



UNIVERSITY OF LEEDS

This is a repository copy of *Combining Coarse and Fine Physics for Manipulation using Parallel-in-Time Integration*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/149219/>

Version: Accepted Version

---

**Proceedings Paper:**

Agboh, WC orcid.org/0000-0002-0242-0215, Ruprecht, D orcid.org/0000-0003-1904-2473 and Dogar, MR orcid.org/0000-0002-6896-5461 (2022) Combining Coarse and Fine Physics for Manipulation using Parallel-in-Time Integration. In: Robotics Research. ISRR 2019. Springer Proceedings in Advanced Robotics. International Symposium on Robotics Research (ISRR) 2019, 06-10 Oct 2019, Hanoi, Vietnam. Springer , pp. 725-740. ISBN 978-3-030-95458-1

[https://doi.org/10.1007/978-3-030-95459-8\\_44](https://doi.org/10.1007/978-3-030-95459-8_44)

---

© 2022 The Author(s), under exclusive license to Springer Nature Switzerland AG. This is an author produced version of an article published in Springer Proceedings in Advanced Robotics. Uploaded in accordance with the publisher's self-archiving policy.

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Combining Coarse and Fine Physics for Manipulation using Parallel-in-Time Integration

Wisdom C. Agboh<sup>1</sup>, Daniel Ruprecht<sup>2</sup>, and Mehmet R. Dogar<sup>1</sup>

<sup>1</sup> School of Computing, University of Leeds, United Kingdom.

<sup>2</sup> School of Mechanical Engineering, University of Leeds, United Kingdom.

**Abstract.** We present a method for fast and accurate physics-based predictions during non-prehensile manipulation planning and control. Given an initial state and a sequence of controls, the problem of predicting the resulting sequence of states is a key component of a variety of model-based planning and control algorithms. We propose combining a coarse (i.e. computationally cheap but not very accurate) predictive physics model, with a fine (i.e. computationally expensive but accurate) predictive physics model, to generate a hybrid model that is at the required speed and accuracy for a given manipulation task. Our approach is based on the Parareal algorithm, a parallel-in-time integration method used for computing numerical solutions for general systems of ordinary differential equations. We adapt Parareal to combine a coarse pushing model with an off-the-shelf physics engine to deliver physics-based predictions that are as accurate as the physics engine but run in substantially less wall-clock time, thanks to parallelization across time. We use these physics-based predictions in a model-predictive-control framework based on trajectory optimization, to plan pushing actions that avoid an obstacle and reach a goal location. We show that with hybrid physics models, we can achieve the same success rates as the planner that uses the off-the-shelf physics engine directly, but significantly faster. We present experiments in simulation and on a real robotic setup. Videos are available here: <https://youtu.be/5e9oTeu4JOU>.

**Keywords:** Physics-based Manipulation, Model-predictive-control

## 1 Introduction

We present a method for fast and accurate physics predictions during non-prehensile manipulation planning and control. Take the case study in Fig.1, where a cylindrical object moves towards the right, pushing a box. We are interested in predicting the motion of the pushed box, in a fast and accurate way. To achieve this, we combine *coarse* physics models with *fine* physics models. By coarse models, we mean computationally cheap but relatively inaccurate predictive physical models. For example in Fig.1a, we use a coarse model to compute the motion of the box. The motion is not completely realistic, but we can compute it extremely fast (7 ms wall-clock time to compute a simulated 8 s push). By fine models, we mean computationally expensive but accurate predictive physical models. In Fig.1d, we use a fine model (in this case, the Mujoco simulator

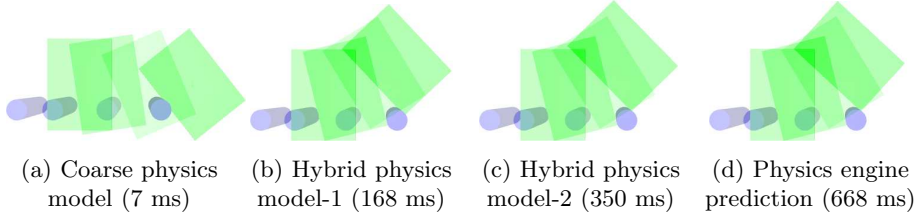


Fig. 1: A spectrum of physics predictions ranging from cheapest and least accurate (a) to expensive and most accurate (d).

[36]) to compute the motion of the same box. The motion is more realistic, but it also requires much more time to compute (668 ms).

We combine these two models to deliver a prediction that is as accurate as the fine model but runs in substantially less wall-clock time. The motion predicted in Fig.1b is similar to the fine model prediction, but is four times faster to compute. The motion predicted in Fig.1c is indistinguishable from the fine model prediction for real world manipulation purposes, and is two times faster to compute.

