a

custom propagation engine developed within the solver code base. Our experience is that developing side-by-side a propagation engine and the planner that uses it, allowed us to identify and address fundamental limitations of the propagation engines. For FAPE, this resulted in very strong propagation engine and compact search space (Bit-Monnot et al. 2020).

ARIES Scheduler

Following a similar route, ARIES relies on its own combinatorial solver that could be captioned as a *hybrid CP-SAT* solver for *optional scheduling*. Its most distinctive features are the following:

- The core of a SAT solver, (clause learning, unit propagation and activity-based search) integrated with finite-domain constraint solver.
- Dedicated reasoner for difference logic (aka. Disjunctive Temporal Networks), that notably enables conflict detection and explanations for temporal constraints.
- First-hand support of optional activities, enabling eager propagation of the potential timing and parameters of activities that are not yet part of the plan

While the solver is still in its early days, preliminary evaluation shows that it has state-of-the-art performance on disjunctive scheduling problems such as the jobshop and open-shop scheduling problems (Bit-Monnot 2023).

Technical Remarks

ARIES is implemented from the ground up in Rust and is freely available under the MIT license at <https://github.com/plaans/aries>. Most of the code base is dedicated to the implementation of our hybrid CP-SAT solver. Comparatively, planning specific code only occupies a small fraction of the code base, mostly dedicated to straightforward parsing and encoding tasks.

Beside its partial support for HDDL, ARIES also provides an integration as a backend for the *unified-planning* library with richer modeling features. Direct encoding of planning problems as chronicles or as a CSP is another way to exploit ARIES in a more flexible way.

Acknowledgments

This work has been partially supported by AIPlan4EU, a project funded by EU Horizon 2020 research and innovation program under GA n.101016442.

References

Bit-Monnot, A. 2018. A Constraint-Based Encoding for Domain-Independent Temporal Planning. In *International Conference on Principles and Practice of Constraint Programming (CP)*.

Bit-Monnot, A. 2023. Enhancing Hybrid CP-SAT Search for Disjunctive Scheduling. In *Under review at ECAI*.

Bit-Monnot, A.; Ghallab, M.; Ingrand, F.; and Smith, D. E. 2020. FAPE: a Constraint-based Planner for Generative and Hierarchical Temporal Planning. arXiv:2010.13121.

De Moura, L.; and Bjørner, N. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*.

Ghallab, M.; and Laruelle, H. 1994. Representation and Control in IxTeT, a Temporal Planner. In *International Conference on Artificial Intelligence Planning and Scheduling (AIPS)*.

Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*.

Godet, R.; and Bit-Monnot, A. 2022. Chronicles for Representing Hierarchical Planning Problems with Time. In *ICAPS Workshop on Hierarchical Planning (HPlan)*.

Laborie, P.; Rogerie, J.; Shaw, P.; and Vilím, P. 2018. IBM ILOG CP Optimizer for Scheduling. *Constraints*.

Schreiber, D. 2021. Lilotane: A Lifted SAT-based Approach to Hierarchical Planning. *Journal of Artificial Intelligence Research (JAIR)*.