Exploiting Contact for Efficient Motion Planning Under Uncertainty

Arne Sieverling   Clemens Eppner   Oliver Brock

***Abstract*—In this paper we want to argue for contact as an enabler for efficient motion planning under uncertainty. Controlled and desired contact can project a high dimensional belief state to a lower-dimensional manifold. A robot can sequence these projections to reduce uncertainty about its state. For realistic applications, these uncertainty reducing actions must be sequenced with uncertainty-increasing free space motion. We present a sampling-based motion planner that searches a belief state over configurations augmented with contact information. The planner finds robust contact-exploiting policies under significant uncertainty in robot and world model. We validate these policies on a seven-DoF robot manipulator in simulation and real world experiments.**

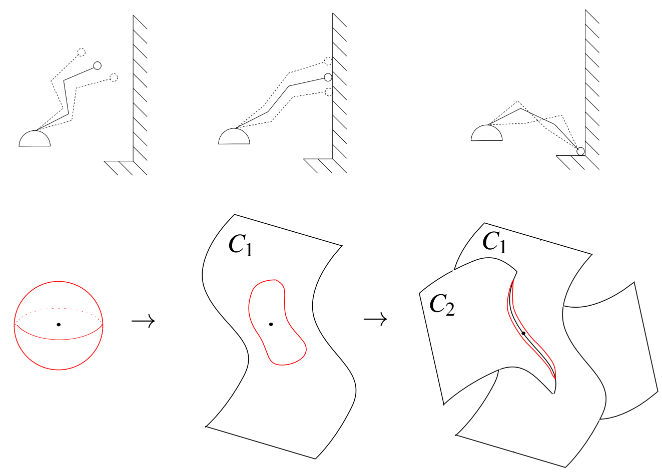

Fig. 1. Contact reduces uncertainty by projecting a high dimensional belief state onto a lower-dimensional contact manifold. Shown above is a 3-DOF manipulator in three different configurations. There is uncertainty in the joint positions which is sketched with dashed lines. Below is a sketch of the configuration space for these three scenes. Shown in red is the distribution of configurations. C_1 and C_2 are the manifolds that describe all configurations in contact with the right or lower wall. In the first step, the robot moves into contact with the right wall, which projects its initial uncertainty onto C_1. In the second step, the robot establishes contact with the lower wall, which projects the distribution on intersection of C_1 and C_2.

I. Introduction

Most state-of-the-art motion planning methods, including sampling-based methods and optimization approaches, are designed solely for contact avoidance. While solutions of these planners are applicable in highly controlled settings, they often fail in unstructured environments due to a high amount of uncertainty. To overcome these limitations, planners include uncertainty into the models of world and robot. The planners then plan in belief space: the space of probability distributions over robot and world state. Belief space planning allows to find contact-free paths, even if the outcome of robot actions or the environment are not known completely. However, planning under uncertainty is a hard problem. If the robot does not have access to its full state but must estimate it using uncertain sensors, the planning problem becomes a POMDP which is intractable to solve completely and hard to approximate for realistic problem sizes.

We will show that an explicit reasoning about contact allows for a tractable solution to the motion planning problem under uncertainty that scales to realistic manipulation tasks in high-dimensional configuration space. The key to the efficiency is the assumption that contact sensing is uncertainty-free. From this assumption follows that measurable contact eliminates uncertainty completely in one dimension. From a belief space planning perspective this is equivalent to a projection of the n-dimensional belief state to a $(n-1)$-dimensional manifold (see Fig. 1). In the ideal case, these projections can be chained until the belief collapses into a single point and uncertainty is reduced completely. This chaining of compliant contact actions is the basis for many real-world manipulation planners such as fine motion planning [12], sensorless manipulation [4], or submodular tactile localization [6].

However, in reality every motion that is not constrained by contact will also increase uncertainty. More specifically, whenever moving, uncertainty increases in those dimensions that are orthogonal to the contact normal. In this paper we propose the *Contact-Exploiting Rapidly Exploring Random Tree (CERRT)*, a manipulation planning algorithm that sequences uncertainty reduction with free-space motion. To combine these two objectives the planner searches in a combined space of contact and configuration. In this work we will show that this search is feasible using sampling-based motion planning on a particle-based state representation.

This paper summarizes our findings presented in a manuscript appearing later this year [16].