Given an initial state and a sequence of controls, the problem of predicting the resulting sequence of states is a key component of a variety of model-based planning and control algorithms [15, 17, 16, 38, 5, 14, 2, 35, 20]. Mathematically, such a prediction requires solving an *initial value problem*. Typically, those are solved through numerical integration over time-steps (e.g. Euler’s method or Runge-Kutta methods) using an underlying physics model. However, the speed with which these accurate physics-based predictions can be performed is still slow [8] and faster physics-based predictions can contribute significantly to contact-based/non-prehensile manipulation planning and control.

There are several ways that could be used to construct coarse models for manipulation planning. Quasi-static physics, which ignores accelerations, is widely used in non-prehensile manipulation planning and control [22, 25, 13], and can be seen as a coarse model. Learning is another method which can be used to generate approximate but fast predictions [27, 19, 33, 40]. Recently, with the advance of deep-learning, there has been multiple attempts at learning approximate “intuitive” physics models which are then used for manipulation planning [3, 11, 10, 32, 26, 7, 4]. Especially when these networks are faced with novel objects that are not in their training data (e.g. consider a network trained with boxes and cylinders, but used to predict the motion of an ellipse) they can generate approximate predictions of motion, and therefore are good candidates as coarse models.

A key question we investigate is whether we can combine such cheap but approximate models, with expensive but more accurate and general models (such as physics engines) to generate a hybrid model that is at the required speed and accuracy for a given manipulation task.

We do this by using a coarse physics model to obtain a rough initial guess of the state at each time point of a trajectory. Then, we evaluate the fine physics

model in parallel across time starting from the initial guesses. Thereafter, we combine the coarse and fine predictions using the iterative Parareal algorithm [24, 21, 34].

In this paper, we use this approach to perform physics-based predictions within a planner for robotic manipulation. Specifically, we consider the task of pushing an object to a goal location while avoiding obstacles. We provide a cheap coarse model and combine it with the Mujoco physics engine as the fine model.<sup>3</sup> The planner performs trajectory optimization to generate a control sequence, executes the first control in the sequence, and then re-runs the trajectory optimizer, in a model-predictive-control fashion. We present this planner in Sec. 3. As a baseline, we use the same planner, with the fine model, Mujoco, as the predictive model. We conduct experiments in simulation and on a real setup and show that the planner with hybrid physics models achieves the same success rates but faster.

To the best of our knowledge, the use of Parareal for contact dynamics (and in general for robotic planning and control) has not been investigated before. When used for contact dynamics, the original Parareal formulation can produce infeasible states where rigid bodies penetrate. We extend Parareal to handle these infeasible state updates through projections to the feasible state space.

### 1.1 Related work

Parareal has been used in different areas, e.g. to simulate incompressible laminar flows [37], or to simulate dynamics in quantum chemistry [23]. It was introduced by Lions et al. [21].

Combining different physics models for robotic manipulation has been the topic of other recent research as well, even though the focus has *not* been improving prediction speed. Kloss et al. [18] focus on the question of accuracy and generalization in combined neural-analytical models. Ajay et al. [4] focus on modeling of the inherent stochastic nature of the real world physics, by combining an analytical, deterministic rigid-body simulator with a stochastic neural network.

We can make physics engines faster by using larger simulation time steps, however this decreases the accuracy and can quickly result in unstable behavior. To generate stable behaviour at large time-step sizes, Pan and Manocha [30] propose an integrator for articulated body dynamics by using only position variables to formulate the dynamic equation. Moreover, Fan et al. [9] propose linear-time variational integrators of arbitrarily high order for robotic simulation and use them in trajectory optimization to complete robotics tasks. Recent work have focused on making the underlying planning and control algorithms faster. For example, Giftthaler et al. [12] introduced a multiple-shooting variant of the trajectory optimizer - iterative linear quadratic regulator (ilqr) which has shown impressive results for real-time nonlinear optimal control of complex robotic systems [29, 31].

<sup>3</sup> We use Mujoco since it is recently the most widely used physics-engine for model-based planning [32, 1, 26, 7].

