

# Hybrid Differential Dynamic Programming for Planar Manipulation Primitives

Neel Doshi<sup>1,2</sup> Francois R. Hogan<sup>2</sup>, Alberto Rodriguez<sup>2</sup>  
<nddoshi, fhogan, albertor>@mit.edu

**Abstract**—We present a hybrid differential dynamic programming (DDP) algorithm for closed-loop execution of manipulation primitives with frictional contact switches. Planning and control of these primitives is challenging as they are hybrid, under-actuated, and stochastic. We address this by developing hybrid DDP both to plan finite horizon trajectories with a few contact switches and to create linear stabilizing controllers. We evaluate the performance and computational cost of our framework in ablations studies for two primitives: planar pushing and planar pivoting. We find that generating pose-to-pose closed-loop trajectories from most configurations requires only a couple (one to two) hybrid switches and can be done in reasonable time (one to five seconds). We further demonstrate that our controller stabilizes these hybrid trajectories on a real pushing system. A video describing our work can be found at <https://youtu.be/YGSe4cUfq6Q>.

## I. INTRODUCTION

Complex manipulation tasks can often be decomposed into a sequence of simpler behaviors. For example, picking a credit-card off a table may consist of a pull to cantilever the card followed by a grasp to acquire the card. In part motivated by this observation, researchers have studied manipulation behavior segmented into *manipulation primitives* such as grasping, pulling, pushing, or pivoting.

These primitives are often used to facilitate planning and control; however, defining these primitives, planning within a primitive, and scheduling primitives are all current areas of research. One approach is to use narrowly defined primitives that are simpler to plan and control at the expense of needing more of them; for example, when every contact mode/type is a primitive [1]. On the other hand, complex, more expressive primitives often incur a higher computational cost and can be challenging to realize on a physical system. An extreme case is when all possible mode sequences are considered in the development of a single behavior [2], [3].

This work proposes a fast planning and control framework that supports a small number of hybrid switches for primitives of moderate complexity with underactuated frictional dynamics. Switching contact formations within a primitive increases its expressiveness, which can reduce the number of primitives needed and, consequently, ease their scheduling.

**Contributions** We develop an algorithm for executing manipulation primitives with frictional contact switches. Our approach extends input-constrained differential dynamic programming (DDP) to handle hybrid dynamics. We plan finite-horizon trajectories while considering a small number of contact switches (up to four) within a reasonable amount

of time (one to five seconds). We also present a numerical study on the convergence properties and computational requirements of our algorithm for two manipulation primitives: planar pushing and planar pivoting. Our experiments find that

- The ability to select and switch contact locations is key to the success of a primitive.
- Only one to two contact location switches are needed to converge from most initial configurations.

Finally, we show that our framework can plan and control hybrid trajectories on a real planar pushing system.

**Paper Structure** We begin by reviewing the basics of DDP, extensions to handle constraints, and our hybrid algorithm in Sec. III. We then derive the motion models for planar pushing and pivoting in Sec. IV. Section V describes simulation studies that evaluate the success rate and computation time of the algorithm for these two primitives. We experimentally validate the algorithm for the planar pushing primitive in Sec. VI, and finally, we summarize the results, limitations, and directions of future work in Sec. VII.

## II. RELATED WORK

In this section we discuss some related research on manipulation primitives and DDP.

**Manipulation Primitives** There is a long history in robotic manipulation of developing the mechanics of and planning algorithms for primitives. Mason [4] introduced the mechanics of planar pushing, which have since been studied by a number of researchers [5], [6], [7]. This line of work has been extended to many other primitives, including prehensile-pushing [8], tumbling [9], pivoting [10], [11], [12], scooping [13], [14], tilting [15], and dynamic in-hand sliding [16].

Researchers have also focused on sequencing primitives to achieve complex manipulations [17], [18], [19], [20], [21]. For example, Toussaint et al. [22] use a few kinematic primitives to realize diverse set of behaviors; however, this approach is only verified in simulation. Woodruff et al. [1] treat each contact formation as a different primitive. They then execute closed-loop dynamic motions with a fixed primitive-schedule using a planar manipulator in a low-gravity environment. Our framework balances these approaches and is similar to that of Hou et al. [12], who develop a planner for two moderate-complexity primitives and demonstrate pose-to-pose re-orientation on a physical system.

**Differential Dynamic Programming** DDP is an iterative, indirect trajectory optimization method that leverages the structure in Bellman's equation to achieve local optimality.

<sup>1</sup>Intelligence Community Postdoctoral Research Fellowship Program, Massachusetts Institute of Technology, Cambridge, MA

<sup>2</sup>Dept. of Mechanical Engineering, Massachusetts Institute of Technology

Originally developed by Jacobson and Mayne [23] for unconstrained systems, it has since been extended to systems with box input constraints [24], linear input constraints [25], and non-linear constraints on input and states [26].

Relevant to this work, Tassa et al. [27] and Mordatch et al. [28] use DDP with smoothed contact models to plan and stabilize trajectories for legged robots. Yamaguchi and Atkeson [29] apply DDP to the problem of planning for graphodynamical systems, and they use a sample based approach to determine the mode sequence. Moreover, Pajarinen et al. [30] consider DDP for planar pushing, and they optimize over a continuous mixture of discrete actions that is forced back into fully discrete actions at convergence. Our work improves on these approaches by considering exact rigid-body frictional contacts, determining the mode schedule using DDP, and retaining the anytime property of the algorithm.