II. Related work

Planning free-space motion and planning contact are two well-established research areas and we will briefly outline our planner's connection to related work in both fields. Our planner balances free-space motion and contact by reasoning

All authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany. We gratefully acknowledge the funding provided by the European Commission (EC, SOMA, H2020-ICT-645599). We thank Felix Wolff for significantly contributing to the planner prototype.

about uncertainty, for which we will review related work in the second half of this section.

Free-space motion: Sampling-based motion planners like the RRT [11] search the collision-free configuration space efficiently but assume no uncertainty and explicitly avoid contact. Sampling-based motion planning can explore the space of configurations in contact [7, 17, 3] but does not reason about the uncertainty reducing capability.

Planning with uncertain actions: Markov Decision Processes (MDP) model actions with uncertain outcome. This framework allows robots to reason about the collision probability of actions and balance short and safe paths [1]. Particle-RRTs [13, 14] represent the outcome of actions as a set of particles, just like our planner. However, the particle-RRT assumes perfect knowledge about robot state which CERRT does not and our method explicitly seeks contact to reduce uncertainty, while the particle-RRT just achieves contact randomly.

POMDPs for manipulation: Once uncertainty exists in action outcome and the robot can not fully observe its own state, the planning problem is a Partially-Observable MDP (POMDP). Sampling-based POMDP solvers [10, 19, 15] were applied to low-dimensional versions of manipulation tasks such as in-hand manipulation to localize an object [8, 9] or pre-grasp manipulation [5, 2]. We will show the uncertainty-reducing capabilities of our planner on the latter application, solving the same problem in Section IV, but with the difference that our method does not assume any a priori discretization of state or action space.

III. Contact-Exploiting RRT (CERRT)

Our planner uses a combined state of belief over configuration and fully-observable contact. When planning with uncertainty and contact, belief states are shaped by the projections on contact manifolds and thus usually non-Gaussian (see Fig. 1). Therefore, we represent the belief over the configuration with a set of N particles. Thus our planner's state is given by $x = \{q_1, \ldots q_N\}$, where each q_i is a robot configuration vector.

The planner finds strategies that combine free-space and contact motion. CERRT assumes that free-space motion always increases uncertainty and therefore must be sequenced with contact motions that reduce uncertainty. Fig. 2 shows an example of a decision the planner must take. The robot can not directly enter the narrow passage but must first contact the wall to reduce uncertainty.

To find such strategies, we grow a tree in the combined space of contact state and belief over configuration. This search is based on a rapidly exploring tree planner, modified with actions that seek contact and slide along surfaces. Particle-based RRT-search has been used before for mobile robot navigation [13], and manipulation [14], although only for uncertain action outcome and not for uncertain robot state.

The basic structure of the CERRT matches the RRT. In each iteration, the planner draws a random sample (a), finds the nearest node in the tree (b), chooses an action that moves the node towards the random sample (c), and simulates the motion towards the sample (d). In the following we will detail how the CERRT implements these four steps by exploiting contact to reason efficiently about uncertainty.

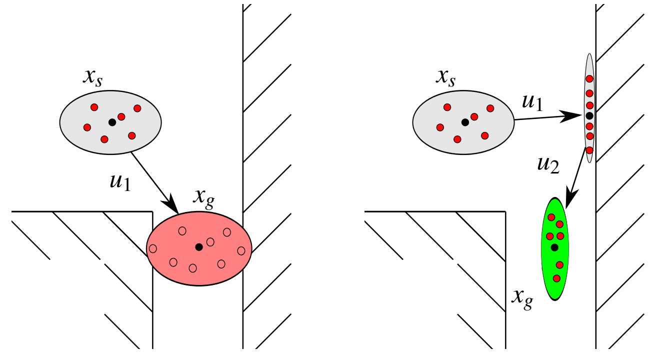

(a) A direct move leads to collision due to high uncertainty. (b) Contact reduces uncertainty enough to fit in the narrow passage.

Fig. 2. Example for the belief space planning problem with contact in 2D. (a) To enter the narrow passage, the robot cannot directly take action u_1 because the resulting uncertainty would lead to collision. (b) By sequencing a contact move u_1 and a free space move u_2, the robot reduces position uncertainty sufficiently to enter the narrow passage.

a) Sampling: The expansion direction for the CERRT is determined in every iteration by a random sample q_{rand} drawn uniformly from C-space. To bias the search towards a goal state, the goal itself is picked 10% of the time.

b) Node selection: Like the RRT, our planner selects the next node to extend x_{near} with minimal distance to a random sample q_{rand}. The distance function is a weighted sum of the Euclidian distance to the mean of x_{near} and the norm of the covariance matrix of x_{near}.

c) Action selection: To explicitly seek contact, the CERRT employs three different action types:

Connect: attempts a straight line connection in configuration space to the sample q_{rand}. *Connect* explores the free space and usually increases position uncertainty (Fig. 3(a)).