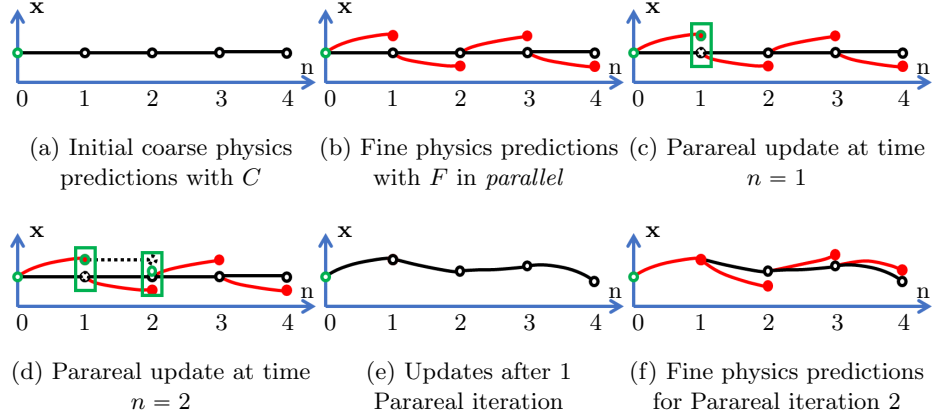


Fig. 2: Combining coarse and fine physics with the Parareal algorithm

## 2 Combining physics models for planning

Given an initial state  $\mathbf{x}_0$  and a sequence of  $N$  controls  $\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$ , we are interested in predicting the resulting sequence of states  $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$  of a physical system. As an example, we consider the problem of pushing an object to a goal location with a cylindrical pusher. The system's state consists of the pose  $\mathbf{q}$  and velocity  $\dot{\mathbf{q}}$  of the pusher  $P$  and slider  $S$ :  $\mathbf{x}_n = [\mathbf{q}_n^P, \mathbf{q}_n^S, \dot{\mathbf{q}}_n^P, \dot{\mathbf{q}}_n^S]^T$ . The slider's pose consists of the translation and rotation of the object on the plane  $\mathbf{q}^S = [q^{S_x}, q^{S_y}, q^{S_\theta}]^T$ . The pusher's pose is:  $\mathbf{q}^P = [q^{P_x}, q^{P_y}]^T$  and control inputs are velocities  $\mathbf{u}_n = [u_n^x, u_n^y]^T$  applied on the pusher for a control duration of  $\Delta t$ .

To predict the next state of the system given an initial state and a control input, we need a physics model  $F$ . We use a general physics engine [36] to model the system dynamics. It solves Newton's equations of motion for the complex multi-contact dynamics problem:

$$\mathbf{x}_{n+1} = F(\mathbf{x}_n, \mathbf{u}_n, \Delta t). \quad (1)$$

Normally, computing all states  $\mathbf{x}_n$  happens in a serial fashion, by evaluating (1) first for  $n = 0$ , then for  $n = 1$ , etc. Instead, we replace this inherently serial procedure by a parallel-in-time integration process. Specifically, we adapt the Parareal algorithm for the contact-based manipulation problem.

Parareal begins with a rough initial guess of the state at each time point  $n$  of the trajectory as shown in Fig. 2a. To get an initial guess, we define a second, coarse physics model:

$$\mathbf{x}_{n+1} = C(\mathbf{x}_n, \mathbf{u}_n, \Delta t) \quad (2)$$

It needs to be computationally cheap relative to the fine model but does not need to be very accurate.

The next step is to evaluate the fine physics model in *parallel* starting from  $N$  initial guesses as shown in Fig. 2b. Thereafter, we do a coarse sweep across the time points. We start from the initial state  $\mathbf{x}_0$  and make a coarse prediction for the next state  $\mathbf{x}_1$  (dotted lines in Fig. 2c). Now, we have 3 approximations

of the state at time  $n = 1$ . We linearly combine these approximations to get an update for  $\mathbf{x}_1$  (in green). Then starting from this new update for  $\mathbf{x}_1$ , we make a coarse prediction for the next state  $\mathbf{x}_2$  (dotted lines in Fig. 2d at  $n = 2$ ). We combine the three approximations to get an update at  $n = 2$ . We continue this coarse sweep for all time points to get the updated trajectory in Fig. 2e. This is the end of the first iteration. We then repeat the whole process iteratively using new updates as initial guesses as shown in Fig. 2f.

In summary, Parareal starts by computing rough initial guesses  $\mathbf{x}_n^{k=0}$  of the system states using the coarse model. The newly introduced superscript  $k$  counts the number of Parareal iterations. In each Parareal iteration, the guess is then refined via

$$\mathbf{x}_{n+1}^{k+1} = C(\mathbf{x}_n^{k+1}, \mathbf{u}_n, \Delta t) + F(\mathbf{x}_n^k, \mathbf{u}_n, \Delta t) - C(\mathbf{x}_n^k, \mathbf{u}_n, \Delta t), \quad (3)$$

for all timesteps  $n = 0, \dots, N - 1$ . The key point in iteration (3) is that evaluating the fine physics model can be done in parallel for all  $n = 0, \dots, N - 1$ , while only the fast coarse model has to be computed serially.

Parareal iterations converge exactly to the fine physics solution after  $k = N$  iterations. After one iteration,  $\mathbf{x}_1^1$  is exactly the fine solution. We can see this in Fig. 2c where the two coarse physics predictions in Eq. 3 are same and cancel out. After two iterations,  $\mathbf{x}_1^1$  and  $\mathbf{x}_2^2$  are exactly the fine solutions and so on.