### III. HYBRID PLANNING AND CONTROL

We first review the basics of DDP (Sec. III-A) and its extension to input-constrained systems (Sec. III-B). We then describe our hybrid DDP algorithm in Sec. III-C.

#### A. DDP Preliminaries

Consider a discrete-time dynamical system of the form

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) \quad (1)$$

where  $\mathbf{f}$  is a smooth function that maps the system's state ( $\mathbf{x} \in \mathbb{R}^n$ ) and control input ( $\mathbf{u} \in \mathbb{R}^m$ ) to the next state. The goal is to find an input trajectory  $\mathbf{U} := \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$  that minimizes an additive cost function,

$$J(\mathbf{x}_0, \mathbf{U}) = l_f(\mathbf{x}_N) + \sum_{k=0}^{N-1} l(\mathbf{x}_k, \mathbf{u}_k). \quad (2)$$

Here  $k$  is the time-step,  $l$  is the running cost,  $l_f$  is the final cost,  $N$  is the time-horizon,  $\mathbf{x}_0$  is the initial state, and  $\mathbf{x}_1 \dots \mathbf{x}_N$  are determined by integrating (1) forward in time. We can define the *optimal cost-to-go* at the  $k$ -th time-step using Bellman's equation [31],

$$V_k(\mathbf{x}_k) = \min_{\mathbf{u}_k} [l(\mathbf{x}_k, \mathbf{u}_k) + V_{k+1}(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k))], \quad (3)$$

with the terminal condition  $V_N(\mathbf{x}_N) = l_f(\mathbf{x}_N)$ .

To handle the non-linearity in (3), DDP iteratively optimizes a quadratic approximation near an initial trajectory. The algorithm starts with a forward pass that integrates (1) from an initial state  $\mathbf{x}_0$  using a current guess of the input trajectory  $\mathbf{U}$ . This is followed by a backward pass that computes a local solution to (3) using a quadratic Taylor expansion to iterate on the value of  $\mathbf{U}$ . This sequence of forward and backward passes is repeated until convergence.

The Taylor expansion of the argument of (3) about a nominal  $(\mathbf{x}, \mathbf{u})$  pair is given by

$$Q(\delta \mathbf{x}, \delta \mathbf{u}) = l(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u}) - l(\mathbf{x}, \mathbf{u}) + V_{k+1}(\mathbf{f}(\mathbf{x} + \delta \mathbf{x}, \mathbf{u} + \delta \mathbf{u})) - V_{k+1}(\mathbf{f}(\mathbf{x}, \mathbf{u})). \quad (4)$$

---

#### Algorithm 1: Input constrained DDP

---

```

1 initialize  $\leftarrow \mathbf{x}_0, \mathbf{U}_0$ 
2 while not converged do
3    $V_N \leftarrow l_f(\mathbf{x}_N)$ 
4   for  $k = N - 1$  to 0 do
5      $Q_k(\delta \mathbf{x}_k, \delta \mathbf{u}_k) \leftarrow (5)$ 
6      $\mathbf{k} \leftarrow$  solve QP (9);  $\mathbf{K} \leftarrow$  see [25]
7      $\delta \mathbf{u}_k \leftarrow \mathbf{k} + \mathbf{K} \delta \mathbf{x}_k$ 
8     Propagate value  $\leftarrow (7)$ 
9   end
10   $\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}_0$ 
11  for  $k = 0$  to  $N - 1$  do
12     $\hat{\mathbf{u}}_k \leftarrow$  ProjectFeasible( $\mathbf{u}_k + \mathbf{k} + \mathbf{K}(\hat{\mathbf{x}}_k - \mathbf{x}_k)$ )
13     $\hat{\mathbf{x}}_{k+1} = \mathbf{f}(\hat{\mathbf{x}}_k, \hat{\mathbf{u}}_k)$ 
14  end
15   $\mathbf{X} \leftarrow \hat{\mathbf{X}}, \mathbf{U} \leftarrow \hat{\mathbf{U}}$ 
16 end
```

---

The quadratic approximation of  $Q$  can be written as:

$$Q(\delta \mathbf{x}_k, \delta \mathbf{u}_k) \approx \frac{1}{2} \begin{bmatrix} 1 \\ \delta \mathbf{x}_k \\ \delta \mathbf{u}_k \end{bmatrix}^T \begin{bmatrix} 0 & \mathbf{q}_x^T & \mathbf{q}_u^T \\ \mathbf{q}_x & \mathbf{Q}_{xx} & \mathbf{Q}_{xu} \\ \mathbf{q}_u & \mathbf{Q}_{xu}^T & \mathbf{Q}_{uu} \end{bmatrix} \begin{bmatrix} 1 \\ \delta \mathbf{x}_k \\ \delta \mathbf{u}_k \end{bmatrix}, \quad (5)$$

where the block matrices are functions of  $V_{k+1}$ ,  $l$ ,  $\mathbf{f}$ , and their first and second derivatives [24]. The control modification is obtained by minimizing (5) with respect to  $\delta \mathbf{u}$  for some state perturbation  $\delta \mathbf{x}$ :

