

Configuration Lattices for Planar Contact Manipulation Under Uncertainty

Michael C. Koval*, David Hsu†, Nancy S. Pollard*, and Siddhartha S. Srinivasa*
mkoval@cs.cmu.edu, dyhsu@comp.nus.edu.sg, {nsp, sidhd}@cs.cmu.edu

*The Robotics Institute, Carnegie Mellon University

†Department of Computer Science, National University of Singapore

I. INTRODUCTION

Our goal is to enable robots to use real-time feedback from contact sensors to reliably manipulate their environment. Contact sensing is an ideal source of feedback for a manipulator because contact is intimately linked with manipulation: the sense of touch is unaffected by occlusion and directly observes the forces that the robot imparts on its environment.

However, contact sensors suffer from a key limitation: they provide rich information while in contact with an object, but little information otherwise. Fully utilizing contact sensors requires *active sensing*. Consider the robot, shown in fig. 1, that is trying to push a bottle into its palm. The robot is uncertain about the initial pose of the bottle, has only an approximate model of physics, and receives feedback from noisy contact sensors on its fingertips. The robot executes an *information-gathering action* by moving its hand laterally to force the bottle into one of its contact sensors. Once the bottle is localized, it moves to achieve the goal.

In this paper, our goal is to autonomously generate policies that allow a robot to use real-time feedback from contact sensors to reliably manipulate objects under uncertainty. We specifically consider the problem of quasistatic [7] tabletop manipulation via pushing. The robot’s goal is to push a movable object into a hand-relative goal region while avoiding collision with stationary obstacles in the environment.

II. METHOD

We formulate the *planar contact manipulation problem* as a *partially observable Markov decision process* (POMDP) [3, 10]. State $s = (q, x) \in S$ is the configuration of the robot $q \in Q$ and pose of the movable object relative to the hand $x \in X$. An *action* $a \in A$ is a position-controlled motion of the robot. After taking an action, the robot receives an *observation* $o \in \{0, 1\}^{n_o} = O$ from its n_o binary contact sensors.

We assume that q is observable and deterministic, but x is latent and evolves according to a stochastic transition model $p(s'|s, a)$ that models the quasistatic motion of the object when pushed. Similarly, we assume that observations are generated by a stochastic observation model $p(o|s, a)$. We encode our goal to push the movable object into the hand-relative goal region $G \subseteq X$ into a reward function that assigns $R(s, a) = 0$ for $s \in G$ and $R(s, a) = -1$ otherwise.

Prior work has successfully used SARSOP [6], an offline POMDP solver, to find robust policies for manipulating an object relative to the hand [2, 5]. These techniques plan for a *hand-relative POMDP* (Rel-POMDP) that neglect q from the state space. As a result, these policies can perform poorly on

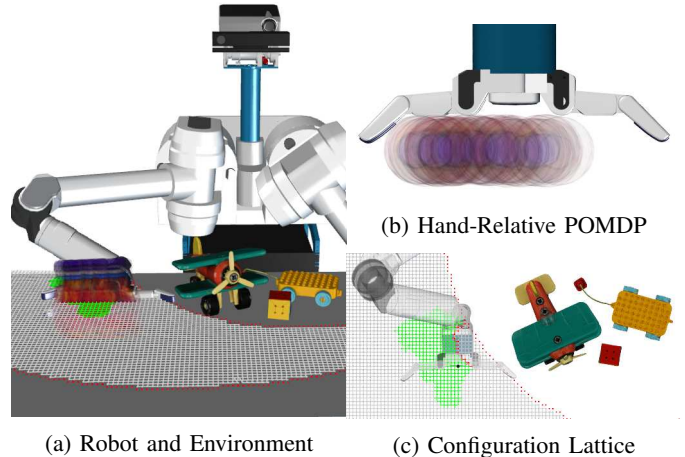


Fig. 1: HERB [12] (a) uses real-time feedback from contact sensors to manipulate a bottle on a cluttered table. (b) Our approach uses policies from the obstacle free, hand-relative contact manipulation problem to a search in the full state space. (c) The configuration of the robot is represented as a point in a state-space lattice consisting of feasible (—), infeasible (—), and unevaluated (—) edges; best viewed in color. The probability distribution over the pose of the bottle is shown as a collection of semi-transparent renders.

robotic manipulators due to reachability limits, collision, and other *kinematic constraints*. In contrast, our approach plans in the joint configuration space $Q \times X$ to consider kinematic constraints during planning.

A. Configuration Lattice POMDP

This planning problem is challenging because the configuration space Q and action space A are continuous and high dimensional, typically at least six dimensions. We address this challenge by planning with a finite set of *action templates*.