The idea is to stop Parareal at much earlier iterations such that it requires significantly less wall-clock time than running  $F$  serially step-by-step. To do this,  $C$  must be computationally much cheaper than  $F$ .

Parareal can be thought of as producing a spectrum of solutions increasing in accuracy and computational cost, from the cheap coarse physics model to the expensive fine physics model — i.e. the  $N$  different approximations after each iteration. An important question is which of these models to choose; i.e. how many iterations of Parareal to use? To decide on the required prediction accuracy, we rely on recent work which analyzes the stochasticity in real-world pushing [39, 6]. We propose to stop Parareal when the approximation error with respect to the fine model is below the real-world pushing stochasticity.

Note that, for the sake of simplicity, we assume here that the number of controls  $N$  and the number of processors used to parallelize in time are identical.

## 2.1 Expected speedup performance of Parareal

We can describe the expected performance of Parareal by a simple theoretical model [28]. Let  $c_c$  and  $c_f$  be the time needed to compute the coarse physics model  $C$  and the fine physics model  $F$  respectively, for a duration of  $\Delta t$ . The speedup of Parareal  $s_p$  over the serial fine model is approximately:

$$s_p = \frac{N \cdot c_f}{(1 + K) \cdot N \cdot c_c + K \cdot c_f} = \frac{1}{(1 + K) \frac{c_c}{c_f} + \frac{K}{N}} \quad (4)$$

This illustrates the importance of finding a cheap coarse model that minimizes the ratio  $c_c/c_f$ . In that case, speedup will be determined mainly by the number

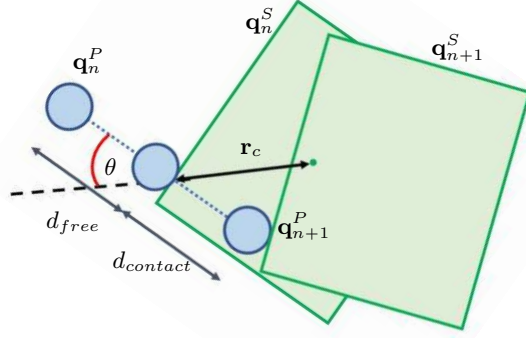


Fig. 3: Coarse physics model

of iterations  $K$ . For example, for a coarse model with negligible cost, after  $K = 1$  Parareal iteration with  $N = 4$  sub-intervals, then the theoretical speedup would be  $s_p \sim N/K = 4$ . That is, we can expect to make physics predictions about *four times faster* than using only the fine physics model in serial.

## 2.2 Coarse physics model

As a case study, we consider the challenging problem of pushing an object. We seek a general coarse physics model with the following requirements:

- It must be significantly cheaper to compute with respect to the fine model.
- It must provide a physics prediction for all possible pusher motions but can be inaccurate.
- It must provide a prediction for sliders of any shape and inertial parameters.

Instead of solving Newton's equation of motion for the multi-contact dynamics problem, we propose a simple kinematic pushing model  $C(\mathbf{x}_n, \mathbf{u}_n, \Delta t)$ . It moves the slider with the same linear velocity as the pusher, as long as there is contact between the two. We also apply a rotation to the slider, based on the position and direction of the contact, with respect to the center of the object. Formally, given the linear velocity of the pusher as the controls  $\mathbf{u}_n = [u_n^x, u_n^y]^T$ , the next state of the system is given by;

$$\mathbf{q}_{n+1}^S = \mathbf{q}_n^S + [u_n^x, u_n^y, \omega]^T \cdot p_c \cdot \Delta t \quad (5)$$

$$p_c = \frac{d_{\text{contact}}}{d_{\text{contact}} + d_{\text{free}}}, \quad \omega = K_\omega \cdot \frac{\|\mathbf{u}_n\| \cdot \sin \theta}{\|\mathbf{r}_c\|} \quad (6)$$

$$\dot{\mathbf{q}}_{n+1}^S = \{[u_n^x, u_n^y, \omega]^T \text{ if } p_c > 0, \dot{\mathbf{q}}_n^S \text{ otherwise}\} \quad (7)$$

$$\mathbf{q}_{n+1}^P = \mathbf{q}_n^P + \mathbf{u}_n \cdot \Delta t, \quad \dot{\mathbf{q}}_{n+1}^P = \mathbf{u}_n. \quad (8)$$

In Eq. 5 the slider’s pose is updated as described above. Here,  $p_c$  is the ratio of the distance  $d_{contact}$  travelled by the pusher when in contact with the slider and the total pushing distance as shown in Fig. 3.  $\mathbf{r}_c$  is a vector from the contact point to the object’s center (green dot) at the current state  $\mathbf{q}_n^S$ ,  $\theta$  is the angle between the pushing direction and the vector  $\mathbf{r}_c$ . Moreover,  $\omega$  is the coarse angular velocity induced by the pusher on the slider, where  $K_\omega$  is a positive constant parameter.