$$\delta \mathbf{u}^* = -\mathbf{Q}_{uu}^{-1}(\mathbf{q}_u + \mathbf{Q}_{xu}^T \delta \mathbf{x}) = \mathbf{k} + \mathbf{K} \delta \mathbf{x}, \quad (6)$$

where  $\mathbf{k} = -\mathbf{Q}_{uu}^{-1} \mathbf{q}_u$  is the feed-forward control and  $\mathbf{K} = -\mathbf{Q}_{uu}^{-1} \mathbf{Q}_{xu}^T$  is the feedback gain. Substituting this for  $\delta \mathbf{u}$  in (5) gives a quadratic model for  $V$

$$\begin{aligned} \Delta V &= \frac{1}{2} \mathbf{k}^T \mathbf{Q}_{uu} \mathbf{k} \\ \mathbf{V}_x &= \mathbf{q}_x - \mathbf{K}^T \mathbf{Q}_{uu} \mathbf{k} \\ \mathbf{V}_{xx} &= \mathbf{Q}_{xx} - \mathbf{K}^T \mathbf{Q}_{uu} \mathbf{K}. \end{aligned} \quad (7)$$

The backward pass initializes the quadratic approximation of  $V$  with  $l_f(\mathbf{x}_N)$  and its derivatives, and then recursively computes (6) and propagates the value approximation (7).

The algorithm then integrates (1) to compute a new trajectory, completing one iteration. The control during this forward pass is set to  $\mathbf{u} + \delta \mathbf{u}^*$  with  $\delta \mathbf{x}_k$  taken as the difference between  $\mathbf{x}_k$  across consecutive iterations. Note that  $\mathbf{Q}_{uu}$  is regularized to ensure that  $\mathbf{Q}_{uu}^{-1}$  exists and a line-search over  $\mathbf{k}$  keeps  $\delta \mathbf{u}$  and  $\delta \mathbf{x}$  small to ensure cost-reduction. These steps enable convergence from an arbitrary initialization [24].

#### B. Input Constrained DDP

Now consider a system where the control inputs are linearly constrained by inequality (or equality) constraints:

$$\mathbf{A}(\mathbf{x}_k) \mathbf{u}_k \geq \mathbf{b}(\mathbf{x}_k). \quad (8)$$

Here  $\mathbf{A}$  and  $\mathbf{b}$  are potentially nonlinear functions of the state. This class of constraints can represent both planar friction and force-balance constraints for a fixed contact




Fig. 1. Planar pushing with four sticking contacts ( $P_0, P_1, P_2, P_3$ ) at the center of each side. Only one contact can be active at a time, and the active contact force ( $\mathbf{f}^i, i \in \{0, 1, 2, 3\}$ ) must lie within its friction cone ( $FC^i$ ).

mode. The DDP algorithm is modified in two ways for these constraints [25]. First, (6) is replaced by a constrained quadratic program (QP) evaluated at the nominal  $\mathbf{x}$ :

$$\begin{aligned} \min_{\delta \mathbf{u}} \quad & Q(\mathbf{0}, \delta \mathbf{u}) \\ \text{s.t.} \quad & \mathbf{A}(\mathbf{x})(\mathbf{u} + \delta \mathbf{u}) \geq \mathbf{b}(\mathbf{x}) \end{aligned} \quad (9)$$

The solution to this QP gives the value of the feed-forward control  $\mathbf{k}$  satisfying the input constraints. We then consider the state variation  $\delta \mathbf{x}$  when solving for the feedback gain  $\mathbf{K}$ . Details on this are given by Murray and Yakowitz [25].

Second, even though  $\mathbf{k}$  satisfies the input constraints, the new control computed using during the forward pass (Line 12, Alg. 1) can violate feasibility. Consequently, it must be projected onto the constraint set. When  $\mathbf{A}$  and  $\mathbf{b}$  have a simple geometric representation (e.g., a box or a cone), we can algebraically project the new control input onto the feasible set. In other cases, we solve another QP detailed by Murray and Yakowitz [25]. The input-constrained DDP algorithm is outlined in Alg. 1.

### C. Hybrid DDP

We extend input-constrained DDP to systems with hybrid switches. We use DDP as a subroutine to (a) explore and rank all feasible mode sequences and (b) optimize the trajectory and feedback law associated with the best mode sequence. In addition to the initial state and input trajectory, the user can specify the maximum number of hybrid switches ( $N_s$ ) and the set of hybrid modes ( $\mathcal{M}$ ) of dimension  $M$ .

Our algorithm first builds a depth  $N_s + 1$  tree of length- $N_s$  possible sequences of contact states. The number of leaves in the tree is upper-bounded by  $N_s^M$ . Each leaf is a trajectory associated with a fixed contact mode sequence with the  $N_s$  switch locations distributed evenly along the planning horizon ( $N$ ). We use input-constrained DDP with a small iteration limit to optimize each leaf in the tree and approximate its cost. We then select the leaf with the lowest cost and fix the mode sequence to that of the selected leaf. Finally, we optimize the trajectory and controller associated with the best leaf using input-constrained DDP (Alg 1).

