Contact-Based Robotics Manipulation Tasks Learning and Execution

Shuai Li[1], Siwei Lyu[2] and Jeff Trinkle[3]

I. Introduction

In our daily lives, we constantly manipulate objects, for example, we open doors, slide books onto shelves, and fix cars. Some of these manipulation tasks are tedious or dangerous and it is more desirable to let robots help us to perform them. However, robotic systems are currently not capable of these kinds of tasks. An important reason for this is that manipulation involves creating contact arrangements that allow the robot to apply certain forces or impart specific motions of the object. If the manipulation moves the object over a long distance or requires significant reorientations, the contact arrangements might need to be modified by removing a few and creating new ones elsewhere, or even by sliding existing ones to new locations.

If we consider the robot and the object as nominally rigid bodies, a manipulation task can be decomposed into a set of finite motions punctuated by a discrete contact state transitions. Although the general dynamics model of rigid body dynamics system is highly nonlinear, if its contact state (the locations of the contacts and whether they are sticking or slipping) is known, then the motion dynamics of the object can be well defined as a linear model [1]. In this case, one can use Kalman filtering to track the object and Linear-Quadratic-Gaussian control to execute the manipulation plan. Therefore, to advance the manipulation capabilities of robotic systems, they should be able to sense the contact locations and slip directions directly (via advanced tactile sensors) or indirectly (through inference as done in [1]).

Many works have been done to incorporate contact information into robotics manipulations. Farahat et al. studied the contact constrains imposed on a system of rigid bodies [2]. They proved that the syetem of contact constraint equations are smooth submanifolds of configuration space. Chalon et al. developed a particle filter that tracks the pose of the object when it is grasped by a robot gripper [3]. Their method incorporates the kinematic constraints of the robot hand into the state transition model to update the pose of the target object. The work by Koval et al. developed the "manifold particle filter" to support their work in planar push-grasping [4]. Their manifold particle filter samples particles on a pre-computed contact manifold, which includes the configuration space of the target object in contact with the robot hand.

Although there are different approaches in the previous works to address the contacts in a manipulation task, they all only considered the contacts as kinematic constraints and failed to cover the dynamic effect of the contacts on the motions of the object. In this work, we designed a Rao-Blackwellised particle filter [5] with a multibody dynamics model at its core to estimate the object's contact states and poses. With the state space of a manipulation task discretized by the object's contact states, we implemented a model-based reinforcement learning algorithm that learns the transition model of the contact states as well as a policy for the manipulation task. Finally, we propose a framework that enables a robot to learn a manipulation task based on contact states and then perform the task with contact states estimated by our designed particle filter.

[1]Shuai Li is with the Dept. of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180, USA lis12@rpi.edu

[2]Siwei Lyu is with the Dept. of Computer Science, University at Albany, SUNY, Albany, NY 12222, USA lsw@cs.albany.edu

[3]Jeff Trinkle is with Faculty of the Dept. of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180, USA trink@cs.rpi.edu

II. Contact State Estimation

The dynamics model of the multibody dynamics system that includes of the robot and the object can be well described as a Linear Complementarity Problem (LCP) model [6]. The LCP model consists of Newton-Euler equations and complementarity conditions as shown in equations (1, 2). M is the mass-inertia matrix, $\dot{v}$ is the derivative of the generalized velocities, λ is vector of the constraint forces generated by contacts, G is corresponding Jacobian, λ_{ext} is the vector of external forces, and Ψ is the gap distance vector. The $\perp$ sign represents the complementarity condition, which indicates that the contact constraint forces can be nonzero if and only if the two objects are in contact.

$$M\dot{v} = G\lambda + \lambda_{ext} \tag{1}$$

$$0 \leq \lambda \perp \Psi \geq 0 \tag{2}$$

As shown in equations (1, 2), the complementarity conditions can be decoupled from the Newton-Euler equation once the contact state is known. In fact, LCP model can be thought of as a piecewise linear model with "switching points" at the contact states. We therefore designed a Rao-Blackwellized particle filter that samples the discrete contact states and propagates the distribution of the object's continuous states, e.g., poses of the object, through multiple Kalman filters, as shown in Fig. 1.

III. Reinforcement Learning With Contact State

As shown above, the motion dynamics model of an object is constrained by its contact states. We therefore propose to discretize the state space of robotics manipulation tasks with the object's contact states. However, the conventional definition of contact states only consider the contacts as kinematic constraints. We treat the contact states as constraints to decouple the complementarity conditions of the LCP model

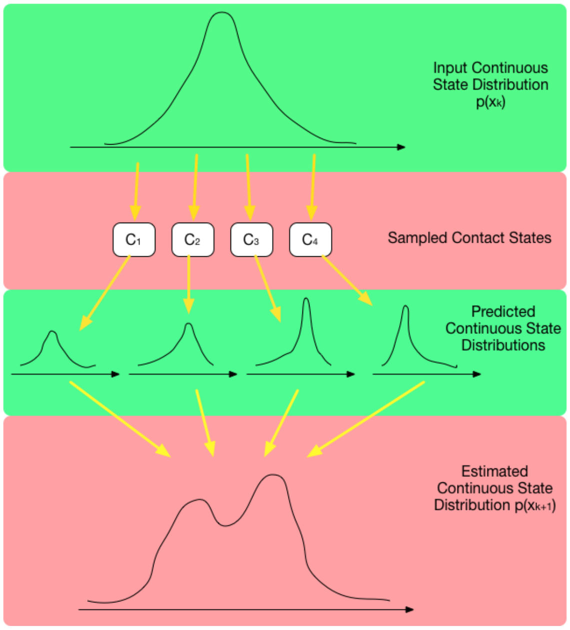

Fig. 1: A diagram of our designed particle filter.