Guarded: implements a guarded move. It moves the robot in the direction of q_{rand} until contact with with the environment is established. *Guarded* always ends in collision and therefore eliminates uncertainty in one dimension (Fig. 3(b)).

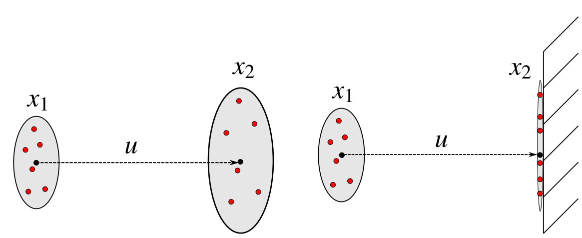

(a) Connect (b) Guarded move

Fig. 3. A free-space move and a move into contact. x_1 and x_2 are the initial and final particle distributions before and after applying action u.

Slide: moves along a surface until the contact state changes, either by colliding with another environmental structure (Fig. 4(a)) or by leaving the sliding surface (Fig. 4(b)). A slide always goes into the direction of q_{rand}, projected onto the sliding surface. Sliding is implemented using a projection method based on the Pseudoinverse of the Jacobian [18]. Thus each slide action moves the robot on the contact manifold until a second contact is reached, which eliminates uncertainty in a second dimension.

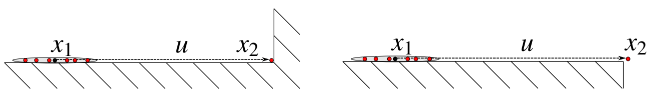

(a) Concave slide (b) Convex slide

Fig. 4. Two sliding actions. (a) the slide moves the distribution along the surface until it achieves contact with another surface. (b) the slides moves until it loses contact with the surface. Both slides reduce uncertainty in two dimensions.

d) Forward simulation: To extend a node x_1, the CERRT samples a particle from $\mathcal{Q}_{x_1}$ and then simulates noisy executions of the action u. The outcome of these executions are added to the new state x_2 if the particle set is in a consistent and measurable contact state.

After inserting a valid node, the planner tries to reach the goal state from the newly inserted node, also using forward simulation. If the resulting distribution is close to the desired goal distribution, the planner returns success. Otherwise it moves to the next iteration and picks another sample.

IV. Experiments

We will now show the capabilities of our planner on manipulation tasks.

2D grasping: This problem models a gripper picking up a square block at unknown location and is inspired by the POMDP literature [5, 2]. The gripper has contact sensors at each jaw and can translate in two dimensions. Because of a large initial uncertainty the gripper must contact the object or the walls first and then, after uncertainty is sufficiently reduced, attempt the grasp from the top.

Fig. 5(a) shows one of the solution paths CERRT found on the simple grasping scenario. All policies first establish contact with wall or object and then slide along the ground until contact with the object is perceived. The planning time for this problem averaged over ten runs is 6.8 s ($\pm$5.1s), while a POMDP version of the problem required an average planning time of 160s [2]. Our approach easily scales to more complex scenarios. Fig. 5(b) shows the result for a multi-step piece (8.2s$\pm$6.9s), Fig. 5(c) a version where the gripper must first navigate through a simple maze (23.4s$\pm$19.3s), Fig. 5(d) a 3D version of the problem with translation and rotation of the gripper.

7D robot arm motion: CERRT can directly be applied to the configuration space of a robot manipulator. In this Section we will show policies generated by CERRT that use contact to reduce uncertainty but also avoid collisions with links that have no contact sensing ability.

We place a 7-DOF Barrett WAM robot in front of the wall depicted in Fig. 6, similar to the scenario from Phillips-Grafflin et al. [14]. The robot model has an initial uncertainty about its own configuration with a standard deviation of $\sigma_{\text{start}} = 0.02$ rad. It also has a motion dependent position uncertainty of $\sigma_\delta = 0.02$. This term models that the robot will accumulate uncertainty of 0.02 rad if it moves a distance of 1 rad. Such a a motion-dependent position error occurs in the real Barrett WAM robot due to stretch of the cables that move the joints. The robot uses a wrist-mounted ATI Gamma force-torque sensor to perceive contact with the end-effector but cannot perceive contact with any other part.