For computational efficiency, we initialize DDP with inputs that result in static equilibrium and prune the tree during exploration by eliminating trajectories that cannot




Fig. 2. Planar pivoting in the gravity plane with the pivot at the lower left corner ( $P_0$ ). We consider three sticking contacts at the other corners ( $P_1, P_2, P_3$ ). Only one contact is active at a time, and the ground-reaction ( $\mathbf{f}^0$ ) and active contact ( $\mathbf{f}^i, i \in \{1, 2, 3\}$ ) forces obey Coulomb friction.

satisfy static equilibrium after a contact switch. The hyper-parameters of our algorithm (explored in Sec. V) are the maximum number of hybrid switches ( $N_s$ ), the set of hybrid modes ( $\mathcal{M}$ ), the planning horizon ( $N$ ), and the maximum number of DDP iterations during tree generation ( $N_i$ ).

In summary, our algorithm can be thought of as an exhaustive tree-search over mode sequences with pruning based on static equilibrium.

## IV. MANIPULATION PRIMITIVES

Here we derive the equations-of-motion (EOM) for the primitives used in this work: quasi-static planar pushing (Sec. IV-A) and dynamic planar pivoting (Sec. IV-B). Our choice of primitives is meant to illustrate that our approach handles manipulation in the horizontal and gravitational planes, as well as, quasi-static and dynamic systems.

### A. Quasi-static Planar Pushing

We consider quasi-static pushing in a horizontal plane with four potential (only one active at a time) sticking point contacts (Fig. 1). The object's state is  $\mathbf{x} = [x, y, \theta]^T$ , where  $x$  and  $y$  are the position of its center-of-mass (COM) and  $\theta$  is its orientation. The discrete-time, quasi-static EOM are

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \dot{\mathbf{x}}_k, \quad (10)$$

where  $\dot{\mathbf{x}}_k$  is the object's twist at time-step  $k$  and  $\Delta t$  is the time step's duration. Using force balance and a ellipsoidal approximation [32] of the *limit surface* [33], we can write

$$\dot{\mathbf{x}}_k = \mathbf{R}(\theta) \mathbf{L}(\mathbf{J}^i)^T \mathbf{f}_k^i. \quad (11)$$

Here  $i \in \{0, 1, 2, 3\}$  is the index of the active contact,  $\mathbf{f}^i$  is the active contact force in the body-frame,  $\mathbf{J}^i$  is the Jacobian for the active contact,  $\mathbf{R}$  is the rotation between body and world frames, and  $\mathbf{L}$  is the gradient of the limit surface w.r.t the support wrench. Hogan et al. [34] give further details.

The  $i$ -th contact force must lie within its friction cone ( $FC^i$ ) for sticking contact, and we upper-bound the normal contact force by  $N_{\max}^i$ . In a frame whose positive  $x$ -axis is aligned with the contact normal (i.e., the contact frame),

$$\begin{aligned} 0 &\leq f_n^i \leq N_{\max}^i \\ -\mu f_n^i &\leq f_t^i \leq \mu f_n^i. \end{aligned} \quad (12)$$




Fig. 3. Pushing trajectories from eight initial conditions. The goal is a solid gray square, contact forces are drawn with black arrows, and the left side of object is shown in blue. Successful (unsuccessful) trajectories are depicted in green (patterned orange). The green shaded trajectory in (b) has a contact switch.

Here  $(f_n^i, f_t^i)$  are the normal and tangential components of  $\mathbf{f}^i$  in the contact frame, and  $\mu$  is the coefficient of friction.

### B. Dynamic Planar Pivoting

We also consider dynamic pivoting in the gravity plane about a sticking frictional pivot with the ground (Fig. 2). The object is rotated about this pivot by sticking contacts at one of the other three corners. Each contact is treated as a point-on-line contact with the line fixed at  $45^\circ$  with respect to both sides of that corner. The object's state is  $\mathbf{x} = [\theta, \dot{\theta}]^T$ , where  $\theta$  is its orientation and  $\dot{\theta}$  is its angular velocity. We write the discrete-time dynamics of the system as

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \dot{\mathbf{x}}_k, \quad (13)$$

where  $\dot{\mathbf{x}}_k = [\dot{\theta}_k, \ddot{\theta}_k]^T$  and  $\ddot{\theta}$  is

$$\ddot{\theta} = \frac{1}{I} (\mathbf{r}^0 \times \mathbf{R}(\theta)^T \mathbf{f}^0 + \mathbf{r}^i \times \mathbf{f}^i) \quad (14)$$

from Newton's angular momentum principle. Here  $i \in \{1, 2, 3\}$  is the index of the active contact,  $\mathbf{f}^i$  is the active contact force in the body-frame,  $\mathbf{f}^0$  is the ground reaction force in the world-frame,  $\mathbf{R}^T$  is the rotation between world and body frames,  $I$  is the object's mass moment of inertia, and  $\mathbf{r}^0$  ( $\mathbf{r}^i$ ) is the vector from the COM to the ground (active) contact. We also constrain the ground reaction force using Newton's linear momentum principle,

$$\dot{\mathbf{p}}_c = \mathbf{f}^i + \mathbf{R}(\theta)(\mathbf{f}^0 + \mathbf{g}), \quad (15)$$

where  $\mathbf{g}$  is the gravitational force and  $\dot{\mathbf{p}}_c \in \mathbb{R}^2$  is the time-derivative of linear momentum of the COM. Note that  $\dot{\mathbf{p}}_c$  can be computed in terms of  $\theta$ ,  $\dot{\theta}$ , and  $\ddot{\theta}$ . Finally, we enforce that all contact forces (active and ground) lie within their friction cones and place an upper bound on the normal contact forces.