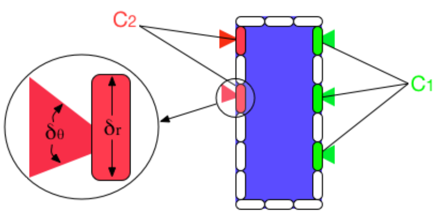

Fig. 2: An example of two contact states for a rectangular object. Each small rectangle represents a contact sensor, and the triangle on each contact sensor represents the range of angles in which the sensor can be triggered. c_1 is defined as a contact state that has all three green contact sensors triggered, and c_2 is a contact state with both red contact sensors triggered.

so that the motion model of the object is fully defined. We define our contact state as a set of contacts whose locations are within certain radius δ_r and contact normals are pointing to the same directions with rolerance δ_θ as shown in Fig. 2. Also, we include the sticking/sliding status in the contact states. To compensate the special contact state, which includes no contacts between the robot and the object, we extend the concept of contact state and define manipulation state m, which includes the relative poses between the robot and the object it manipulates, as well as the contact states of the object with the robot and the environment. Similar to the concept of contact graph [7], we also define the manipulation graph to describe the transition relations among manipulation states.

We implemented a Dyna-Q reinforcement learning algorithm [8] to learn policy of a robotics manipulation task as well as a manipulation graph.

IV. Contact Aware Manipulation

As the policy is learned based on the contact states of the object, in order to let the robot execute the policy in a manipulation task, we applied our designed particle filter to estimate the pose and contact states of the object. To sample the contact states, our particle filter will require a graph from which it can obtain the distributions of the contact states for the next time step given the current contact state. On the other hand, our reinforcement learning algorithm builds the manipulation graph as it learns the policy, and therefore it provides exactly the key piece for our filter.

In order to select the best action during the task execution, we adopt the QMDP method [9] to combine the state estimation from our filter with the policy learned. Specifically, in the particle filter, each particle outputs an estimation of the manipulation state m_i and the weight of its corresponding particle is w_i. If the Q value for the pair (m_i, a_j) is Q_{ij}, the best action a^* is calculated as in equation (3), where N is the number of particles.

$$a^* = \arg\max_{a_j} \sum_{i}^{N} w_i Q_{ij} \tag{3}$$

V. Experiments

We tested our approaches in both simulation and physical experiments. The configuration space of our experiments ranges from $SE(2)$ to $SE(3)$. In order to test our particle filter, we have conducted simple experiments such as a parallel jaw gripper grasping a triangle object, and also complex grasping experiments in $SE(3)$ such as a Barrett hand attempting grasps on a rectangular cube that can be tipped over. We tested our reinforcement learning algorithm in two-dimensional simulation peg-in-hole experiments, and we also attempted physical experiments with similar setups with our WAM arm and a Kinect as the vision sensor.

The results showed that our particle filter is able to track both the pose of the object and its contact states accurately during manipulation tasks [1]. We also verified that our framework enables robots to learn and execute simple peg-in-hole tasks.

VI. Conclusion and Future Work

In this work, we present a framework that incorporates the contact information into the robotic manipulations through a model of multibody dynamics systems. We designed a particle filter that estimates the object's poses as well as its contact states. A reinforcement learning algorithm was implemented to let a robot learn the policy for finishing a manipulation task based on the contact states. We showed that the reinforcement learning algorithm can provide the transition model for contact states, which will be used in our filtering method, and the poses and contact states estimation from our filter can also be used as feedback during the manipulation task.

In the future, we plan to apply our methods to more complex robotics manipulation tasks. We would also like to improve the efficiency of our particle filter and reinforcement learning algorithm.

VII. Acknowledgements

This work was partially supported by the National Science Foundation through the National Robotics Initiative (grants NRI-153702 and NRI-1537257) and the Independent Research and Development program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

[1] S. Li, S. Lyu, and J. Trinkle. State estimation for dynamic systems with intermittent contact. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3709–3715. IEEE, 2015.

[2] A. O. Farahat, P. F. Stiller, and J. C. Trinkle. On the algebraic geometry of contact formation cells for systems of polygons. In *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*, pages 3028–3033 vol.4, May 1994.

[3] M. Chalon, J. Reinecke, and M. Pfanne. Online in-hand object localization. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2977–2984, Nov 2013.

[4] Michael Koval, Nancy Pollard , and Siddhartha Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *International Journal of Robotics Research*, 34(7):922–945, June 2015.

[5] Arnaud Doucet, Nando De Freitas, Kevin Murphy, and Stuart Russell. Rao-blackwellised particle filtering for dynamic bayesian networks. In *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*, pages 176–183. Morgan Kaufmann Publishers Inc., 2000.

[6] D.E. Stewart and J.C. Trinkle. An Implicit Time-Stepping Scheme for Rigid Body Dynamics with Inelastic Collisions and Coulomb Friction. *International Journal of Numerical Methods in Engineering*, 39, 1996.

[7] J. Xiao and X. Ji. Automatic generation of high-level contact state space. *The International Journal of Robotics Research*, 20(7):584–606, 2001.

[8] Richard S. Sutton. Planning by incremental dynamic programming. In *The Eighth International Workshop on Machine Learning*, pages 353–357, 1991.

[9] Anthony R Cassandra and Leslie Pack Kaelbling. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995: Proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, California, July 9-12 1995*, page 362. Morgan Kaufmann, 2016.