Conformant-FF

Jörg Hoffmann
Max Planck Institute for CS
Saarbrücken, Germany
hoffmann@mpi-sb.mpg.de

The Conformant-FF planner, as entered into (the conformant track of) IPC-5, is exactly the system as described by (Brafman & Hoffmann 2004; Hoffmann & Brafman 2006). The only differences amount to a slightly changed input syntax, as used in IPC-5. The planner is an extension of the FF system (Hoffmann & Nebel 2001). The main trick in the extension is the use of SAT-based techniques to reason about uncertainty.

Conformant-FF, in its current implementation, can deal with initial state sets given as CNF formulas, actions with conditional effects, full ADL in effect conditions, and a subset of ADL in action preconditions and the goal formula. Constructs not supported are disjunction in preconditions and the goal (unless the disjunction disappears when processing static facts), and, most importantly, nondeterministic action effects.

Conformant-FF performs a search in belief space, as suggested first with the GPT system (Bonet & Geffner 2000), and also done in the various versions of the MBP system (Cimatti & Roveri 2000; Bertoli & Cimatti 2002). A belief state is the set of world states that are possible at some time point. The belief space is the space of all belief states reachable from the initial belief state. There are two key differences between the three systems: (1) their *representation* of belief states; (2) the *heuristic* used to guide the search.

GPT represents the belief states explicitly, enumerating the respective world states. Standard heuristic functions (from deterministic planning) can then be aggregated appropriately. MBP represents the belief states symbolically, i.e., each belief state is now a BDD. The heuristic simply prefers belief states with less uncertainty, i.e., BDDs that represent a smaller state set. Conformant-FF uses a very lazy representation of belief states, including only a *partial* knowledge: for a belief state B, it just computes the facts $T(B)$ that are true in *all* world states $s \in B$. This knowledge suffices to do STRIPS-style conformant planning: an action precondition *pre* (a conjunction of facts) is satisfied in B iff $pre \subseteq T(B)$; the goal G (a conjunction of facts) is satisfied in B iff $G \subseteq T(B)$. The facts $T(B)$ are computed by encoding the semantics of the action sequence leading to B as a "time-stamped" CNF formula ϕ, defining how fact values change over the action sequence, in a straightforward way. Conjoining ϕ with the initial state formula ϕ^I, one gets that $p \in T(B)$ iff $\phi^I \wedge \phi \models p$.

Conformant-FF's lazy or "implicit" representation can be seen as a way of trading space for time: on the positive side, we do not need to keep full detail about each B in memory; on the negative side, not having full detail about B forces us to reason all the way back to the initial state (in building the formulas ϕ) when computing the successors to B. In practice, we found even naive SAT solvers to be extremely efficient in solving the formulas arising in this context, so that the runtime price to pay is, in most cases, low.

The probably more crucial novelty in Conformant-FF is its heuristic function. This is an extension of FF's "relaxed plan heuristic" to the conformant setting, i.e., to initial states given as CNF formulas. The underlying relaxation is still to ignore the delete lists. Relaxed planning, however, is still **co-NP**-hard when the initial "state" is a CNF. We get around this by making another relaxation: we ignore all but one of the (unknown) effect conditions of each effect. This corresponds to a 2-projection of the CNF formula that would encode the semantics of the relaxed actions. To obtain a polynomial worst-case behavior, one would also have to 2-project the initial state formula. We tried this, and found it to produce very bad heuristic values in many examples. So, instead, we keep the initial state formula unchanged, investing the effort to reason about it with a SAT solver. This produces good heuristic values in many cases, with a tolerable overhead since the initial state formula is typically neither overly large nor overly complicated.

In the traditional conformant benchmarks (Bombs, Ring, Cube, ...) Conformant-FF is sometimes competitive with GPT and MBP, sometimes outperformed vastly (particularly in Ring). In benchmarks created as classical benchmarks enhanced with uncertainty, however, Conformant-FF is many orders of magnitude superior to both GPT and (all variants of) MBP, since the heuristic function inherits, to a large extent, the quality of FF's heuristic function in the classical setting.

References

Bertoli, P., and Cimatti, A. 2002. Improving heuristics for planning as search in belief space. In Ghallab, M.; Hertzberg, J.; and Traverso, P., eds., *Proceedings of the 6th International Conference on Artificial Intelligence Planning and Scheduling (AIPS-02)*, 143–152. Toulouse, France: Morgan Kaufmann.

Bonet, B., and Geffner, H. 2000. Planning with incomplete information as heuristic search in belief space. In Chien, S.; Kambhampati, R.; and Knoblock, C., eds., *Proceedings of the 5th International Conference on Artificial Intelligence Planning Systems (AIPS-00)*, 52–61. Breckenridge, CO: AAAI Press, Menlo Park.

Brafman, R., and Hoffmann, J. 2004. Conformant planning via heuristic forward search: A new approach. In Koenig, S.; Zilberstein, S.; and Koehler, J., eds., *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS-04)*, 355–364. Whistler, Canada: Morgan Kaufmann.

Cimatti, A., and Roveri, M. 2000. Conformant planning via symbolic model checking. *JAIR* 13:305–338.

Hoffmann, J., and Brafman, R. 2006. Conformant planning via heuristic forward search: A new approach. *Artificial Intelligence* 170(6–7):507–541.

Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.