## V. NUMERICAL STUDIES

We use our algorithm to plan pose-to-pose trajectories for both primitives. We present a number of representative trajectories in Sec. V-A and conduct ablation studies in Sec. V-B. We use a simple quadratic total cost of the form

$$J(\mathbf{x}_0, \mathbf{U}) = \Delta \mathbf{x}_N^T \mathbf{Q}_N \Delta \mathbf{x}_N + \sum_{k=0}^{N-1} \Delta \mathbf{x}_k^T \mathbf{Q} \Delta \mathbf{x}_k + \mathbf{u}^T \mathbf{R} \mathbf{u}, \quad (16)$$

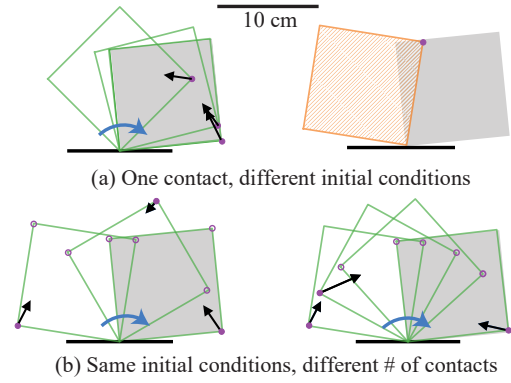


Fig. 4. Pivoting trajectories from two representative initial conditions. The goal is the gray square, the direction of rotation is clockwise (blue arrow), the contact forces are drawn with black arrows, and enabled corner-contacts are marked with purple circles with active contacts filled in. Successful (unsuccessful) trajectories are depicted in green (patterned orange).

to generate all trajectories. Here  $\Delta \mathbf{x}$  is the distance to the goal and  $\mathbf{Q}_N$ ,  $\mathbf{Q}$ , and  $\mathbf{R}$  are positive definite diagonal matrices.

### A. Simulated Trajectory Planning

Representative planar pushing and pivoting trajectories are shown in Fig. 3 and Fig. 4, respectively.

**Planar Pushing** We compute trajectories from eight initial conditions for enabled contact sets of size one and three (Fig. 3). The goal is  $(x, y, \theta) = (0, 0, 0)$ . Trajectories are considered successful if the final errors in  $x$ ,  $y$ , and  $\theta$  are less than 5 cm, 5 cm, and  $5^\circ$ , respectively. We set the maximum number of hybrid switches ( $N_s$ ) to 1, the maximum iterations during tree generation ( $N_i$ ) to 10, and the planning horizon ( $N$ ) to 24. Moreover, we use a time-step ( $\Delta t$ ) of 0.5 s, a coefficient of friction ( $\mu$ ) of 0.3 at both frictional contacts, and allow a maximum normal force ( $N_{\max}$ ) of 0.5 N.

With only the left contact enabled (Fig. 3a), as expected, the algorithm finds solutions for initial conditions that are to the left of the goal. Note that this corresponds to pure input-constrained DDP. With three contacts enabled (Fig. 3b), the algorithm finds trajectories to the goal from all initial conditions. The algorithm usually only needs to select the best contact; however, it needs a hybrid switch for one trajectory (solid green). The mean planning time is 0.40 s and 0.70 s for one and three enabled contacts, respectively.

**Planar Pivoting** We compute trajectories for enabled contact sets of size one, two and three from two initial conditions (Fig. 4). The goal is  $\theta = 10^\circ$  and  $\dot{\theta} = 0^\circ \text{s}^{-1}$ . Successful trajectories have final errors in  $\theta$  and  $\dot{\theta}$  that are less than  $10^\circ$  and  $10^\circ \text{s}^{-1}$ , respectively. The object's mass is 0.1 kg and its density is uniform. We use  $N_s = 2$ ,  $N_i = 10$ ,  $N = 16$ ,  $\Delta t = 0.05$  s,  $\mu = 0.5$ , and  $N_{\max} = 10$  N.

For pivoting, we observe that the ability to reason about contact switches is important. For example, we cannot pivot the object from  $80^\circ$  to  $10^\circ$  with only a single contact enabled using pure input-constrained DDP (Fig. 4a). Moreover, the planner finds different mode sequences with more than one enabled contacts (Fig. 4b). Finally, the mean planning time is 0.67, 3.12, and 7.30 s for the trajectories with one, two, and three enabled contacts, respectively.