Also note that, even though Fig. 3 shows the pusher and slider in contact at the next time step, this does not have to be so; i.e. the coarse model can leave the two in separation.

In Eq. 7 the velocity of the slider is updated to be the same as the current pusher velocity if there is any contact. In Eq. 8 the pusher position and velocity are updated.

### 2.3 Infeasible states

The new iterate  $\mathbf{x}_{n+1}^{k+1}$  given by the Parareal iteration (Eq. 3) can be an infeasible state where the pusher and slider penetrate each other. Contact dynamics is not well-defined for such states. It can lead to infinitely large object accelerations and an unstable fine physics model. We have not encountered such a problem of infeasible (or unallowed) states in other dynamics domains that Parareal has been applied to.

To handle these cases, we project the infeasible state to the nearest feasible state. We write the following optimization problem:

$$\mathbf{q}_{n+1}^{S*} = \arg \min_{\mathbf{q}_{n+1}^S} \|\mathbf{q}_{n+1}^S - \mathbf{q}_{n+1}^{S_{infeasible}}\|, \quad s.t. \quad d_p \leq 0 \quad (9)$$

where  $\mathbf{q}_{n+1}^{S_{infeasible}}$  is the infeasible slider’s pose, and  $d_p$  is the penetration depth. The goal is to find the nearest slider pose  $\mathbf{q}_{n+1}^{S*}$  that satisfies the no-penetration constraint  $d_p \leq 0$ .

We can use an off-the-shelf optimizer to find a solution rather efficiently. However, for simple systems we can analytically find the penetration depth and move the slider along the contact normal to resolve penetration.

In Sec. 4.2, we evaluate the open-loop pushing performance of our hybrid physics models. Our goal is to use these hybrid models for planning and control.

## 3 Push Planning and Control

We use the predictive models described above in a planning and control framework for pushing an object on a table to a goal location, avoiding a static obstacle. This task retains many challenges of general robotic control through contact such as impulsive contact forces, under-actuation and hybrid dynamics (separation, sticking, sliding, e.t.c.). We present an example scene and execution in Fig. 4.

To solve this problem, we take an optimization approach. Given the obstacle and goal position and geometry, the current state of the pusher and slider  $\mathbf{x}_0$ , and



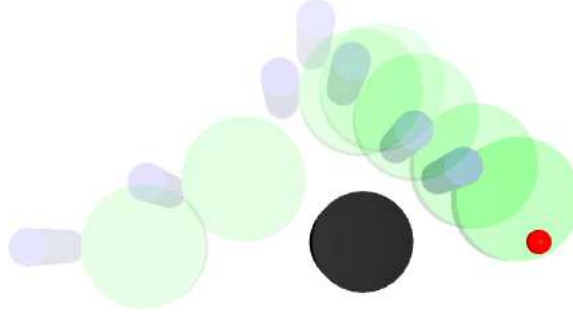


Fig. 4: Push planning with a hybrid physics model to avoid an obstacle (in black) while pushing the cylindrical slider to a goal location (in red).

an initial candidate sequence of controls  $\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$ , the optimization procedure outputs an optimal sequence  $\{\mathbf{u}_0^*, \mathbf{u}_1^*, \dots, \mathbf{u}_{N-1}^*\}$  according to some defined cost. We explain this optimization process, and the cost formulation that is optimized, below (Sec. 3.1).

The predictive models that we have developed earlier in the paper are used within this optimizer to *roll-out* a sequence of controls, to predict the states  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , which are then used to compute the cost associated with those controls.

Once the optimization produces a sequence of controls, we use it in a model-predictive-control (MPC) fashion, by executing only the first control in the sequence. Afterwards, we update  $\mathbf{x}_0$  with the observed state of the system, and repeat the optimization to generate a new control sequence. This is repeated until task completion. We consider the task completed if the slider reaches the goal region (success), if it hits the obstacle (failure), if it falls off the edge of the table (failure) or if a maximum number of controls are executed before failure occurs.

When we repeat the optimization within MPC, we warm-start it by using the previously optimized control sequence as the initial candidate sequence. For the very first optimization, the initial candidate sequence is generated as a straight line push towards the goal (which collides with the obstacle in all our scenes).

Such an optimization-based MPC approach to pushing manipulation is frequently used [5, 15, 18, 2]. Here, our focus is to evaluate the performance of different predictive physics models described before in the paper within such a framework.

### 3.1 Trajectory Optimization

In this section we use the shorthand  $\mathbf{u}_{0:N-1}$  to refer to the control sequence  $\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$ . Similarly for states we use  $\mathbf{x}_{0:N}$ .

