

Anchoring abstractions for near-optimal task and motion planning

William Vega-Brown and Nicholas Roy
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
Email: {wrvb,nickroy}@csail.mit.edu

Abstract—We describe anchored angelic A* (AAA*), an algorithm for efficient and asymptotically near-optimal planning for systems governed by piecewise-analytic differential constraints. This class of systems includes many task and motion planning domains, such as quasi-static manipulation planning. AAA* uses an abstract representation of the problem domain to extract upper and lower bounds on the cost of concrete motion plans. It then uses those bounds to guide the search for a plan while maintaining performance guarantees. We demonstrate the use of AAA* on a simulated continuous Sokoban problem and show that it accelerates planning by an order of magnitude relative to our prior work (which used purely metric heuristics).

We consider the problem of integrated task and motion planning through the lens of optimal motion planning under piecewise-analytic differential constraints. This formalism describes a broad family of task and motion planning problems, including navigation among movable obstacles and quasistatic manipulation planning.

As a motivating example, we consider a toy problem, illustrated in fig. 1. A circular robot must push K circular objects into specified goal regions. The pose of each object is its position in the plane; the configuration space for the problem is then $\mathbb{R}^{2(K+1)}$. We neglect inertia and assume first-order dynamics; the robot can move an object only by making contact and pushing, but it cannot pull an object or push multiple objects at once. This problem is the continuous generalization of the warehouse problem, also known as Sokoban.

Continuous Sokoban and similar problems combine the challenges of non-convex optimization and combinatorial search, as the robot must decide both the sequence of objects to manipulate and where to move each object. Many authors [3, 9, 10, 12] have studied similar problems; nearly all research effort has concentrated on satisficing planning, where the objective is to return any plan that satisfies some set of objectives. Our focus is instead on optimizing planning, where the objective is to return the best plan that satisfies an objective function like path length or execution time. The few authors who have attempted optimal planning have either focused on restricted cases [4, 14] or local optimality [11]. In contrast, our previous work [13] described a globally asymptotically optimal planning algorithm called FORGG.

The core idea of FORGG is to construct a set of random graphs on subsets of the configuration space—called *modes*—where the dynamics are analytic, which is a condition sufficient to ensure that a large enough random graph will contain

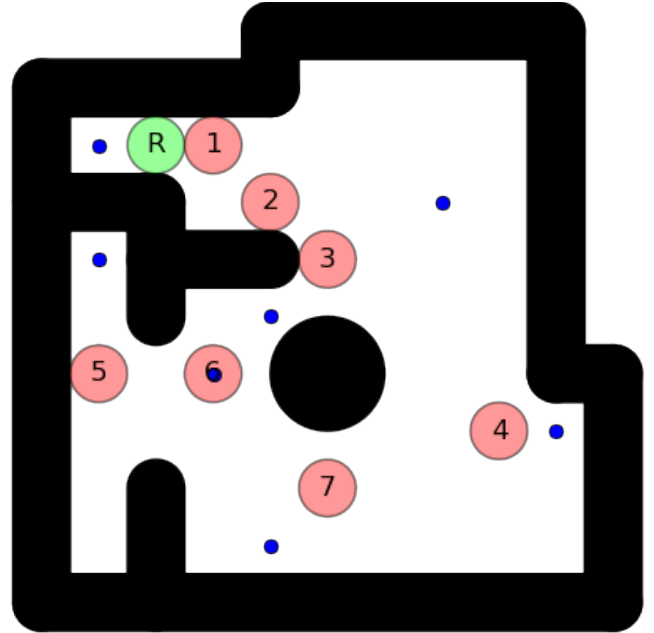


Fig. 1. An example of the continuous sokoban problem. The robot (in green, labeled 'R') must push the seven red, numbered objects to the goal locations marked with blue dots.

a near-optimal plan. We can then glue these graphs together into a single planning graph covering the entire configuration space. By taking advantage of the product structure of the differential constraints, we are able to build this graph efficiently, even for large problems. However, the FORGG approach still relies on a best-first search to find the best plan in the graph, and that search is computationally intractable for large problems.

In this abstract, we describe the anchored angelic A* (AAA*) search algorithm, which extends our prior work by using an abstraction to plan efficiently while preserving near-optimality guarantees. There is a rich body of work on using abstraction to accelerate search [2, 5, 6, 7] in the context of satisficing planning, but little attention has been paid to optimizing planning with abstraction.

Our approach combines the insights of the multi-heuristic A* (MHA*) algorithm developed by Aine et al. [1] and the angelic hierarchical planning framework developed by Marthi et al. [8]. The core idea of multi-heuristic A* is *anchoring*: by simultaneously performing searches with both admissible and

inadmissible heuristics, and sharing information between the searches, MHA* can accelerate planning while still ensuring that the plan returned will be near-optimal.

Although powerful, this idea is insufficient in task and motion planning problems like continuous Sokoban, as even crude heuristic estimates of the cost to reach a goal state are computationally expensive to compute. We therefore extend the MHA* algorithm to reason over *abstract plans*, such as “make contact with object 1 and push it to a goal state, then push object 2 to a goal state.” AAA* uses the *angelic semantics* introduced by Marthi et al. [8] to describe bounds on the cost of high-level actions.