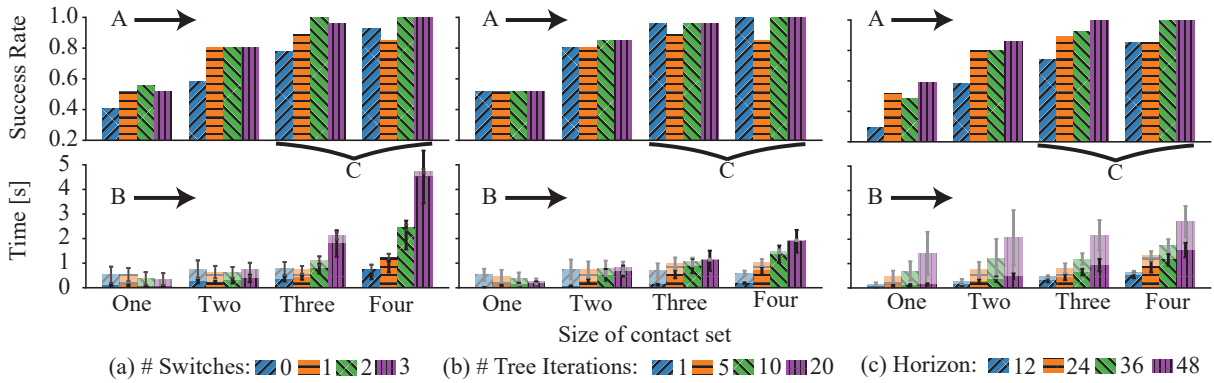


Fig. 5. Ablation studies for planar pushing. The top row shows the success rate, and the bottom row shows the planning time for generating the trajectory-tree (solid) and optimizing the best trajectory (transparent). The error bars depict  $\pm$  one standard deviation. On the x-axis we show the size of contact set, and the different colors indicate (a) the maximum number of hybrid switches, (b) the number of DDP iterations used when generating the trajectory-tree and (c) the total horizon of the trajectory in discrete steps. The labels A-C highlight important trends that are discussed in the text.

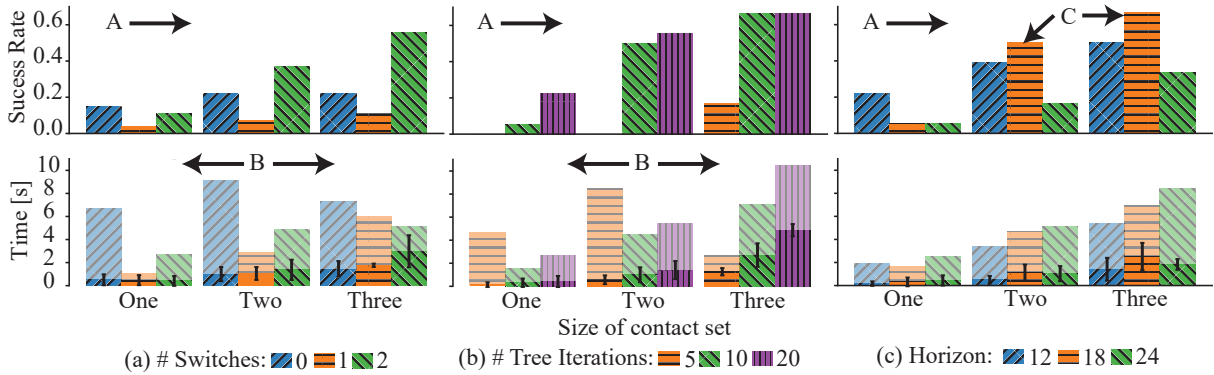


Fig. 6. Ablation studies for planar pivoting following the convention in Fig. 5.

## B. Ablation Studies

We also conduct one-dimensional ablation studies that explore how the hyper-parameters of our algorithm affect success rate (defined above) and planning time.

**Planar Pushing** In Fig. 5, we depict the effect of the number enabled contacts and one other ablation parameter: (a) number of hybrid switches, (b) number of DDP iterations during mode selection (tree generation), and (c) the horizon of the trajectory. When not varied, these parameters are fixed to  $N_s = 1$ ,  $N_i = 5$ , and  $N = 24$ . For each parameter, we consider all active contact combinations and plan trajectories from 27 initial conditions for each contact combination.

We find that across all parameters, success rate increases with the number of enabled contacts (A, Fig. 5). This is intuitive as allowing for more contact normal directions increases controllability. This success, however, comes with an increased planning time (B, Fig. 5), though planning time is most affected by the number of hybrid switches (Fig. 5a). We also find that success rate is insensitive to the choice of  $N_i$  (Fig. 5b), and is robust to all hyper-parameter changes with three or four enabled contacts (C, Fig. 5). Finally, we can achieve a success rate of 100% with a planning time of  $\sim 1$  s for a number of different hyper-parameter combinations.

**Planar Pivoting** We present the effects of varying the same hyper-parameters as above for pivoting in Fig. 6. When not varied, these parameters are fixed to  $N_s = 2$ ,  $N_i = 10$ , and

$N = 18$ . For each parameter combination, we consider all contact combinations and plan trajectories from two initial conditions for object with aspect ratios of 0.5, 1.0, and 1.5.

Similar to the pushing primitive, we find that success rate increases with the number of enabled contacts (A, Fig. 6); however, we are only able to reach a maximum success rate of 0.6-0.7. Interestingly, there is not a corresponding increase in planning time (B, Fig. 6), though the overall planning time is higher than for planar pushing due to the more complex dynamics of pivoting. Our results suggest that the planner is also more sensitive to the choice of hyper-parameters; for example, planning over an 18-step horizon outperforms planning over 12 and 24-step horizons (C, Fig. 6).

## VI. EXPERIMENTAL RESULTS

We evaluate our approach on a real planar pushing system.

**Experimental Set-up** We use an industrial robotic manipulator (ABB IRB 120, Fig. 7). The object rests on a flat plywood surface and is moved by a metallic rod attached to the robot. The feedback controller (6) runs at  $\sim 250$  Hz, and the object pose is tracked using a motion capture system (Vicon, Bonita) at 300 Hz. The object’s physical properties are described in Sec. V-A, and it has a length of 0.09 m.