Our goal is to find an optimal control sequence  $\mathbf{u}_{0:N-1}^*$  for a planning horizon  $N$ , given an initial state  $\mathbf{x}_0$ , and an initial candidate control sequence  $\mathbf{u}_{0:N-1}$ .

We define the cost function  $J$ , for a given control sequence and the corresponding state sequence:

$$J(\mathbf{x}_{0:N}, \mathbf{u}_{0:N-1}) = \sum_{n=1}^{N-1} J_n(\mathbf{x}_n, \mathbf{u}_{n-1}, \mathbf{u}_n) + w \cdot J_N(\mathbf{x}_N), \quad (10)$$

where  $J_n$  is the running cost at each step,  $w$  is a positive weighting constant,  $J_N$  is the terminal (final) cost function.

The output of optimization is the minimizing control sequence:

$$\mathbf{u}_{0:N-1}^* = \arg \min_{\mathbf{u}_{0:N-1}} J(\mathbf{x}_{0:N}, \mathbf{u}_{0:N-1}), \quad s.t. \quad \mathbf{x}_{n+1} = f(\mathbf{x}_n, \mathbf{u}_n, \Delta t), \quad \mathbf{x}_0 \text{ fixed}. \quad (11)$$

Here,  $f$  is the system dynamics constraint that must be satisfied at all times.

We define the running cost for our pushing around an obstacle problem as:

$$J_n(\mathbf{x}_n, \mathbf{u}_{n-1}, \mathbf{u}_n) = w_s \cdot (1/||[q^{sx}, q^{sy}]^T - \mathbf{p}_{obs}||^2) + w_p \cdot (1/||\mathbf{q}^p - \mathbf{p}_{obs}||^2) \\ + w_u \cdot ||\mathbf{u}_t - \mathbf{u}_{t-1}||^2 + W_E$$

where  $w_s, w_p, w_u$  are positive constant weights.  $\mathbf{p}_{obs}$  is the position vector of the static obstacle to be avoided, and  $[q^{sx}, q^{sy}]$  are the x,y positions of the slider respectively. The above formula associates high cost for the slider or the pusher to approach the obstacle. It has a smoothness cost, to prevent high accelerations. Additionally, it has a constant edge cost  $W_E$  for the slider falling off the table.

We define the final cost  $J_N$  as:  $J_N(\mathbf{x}_N) = ||[q^{sx}, q^{sy}]^T - \mathbf{p}_{goal}||^2$ , where  $\mathbf{p}_{goal}$  is the position vector of the target/goal location.

There exists different optimization methods to solve this problem [20, 17, 38, 15, 2]. The main difference lies in the way the cost gradient is computed for a given sequence of controls. For ease of implementation, here we use derivative-free stochastic sampling-based methods [17, 38, 2]. Particularly, we use the algorithm presented in Agboh and Dogar [2]. In each optimization iteration, to find the cost gradient at the current control sequence, these stochastic sampling methods generate multiple noisified versions of the current control sequence, they roll-out these noisy controls to find the cost associated with each one, and use these costs to compute a numerical gradient, which is then used to update the control sequence to minimize the cost. The roll-out of these noisy control sequences to compute the resulting states and the cost is where we use the physics models.

### 3.2 Parareal and MPC

The Parareal framework for generating hybrid physics models yields itself well to model-predictive control. Recall from Fig. 2e that after 1 Parareal iteration, physics prediction for the first state  $\mathbf{x}_1$  is exactly the same as the fine model. This means planning is accurate at least for the first action for all our hybrid models. This aligns well with our MPC framework since we execute only the first action and then re-plan.

## 4 Experiments and Results

In our experiments, we address three key issues. First, we investigate how fast Parareal converges to the fine physics solution for pushing tasks. Second, we

investigate the open-loop pushing performance of different physics models and compare it with real-world data. Finally, we investigate how the different physics models generated by Parareal (at different iterations) can be used within a planning and control framework to complete non-prehensile manipulation tasks.

The goal of Parareal for push planning is to simulate physics faster than any given fine physics model. Therefore, we must stop Parareal at much earlier iterations to achieve meaningful speedup. However, we must also understand how different Parareal’s predictions are with respect to the fine physics predictions at different iterations. We investigate Parareal’s convergence for specific pushing examples in Sec. 4.1. In addition, in Sec. 4.1, we measure the empirical speedup we get from Parareal, and compare it with the theoretical speedup that was presented in Sec. 2.1.

During push planning and control, we must decide on a Parareal iteration that gives us an acceptable approximation to the fine solution. To this end, in Sec 4.2, we conduct open-loop pushing experiments. We statistically investigate Parareal’s approximation error with respect to the fine solution, for a particular number of Parareal iterations. To do that, we start from a large number of random initial states and apply different control sequences open-loop. We then analyze how these statistical approximation errors compare with the standard deviation of the real-world uncertainty during similar pushing tasks [39].