An *action template* $a_{\text{lat}} \in A_{\text{lat}}$ defines a Cartesian trajectory of the end-effector relative to its current pose. Given an action template, we use a *local planner* $\phi(q, a_{\text{lat}}) \mapsto \xi$ to search for a trajectory ξ in the robot’s configuration space that instantiates a_{lat} in configuration q . If ϕ succeeds, a_{lat} is feasible and the robot may execute ξ . Otherwise, it is deemed *infeasible* and the robot is penalized by receiving -1 reward.

Given an initial configuration $q_0 \in Q$, the set of all configurations reachable by instantiating A_{lattice} forms a discrete *configuration space lattice* $Q_{\text{lat}} \subseteq Q$. The robot’s configuration corresponds to a lattice point $q \in Q_{\text{lat}}$ and an action template $a_{\text{lat}} \in A_{\text{lat}}$ corresponds to an edge between two points. Note that the connectivity of the lattice—and the feasibility of its

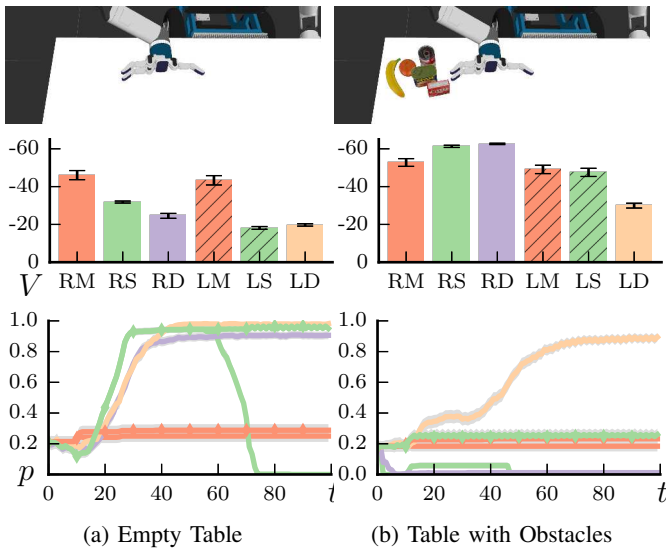


Fig. 2: Performance of the Rel-QMDP (RM \square \rightarrow), Rel-SARSOP (RS \square \rightarrow), Rel-DESPOT (RD \square \rightarrow), Lift-QMDP (LM \square \rightarrow), Lift-SARSOP (LS \square \rightarrow), and Lat-DESPOT (LD \square \rightarrow) policies on Lat-POMDP on two environments. (Top) Value V achieved by each policy after 100 timesteps. Note that the y -axis is inverted; lower (less negative) is better. (Bottom) Probability $p = \Pr(s_t \in G)$ that the movable object is in the goal region at each timestep. All results are averaged over problem 500 instances and error bars denote a 95% confidence interval. This figure is best viewed in color.

edges—may be lazily constructed as the planner explores the robot’s configuration space. This is illustrated in fig. 1.

B. Heuristic-Guided Online Search

It is challenging to plan for this *configuration lattice POMDP* (Lat-POMDP) with an offline method, e.g. SARSOP [6], because the optimal policy may dramatically change when an obstacle is added to, removed from, or moved within the environment. Instead, we use DESPOT—a state-of-the-art online POMDP solver—to find a solution [11].

DESPOT’s performance depends heavily on the quality of the upper and lower bounds used to guide its search. Through construction, the optimal value function of Rel-POMDP is an upper bound on that of Lat-POMDP. As a result, we can use any upper bound for Rel-POMDP as an upper bound on Lat-POMDP. We derive a lower bound by performing rollouts of a policy that is greedy with respect to that value function [9, 11].

The key advantage of using heuristics from Rel-POMDP is that they do not depend on the configuration of the environment and, thus, remain unchanged across problem instances.

III. EXPERIMENTAL RESULTS

We validate the efficacy of the proposed algorithm on a simulation of HERB [12], a mobile manipulator equipped with a 7-DOF Barrett WAM arm [8] and the BarrettHand [13] end-effector, manipulating the bottle shown in fig. 1. The robot begins in a known initial configuration q_0 with the initial pose of the bottle x_0 drawn from a Gaussian distribution centered in front of the palm with variance $\Sigma^{1/2} = \text{diag}[5 \text{ mm}, 10 \text{ cm}]$. At each timestep, the robot chooses one of four action templates A_{lat} that move 1 cm at a constant Cartesian velocity in the

xy -plane, receives an observation o from its $n_o = 2$ fingertip contact sensors, and receives a reward.

The motion of the object is simulated by injecting noise into the Box2D physics simulator [1]. We simulate contact sensors on HERB’s fingertips that perfectly discriminate between contact and no-contact [4, 5], but provide no additional information about where contact occurred along the sensor.