We convert the inputs of our model (applied force and time-step length) into position commands for the robot manipulator by integrating the following kinematic relation,

$$\mathbf{x}_{k+1}^p = \mathbf{x}_k^p + \Delta t_k \mathbf{R} \mathbf{J}^T \mathbf{L} (\mathbf{J}^i)^T \mathbf{f}_k^i. \quad (17)$$




Fig. 7. Robot arena for the planar pushing experiments.




Fig. 8. Open-loop (orange,  $n=5$ ) and closed-loop (green,  $n=5$ ) straight-line pushes. The light-gray box is the initial condition, and the black-outlined gray box is the goal. The controller significantly reduces error.

Here  $\mathbf{x}^p$  is the Cartesian position of the contact in the world frame, and  $\mathbf{R}$ ,  $\mathbf{J}^i$ , and  $\mathbf{L}$  are defined in Sec. IV-A.

**Straight Line Pushing** We evaluate the performance of our controller on straight-line pushing (Fig. 8). We execute five open-loop and five closed-loop 40 cm pushes. The open-loop standard deviation (s.d.) in error is  $x=3.5$  cm,  $y=4.2$  cm, and  $\theta=85^\circ$ . The controller significantly reduces the closed-loop s.d. in error to  $x=0.5$  cm,  $y=0.1$  cm, and  $\theta=1.25^\circ$ .

**Hybrid Pushing** We also validate our framework for three pushes starting from more challenging initial conditions, with zero, one, and two contact switches (Fig. 9). Our planner finds pushing trajectories that reach the goal and are effectively stabilized by the controller. However, slipping between the pusher and the object results in slightly higher final pose error than in the straight-line pushing scenario.

## VII. DISCUSSION

We summarize the major findings of this work (Sec. VII-A), discuss some important limitations (Sec. VII-B), and propose directions for future work (Sec. VII-C).

### A. Conclusions

We develop a hybrid DDP algorithm for dynamical systems with frictional contact and discontinuous switches. Our approach reasons over a finite horizon, supports a small number of contact switches, and generates a linear stabilizing controller. Our approach can quickly generate closed-loop trajectories that drive most initial conditions to the goal for the two planar manipulation primitives considered. Finally, we demonstrate our controller's ability to track planned trajectories on a real pushing system.

### B. Limitations

Though we can drive any initial condition to the origin for planar pushing, this is not the case for planar pivoting. We believe this is due to poor initialization; while static

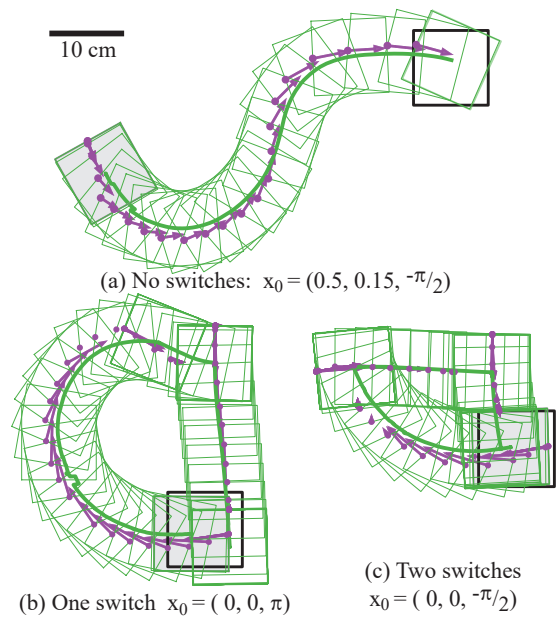


Fig. 9. Closed-loop pushes with (a) no contact switches, (b) one contact switch, and (c) two contact switches. The object pose and Cartesian trajectory is shown in green. Nominal contact locations and applied forces are shown with purple circles and arrows, respectively. The light-gray box is the initial condition, and the black-outlined box is the goal.

equilibrium is a fixed-point for quasi-static planar pushing, it is not the same for dynamic planar pivoting. Furthermore, though we can accurately track straight-line pushes on a real system, tracking errors are larger for more complex trajectories. This is likely a result of slipping between the pusher and object, and a lower-level controller that enforces sticking (e.g., [21]) would complement our approach. Finally, the computational cost of our algorithm increases combinatorially with the number of allowed switches; however, we show that a small number of switches is sufficient for executing for planar pushing and planar pivoting.

### C. Future Work

One extension is to apply our approach on a wider range of primitives, including pulling, prehensile pushing, rolling, tilting, etc. This will require both identifying appropriate mechanics models and adapting the hybrid DDP framework. In particular, we would like to improve both the initialization and pruning procedure for our algorithm and to reduce its dependence on user defined hyper-parameters. We would also like to explore more sophisticated controllers, as detailed by Hogan and Rodriguez [6], to reason about contact-sliding relative to the object.

## ACKNOWLEDGEMENTS

This research was supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at the Massachusetts Institute of Technology, administered by Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence.