In many domains, determining the feasibility or cost of an abstract action is computationally hard. Angelic semantics rely on cheaply computed *bounds* on the cost or feasibility of higher level actions, rather than requiring exact feasibility checks. These bounds take the place of an inadmissible heuristic in our algorithm; in effect, AAA* simultaneously and incrementally computes the feasibility and minimal cost of abstract plans. The only exact cost and feasibility computations required are for primitive, local actions such as the path of a robot between nearby configurations.

In general, the choice of abstraction and cost bounds are domain specific. We use a generic abstraction suitable for many object manipulation problem domains. We consider the set of *monotonic* plans, which are plans that touch each object at most once. We use the Euclidean length of each action in an abstract plan to lower-bound the cost of any concrete realization of the plan. In easy manipulation problems, when there are good monotonic plans, this allows for very fast planning. In harder problems, we can fall back on the admissible search to find an acceptable plan.

AAA* combines the graph construction of FORGG with an abstract search anchored by our admissible heuristic, and is guaranteed to return a $(1 + \epsilon)$ -optimal plan for any specified $\epsilon > 0$. We implemented AAA* in Python and performed qualitative comparisons with FORGG on several small instances of continuous Sokoban. In each case, we searched a graph with 100 poses per object, leading to graph sizes of approximately 10^4 , 10^6 , and 10^8 for problems with 1, 2, and 3 objects respectively. The parameters of AAA* were chosen to ensure a plan within a factor of 2 of the optimal plan. Results are presented in table I; we observe that AAA* is one to two orders of magnitude faster yet results in only a small decrease in solution quality. In particular, AAA* was able to solve problems with three or more instances, while FORGG failed to find a solution to a problem with more than two objects within twenty minutes of search time. These results are of course preliminary, and further evaluation is planned, both on larger problems and in additional domains.

AAA* has strong theoretical properties, and although we have yet to conduct a complete experimental study, early results indicate that it offers a significant speedup over FORGG, the only other asymptotically optimal planning algorithm of which we are aware. Because it provides a way to take advantage of an incomplete abstraction or an inadmissible

TABLE I
PERFORMANCE COMPARISON

Number of objects		1	2	3
Expansions	FORGG	237	3197	—
	AAA*	115	753	12125
Time	FORGG	3.30s	58.70s	1200
	AAA*	1.38s	2.55s	30.58s
Cost	FORGG	27.78	45.42	—
	AAA*	35.64	51.87	163.96

heuristic, AAA* represents an promising avenue for bringing the many years of research on discrete planning and decision making into a new space of continuous planning problems.

REFERENCES

- [1] Sandip Aine, Siddharth Swaminathan, Venkatraman Narayanan, Victor Hwang, and Maxim Likhachev. Multi-heuristic A*. *IJRR*, 35(1-3), 2016.
- [2] Esra Erdem, Kadir Haspalamutgil, Can Palaz, Volkan Patoglu, and Tansel Uras. Combining high-level causal reasoning with low-level geometric reasoning and motion planning for robotic manipulation. In *ICRA*, 2011.
- [3] Fabien Gravot, Stephane Cambon, and Rachid Alami. asymov: a planner that deals with intricate symbolic and geometric problems. In *ISRR*, 2005.
- [4] Kris Hauser. Minimum constraint displacement motion planning. In *RSS*, 2013.
- [5] Kris Hauser and Victor Ng-Thow-Hing. Randomized multimodal motion planning for a humanoid robot manipulation task. *IJRR*, 30(6):678–698, 2011.
- [6] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Hierarchical task and motion planning in the now. In *ICRA*, 2011.
- [7] Fabien Lagriffoul, Dimitar Dimitrov, Julien Bidot, Alessandro Saffiotti, and Lars Karlsson. Efficiently combining task and motion planning using geometric constraints. *IJRR*, 2014.
- [8] Bhaskara Marthi, Stuart Russell, and Jason Wolfe. Angelic hierarchical planning: Optimal and online algorithms. In *ICAPS*, 2008.
- [9] Dennis Nieuwenhuisen, A Frank van der Stappen, and Mark H Overmars. An effective framework for path planning amidst movable obstacles. In *WAFR*. 2008.
- [10] Mike Stilman and James Kuffner. Planning among movable obstacles with artificial constraints. *IJRR*, 27(11-12):1295–1307, 2008.
- [11] Marc Toussaint. Logic-geometric programming: An optimization based approach to combined task and motion planning. In *IJCAI*, 2015.
- [12] Jur Van Den Berg, Mike Stilman, James Kuffner, Ming Lin, and Dinesh Manocha. Path planning among movable obstacles: a probabilistically complete approach. In *WAFR*. 2009.
- [13] William Vega-Brown and Nicholas Roy. Asymptotically optimal planning under piecewise-analytic constraints. In *WAFR*, 2016.
- [14] Jason Wolfe, Bhaskara Marthi, and Stuart Russell. Combined task and motion planning for mobile manipulation. In *ICAPS*, 2010.