In Sec 4.3 we investigate the performance of different physics models produced by Parareal, when used within the MPC framework described in Sec. 3 to push an object to a goal region, avoiding an obstacle. We compare the success rates and total task completion times for the different physics models (the coarse model, Parareal at different iterations, and the fine model). We perform these experiments on a real robot setup Sec 4.3.

#### 4.1 Parareal convergence for pushing

**Parareal convergence for specific pushing examples:** We consider simulating the results of applying a control sequence starting from an initial state for a box and a cylinder as shown in Fig. 5. We consider four cases: pushing a cylinder from the center (Fig. 5a), pushing a cylinder from the side, pushing a box from the center, and finally pushing a box from the side (Fig. 5b). The control sequence used here is  $\mathbf{u}_{0:3} = \{[25, 0], [25, 0], [25, 0], [25, 0]\}mm/s$  where each control input is applied for a control duration  $\Delta t = 1s$  such that the total pushing distance is  $100mm$ .

We use the Mujoco [36] physics engine as the fine physics model. To make the fine model as fast as possible, we run it at the largest possible simulation time-step ( $1ms$ ) for our model. Beyond this time-step, the physics engine becomes unstable and breaks down. In addition, all experiments are run on a desktop computer (Intel(R) Xeon(R) CPU E3-1225 v3) with  $N = 4$  cores. At each iteration of Parareal, we calculate the root mean square (RMS) error between Parareal’s predictions and the physics engine’s predictions of the corresponding sequence of states. These RMS errors can be seen in Fig. 5 for two different cases. The results are similar for others. Note that the errors are given in log

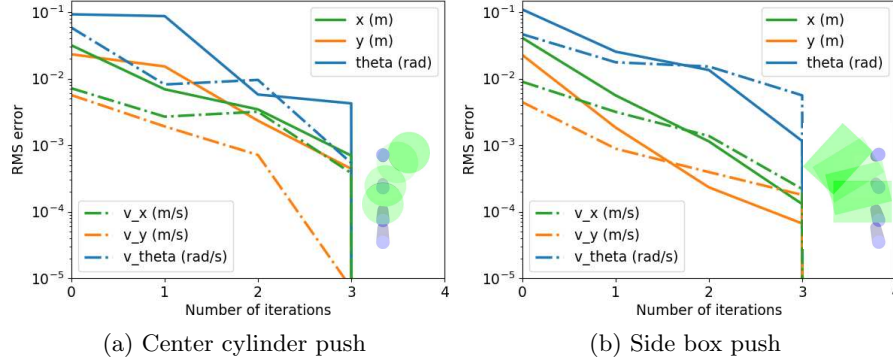


Fig. 5: Root mean square error (in log scale) along the full trajectory for pushing a cylinder (a) and box (b) from the center and side respectively, for increasing Parareal iterations. The motions are illustrated lower-right in each plot.

scale for the full state of the slider (pose and velocities). In general, we see a quick decrease in the error along the full trajectory starting from the large error of the coarse model at iteration 0. In addition, at the final iteration, we verify that the Parareal solution is exactly the same as using the fine model since the errors go to 0.

**Parareal speedup for specific pushing examples:** Using each hybrid physics model, we repeatedly predict the sequence of states (which is deterministic for a given physics model) and record the total time it takes (which varies slightly depending on computer load). In Fig. 6, we see the average prediction time over 100 runs for each physics model for a box side push. These actual prediction times are close to the expected prediction time (Eq. 4) for the different physics models. For example, at 1 Parareal iteration we spent 28% of the time spent by the full physics engine, i.e. about *four times faster*. The results for the cylinder side push, cylinder center push and box center push are similar to those in Fig. 6.

## 4.2 Open-loop pushing experiments

We compare the predictions of different physics models (Parareal iterations) for open-loop pushing. We start from 100 randomly sampled initial states<sup>4</sup>. At each initial state, we used three different control sequences, giving 300 different slider trajectories. The three control sequences  $\mathbf{u}^{\{1,2,3\}}$  that we used at each initial state were fixed, and are given by:

$$\mathbf{u}_{0:3}^i = v_c \cdot \{v, v, v, v\}, \quad v = [\cos(\alpha_i), \sin(\alpha_i)], \quad \alpha_i = \{0^\circ, 15^\circ, -15^\circ\}, i \in 1, 2, 3$$

$v_c$  is a constant pushing speed:  $v_c = 25 \text{ mm/s}$ .

We applied each control input for a duration of  $\Delta t = 1.5 \text{ s}$  such that the total distance travelled by the pusher is  $150 \text{ mm}$  in all cases. We calculated for

<sup>4</sup> We change the pusher’s position along the direction perpendicular to the direction of motion. We sample uniformly within the edges of the slider (rectangle).