We compute the kinematic structure of the lattice by selecting a configuration $q_{\text{lat}}(x_r) \in Q_{\text{lat}}$ for each lattice point using an iterative inverse kinematics solver. Next, we use a Cartesian motion planner to find a trajectory that connects adjacent points while satisfying the constraints imposed by the action templates. We defer collision checking until DESPOT queries the feasibility of an action template online.

Figure 2 shows the performance of our approach, Lat-DESPOT (LD \square \rightarrow), compared to that of five baseline algorithms averaged over 500 simulation experiments. Rel-QMDP (RM \square \rightarrow) is a baseline that does not take multi-step information-gathering actions, Rel-SARSOP (RS \square \rightarrow) is a policy computed by running SARSOP for 300 s offline, and Rel-DESPOT (RD \square \rightarrow) is a policy found by running DESPOT on Rel-POMDP. All three of these algorithms ignore kinematic constraints during planning. Lift-QMDP (LM \square \rightarrow) and Lift-SARSOP (LS \square \rightarrow) heuristically lift the Rel-QMDP and Rel-SARSOP policies into configuration space by myopically executing the feasible action with highest value.

Our results show that the SARSOP and DESPOT policies all outperform the QMDP policies on the empty table (fig. 2a) in both value and success rate. This confirms our hypothesis that active information gathering is important when using contact sensors. Our results also show that Rel-QMDP, Rel-SARSOP, and Rel-DESPOT perform poorly on cluttered scenes (fig. 2b). This occurs because these policies do not consider kinematic constraints and quickly collide with the environment. *Our approach, Lat-DESPOT, performs best across both scenes because it considers both kinematic constraints and object pose uncertainty during planning.*

IV. CONCLUSION

We formulated planar contact manipulation under uncertainty as a POMDP in the joint space of robot configurations and poses of the movable object. We simplify the problem by constructing a lattice in the robot’s configuration space that captures kinematic constraints and guides DESPOT, a state-of-the-art online POMDP solver [11], towards an optimal solution using heuristics derived from a relaxation of the problem. Our results demonstrate that Lat-POMDP outperforms baselines that do not consider kinematic constraints during planning.

In future work, we plan to scale Lat-POMDP up to larger problems, e.g. more action templates and non-planar motion, and evaluate its policies on a real-robot.

This work was supported by a NASA Space Technology Research Fellowship (award NNX13AL62H), the National Science Foundation (awards IIS-1218182 and IIS-1409003), the U.S. Office of Naval Research, and the Toyota Motor Corporation.

REFERENCES

- [1] E. Catto. Box2D. <http://box2d.org>, 2010.
- [2] M. Horowitz and J. Burdick. Interactive non-prehensile manipulation for grasping via POMDPs. In *IEEE International Conference on Robotics and Automation*, 2013. doi: 10.1109/ICRA.2013.6631031.
- [3] L.P. Kaelbling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 1998. doi: 10.1016/S0004-3702(98)00023-X.
- [4] M.C. Koval, N.S. Pollard, and S.S. Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *International Journal of Robotics Research*, 34(7):922–945, 2015. doi: 10.1177/0278364915571007.
- [5] M.C. Koval, N.S. Pollard, and S.S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. *International Journal of Robotics Research*, 2015. doi: 10.1177/0278364915594474. In press.
- [6] H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems*, 2008.
- [7] K.M. Lynch, H. Maekawa, and K. Tanie. Manipulation and active sensing by pushing using tactile feedback. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1992. doi: 10.1109/IROS.1992.587370.
- [8] K. Salisbury, W. Townsend, B. Eberman, and D. DiPietro. Preliminary design of a whole-arm manipulation system (WAMS). In *IEEE International Conference on Robotics and Automation*, 1988. doi: 10.1109/ROBOT.1988.12057.
- [9] D. Silver and J. Veness. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*, 2010.
- [10] R.D. Smallwood and E.J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21(5):1071–1088, 1973. doi: 10.1287/opre.21.5.1071.
- [11] A. Somani, N. Ye, D. Hsu, and W.S. Lee. DESPOT: Online POMDP planning with regularization. In *Advances in Neural Information Processing Systems*, 2013.
- [12] S.S. Srinivasa, D. Berenson, M. Cakmak, A. Collet, M.R. Dogar, A.D. Dragan, R.A. Knepper, T. Niemueller, K. Strabala, and M. Vande Weghe. HERB 2.0: Lessons learned from developing a mobile manipulator for the home. *Proceedings of the IEEE*, 100(8):1–19, 2012.
- [13] W. Townsend. The BarrettHand grasper—programmably flexible part handling and assembly. *Industrial Robot: An International Journal*, 27(3):181–188, 2000. doi: 10.1108/01439910010371597.