## REFERENCES

- [1] J. Z. Woodruff and K. M. Lynch, "Planning and control for dynamic, nonprehensile, and hybrid manipulation tasks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4066–4073.
- [2] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *International Journal of Robotics Research*, vol. 33, no. 1, pp. 69–81, Jan. 2014.
- [3] Z. Manchester, N. Doshi, R. J. Wood, and S. Kuindersma, "Contact-implicit trajectory optimization using variational integrators," *The International Journal of Robotics Research*, p. 0278364919849235, 2019.
- [4] M. T. Mason, "Mechanics and planning of manipulator pushing operations," *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 53–71, 1986.
- [5] K. M. Lynch and M. T. Mason, "Stable pushing: Mechanics, controllability, and planning," *The International Journal of Robotics Research*, vol. 15, no. 6, pp. 533–556, 1996.
- [6] F. Hogan and A. Rodriguez, "Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics," in *WAFR*, 2016.
- [7] J. Zhou and M. T. Mason, "Pushing revisited: Differential flatness, trajectory planning and stabilization," in *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2017.
- [8] N. Chavan-Dafle, R. Holladay, and A. Rodriguez, "In-hand manipulation via motion cones," in *Robotics: Science and Systems*, 2018.
- [9] N. Sawasaki and H. Inoue, "Tumbling objects using a multi-fingered robot," *Journal of the Robotics Society of Japan*, vol. 9, no. 5, pp. 560–571, 1991.
- [10] A. Holladay, R. Paolini, and M. T. Mason, "A general framework for open-loop pivoting," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3675–3681.
- [11] Y. Karayiannidis, C. Smith, D. Kragic, et al., "Adaptive control for pivoting with visual and tactile feedback," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 399–406.
- [12] Y. Hou, Z. Jia, and M. T. Mason, "Fast planning for 3d any-pose-reorienting using pivoting," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1631–1638.
- [13] J. C. Trinkle, J. M. Abel, and R. P. Paul, "An investigation of frictionless enveloping grasping in the plane," *The International journal of robotics research*, vol. 7, no. 3, pp. 33–51, 1988.
- [14] J. C. Trinkle, "On the stability and instantaneous velocity of grasped frictionless objects," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 5, pp. 560–572, 1992.
- [15] M. A. Erdmann and M. T. Mason, "An exploration of sensorless manipulation," *IEEE Journal on Robotics and Automation*, vol. 4, no. 4, pp. 369–379, 1988.
- [16] J. Shi, J. Z. Woodruff, P. B. Umbanhowar, and K. M. Lynch, "Dynamic in-hand sliding manipulation," *IEEE Transactions on Robotics*, vol. 33, no. 4, pp. 778–795, 2017.
- [17] J. C. Trinkle and J. J. Hunter, "A framework for planning dexterous manipulation," in *1991 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1991, pp. 1245–1251.
- [18] H. Terasaki and T. Hasegawa, "Motion planning of intelligent manipulation by a parallel two-fingered gripper equipped with a simple rotating mechanism," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 2, pp. 207–219, 1998.
- [19] J. Barry, K. K. Hsiao, L. P. Kaelbling, and T. Lozano-Pérez, *Manipulation with Multiple Action Types*. Heidelberg: Springer International Publishing, 2013, pp. 531–545.
- [20] F. Suárez-Ruiz and Q.-C. Pham, "A framework for fine robotic assembly," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 421–426.
- [21] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, "Tactile dexterity: Manipulation primitives with tactile feedback," in *2020 IEEE international conference on robotics and automation (ICRA)*, 2020.
- [22] M. Toussaint, K. Allen, K. A. Smith, and J. B. Tenenbaum, "Differentiable physics and stable modes for tool-use and manipulation planning," in *Robotics: Science and Systems*, 2018.
- [23] D. H. Jacobson and D. Q. Mayne, *Differential dynamic programming*. North-Holland, 1970.
- [24] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1168–1175.
- [25] D. M. Murray and S. J. Yakowitz, "Constrained differential dynamic programming and its application to multireservoir control," *Water Resources Research*, vol. 15, no. 5, pp. 1017–1027, 1979.
- [26] Z. Xie, C. K. Liu, and K. Hauser, "Differential dynamic programming with nonlinear constraints," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 695–702.
- [27] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and Stabilization of Complex Behaviors through Online Trajectory Optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [28] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 43–43:8, Aug. 2012.
- [29] A. Yamaguchi and C. G. Atkeson, "Differential dynamic programming for graph-structured dynamical systems: Generalization of pouring behavior with different skills," *IEEE-RAS International Conference on Humanoid Robots*, no. November 2016, pp. 1029–1036, 2016.
- [30] J. Pajarinen, V. Kyrki, M. Koval, S. Srinivasa, J. Peters, and G. Neumann, "Hybrid control trajectory optimization under uncertainty," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 5694–5701.
- [31] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [32] R. D. Howe and M. R. Cutkosky, "Practical force-motion models for sliding manipulation," *The International Journal of Robotics Research*, vol. 15, no. 6, pp. 557–572, 1996.
- [33] S. Goyal, A. Ruina, and J. Papadopoulos, "Planar sliding with dry friction part I. limit surface and moment function," *Wear*, vol. 143, no. 2, pp. 307–330, 1991.
- [34] F. R. Hogan, E. R. Grau, and A. Rodriguez, "Reactive planar manipulation with convex hybrid mpc," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 247–253.