The outcome of the planner can be seen in Fig. 6. From ten attempts, the planner solved this problem six times within 180 s. The six successful searches required an average time of 23.8 s$\pm$29.3 s. To validate the robustness of the plan, we introduce an unexpected disturbance. We raise the wall including all obstacles by 7 cm and execute the motion on the robot. The contact with the cyan and red boxes reduces uncertainty and the robot reaches the target with an error of 2 cm. This shows that the exploitation of contact allowed the robot to reduce position error by 5 cm.

V. Conclusion

This paper showed how to plan robust manipulation strategies under significant uncertainty in robot state, action, and world model by exploiting contact. We have shown that a simple kinematic model of contact can make belief space planning in high-dimensional continuous space tractable. The reason for the efficiency of contact is that measurable, uncertainty-free contact projects a high-dimensional state distribution to a lower-dimensional manifold. We presented a planner that sequences such projections with free-space motion to balance uncertainty reduction and progress towards a goal. Our results stand in contrast to a more traditional view on motion planning that sees contact as a problem due to difficult to model friction effects and discontinuities. In this paper we have shown that contact enables planning of simple and safe policies for manipulation problems under significant uncertainty.

References

[1] Ron Alterovitz, Thierry Siméon, and Kenneth Y Goldberg. The stochastic motion roadmap: A sampling framework for planning with markov motion uncertainty. In *Robotics: Science and systems*, volume 3, pages 233–241, 2007.

[2] Haoyu Bai, David Hsu, Wee Sun Lee, and Vien A Ngo. Monte carlo value iteration for continuous-state POMDPs. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 175–191. Springer, 2010.

[3] Nassime Blin, Michel Taïx, Philippe Fillatreau, and Jean-Yves Fourquet. I-rrt-c: Interactive motion planning with contact. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 4244–4249. IEEE, 2016.

[4] Michael A. Erdmann and Matthew T. Mason. An exploration of sensorless manipulation. *IEEE Journal on Robotics and Automation*, 4(4):369–379, 1988.

[5] Kaijen Hsiao, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Grasping POMDPs. In *International Conference on Robotics and Automation (ICRA)*, pages 4685–4692. IEEE, 2007.

[6] Shervin Javdani, Matthew Klingensmith, J Andrew Bagnell, Nancy S Pollard, and Siddhartha S Srinivasa. Efficient touch based localization

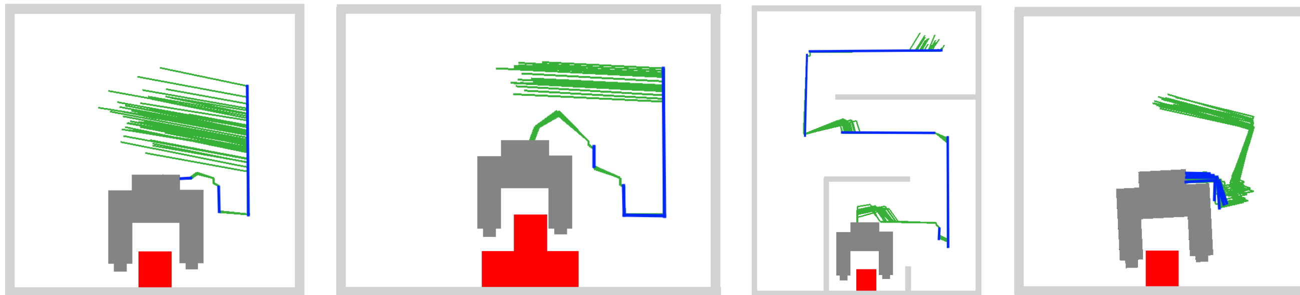

(a) A simple grasping experiment from the POMDP literature [5]

(b) Grasping a more complex shape. The gripper slides up the steps of the object.

(c) Solving a simple maze before grasping.

(d) Planning in $SE(2)$ (2D translation and gripper rotation)

Fig. 5. Solution strategies of CERRT for a grasping scenario. The gripper shows the final configuration of the path. The lines show 20 sampled trajectories, free-space motions are shown in green and slides in blue. The beginning of the paths is always in free space and the end is before grasping.

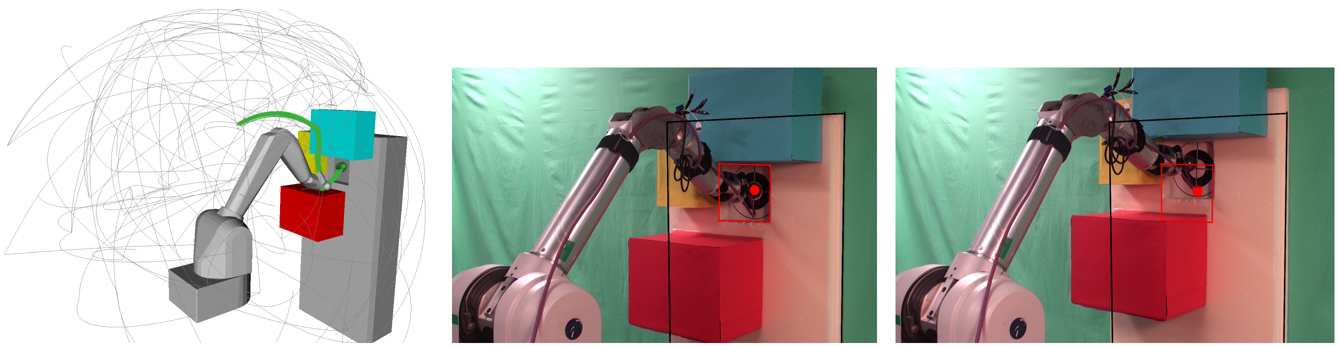

(a) The planner output. Gray lines show all explored motions. The green line is the found path.

(b) First execution

(c) Second Execution. The black overlay shows the wall position from (b)

Fig. 6. The manipulator must touch the target in the square opening of the wall. (a) Our planner finds a strategy that moves to the cyan box, slides down until it loses contact, does a guarded move to the top of the red box, and moves to the target. (b) We execute this strategy on the real robot and reach the target precisely. (c) We now raise the obstacles by 7 cm and execute the policy from (b) again. The robot uses the contact to reduce uncertainty and reaches the target with an error of 2 cm.

through submodularity. In *International Conference on Robotics and Automation (ICRA)*, pages 1828–1835. IEEE, 2013.

[7] Xuerong Ji and Jing Xiao. Planning motions compliant to complex contact states. *The International Journal of Robotics Research*, 20(6): 446–465, 2001.

[8] Michael Koval, David Hsu, Nancy Pollard , and Siddhartha Srinivasa. Configuration lattices for planar contact manipulation under uncertainty. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[9] Michael C Koval, Nancy S Pollard, and Siddhartha S Srinivasa. Pre-and post-contact policy decomposition for planar contact manipulation under uncertainty. *The International Journal of Robotics Research*, 35 (1-3):244–264, 2016.

[10] Hanna Kurniawati, David Hsu, and Wee Sun Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems*, 2008.

[11] Steven M. Lavalle. Rapidly-exploring random trees: A new tool for path planning. Technical report, Department of Computer Science, Iowa State University, 1998.

[12] Tomás Lozano-Pérez, Matthew T Mason, and Russell H Taylor. Automatic synthesis of fine-motion strategies for robots. *The International Journal of Robotics Research*, 3(1):3–24, 1984.

[13] Nik A Melchior and Reid Simmons. Particle RRT for path planning with uncertainty. In *International Conference on Robotics and Automation (ICRA)*, pages 1617–1624. IEEE, 2007.

[14] Calder Phillips-Grafflin and Dmitry Berenson. Planning and resilient execution of policies for manipulation in contact with actuation uncertainty. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[15] Josep M Porta, Nikos Vlassis, Matthijs TJ Spaan, and Pascal Poupart. Point-based value iteration for continuous POMDPs. *Journal of Machine Learning Research*, 7:2329–2367, 2006.

[16] Arne Sieverling, Clemens Eppner, and Oliver Brock. Interleaving motion in contact and in free space for planning under uncertainty. In *to appear in: International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, 2017.

[17] Thierry Siméon, Jean-Paul Laumond, Juan Cortés, and Anis Sahbani. Manipulation planning with probabilistic roadmaps. *The International Journal of Robotics Research*, 23(7-8):729–746, 2004.

[18] Mike Stilman. Task constrained motion planning in robot joint space. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 3074–3081. IEEE/RSJ, 2007.

[19] Sebastian Thrun. Monte carlo POMDPs. In *NIPS*, volume 12, pages 1064–1070, 1999.