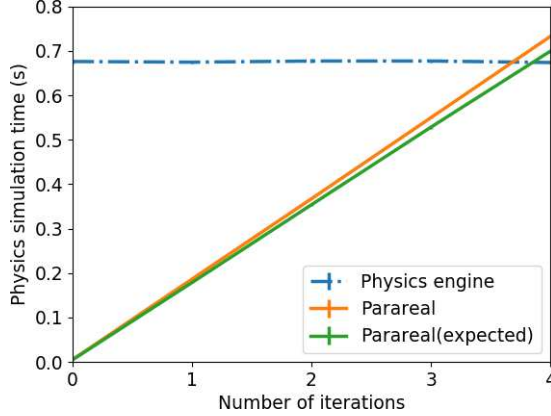


Fig. 6: Physics simulation time averaged over 100 runs for a box side push within 95 % confidence interval of the mean.

each physics model (Parareal iteration), the RMS difference of the final state (in comparison with the final state prediction of the fine physics model) for the 300 trajectories.

Our results are shown in Table. 1. We see that on the average the coarse physics model is quite inaccurate but with increasing Parareal iterations, the mean difference from the fine physics model goes to zero. However, to decide on how much error with respect to the physics engine is appropriate for pushing tasks, we look at the real-world’s uncertainty for pushing dynamics.

Yu et al. [39] provide real-world pushing data for a similar pusher-slider system. Starting at the same initial state, they push a box repeatedly in the real-world with a cylindrical pusher and record the resulting final positions. The pushing distance is  $150mm$  with a quasi-static pushing speed of  $20mm/s$ . As shown in Table 1, they record a translation standard deviation of  $8.10mm$  and a rotation standard deviation of  $4.20^\circ$  on a plywood surface. Notice that for Parareal after 2 iterations, we see a mean translation difference of  $6.39mm$  and a mean rotation difference of  $3.82^\circ$  when compared to the fine model (physics engine) predictions.

We conclude that, for real-world purposes, it should not be necessary to run Parareal for more than 2 iterations, as approximating the physics engine more accurately than the inherent uncertainty in real-world pushing should not contribute to real-world performance. Note that, 2 Parareal iterations here corresponds to a) exactly the fine physics predictions for your first two actions and b) a model that is two times faster than the physics engine.

### 4.3 Push planning and control with hybrid physics models

We measure the performance of the different physics models when used within the optimization and MPC framework, described in Sec. 3, to push an object to a goal region, while avoiding an obstacle.

Table 1: Open-loop pushing

	Mean trans. diff. (mm)	Mean rot. diff. (deg)
Coarse Physics	62.67	17.63
Parareal-1 iter.	28.43	6.30
Parareal-2 iter.	6.39	3.82
Parareal-3 iter.	2.47	0.79
Parareal-4 iter.	0.00	0.00
	Trans. std. (mm)	Rot. std. (deg)
Push dataset [39]	8.10	4.20

Our real robot setup is shown in Fig. 8 where we have a Robotiq two-finger gripper holding the cylindrical pusher. We place markers on the pusher and slider to sense their full pose in the environment with the OptiTrack motion capture system. We consider 5 randomly generated scenes (one shown in Fig. 8) where the pusher must avoid the obstacle at the center of the table before bringing the slider to the goal location. For each of the 5 scenes, we used the coarse physics model, Parareal iterations (1,2, and 3), and the full physics engine for push planning and control. That is a total of 25 planning and control runs with the real robot. We plan using our various physics models and execute actions in the real world in an MPC fashion. For the stochastic trajectory optimizer, our control sequences contain 4 control inputs each applied for a control duration of  $\Delta t = 1s$ . In addition, we use 20 noisy control sequence samples (as explained in Sec. 3.1) per optimization iteration for the trajectory optimizer with an exploration variance of  $10^{-4}$ . Normally, each noisy control sequence of the optimizer is rolled out independently. However, since we use a standard quad-core desktop PC, the *parallelization is across time only*.

As the pusher attempts to bring the slider to a desired goal location, there are three possible failure modes. First, we declare failure when the slider collides with the static obstacle. Second, we declare failure when the pusher is unable to bring the slider to the goal location after executing 20 actions (5 times the number of actions in a given control sequence). Third, we declare failure when the slider falls off the edge of the table.

For each hybrid physics model, we achieved 100% success rate on the real robot. This is same as using a full physics engine, only significantly faster Fig. 7. For instance, using 1 Parareal iteration, we can complete the push planning task about *four times faster* than the physics engine. Note that the planning times here can be reduced by parallelizing the stochastic trajectory optimizer on a PC with more cores. Currently, the 4 cores of our PC is used to *parallelize across time*. Furthermore, in all the 5 scenes considered, the robot was unable to complete the push planning task by using only the coarse model. We present snapshots from the experiments in Fig. 8.

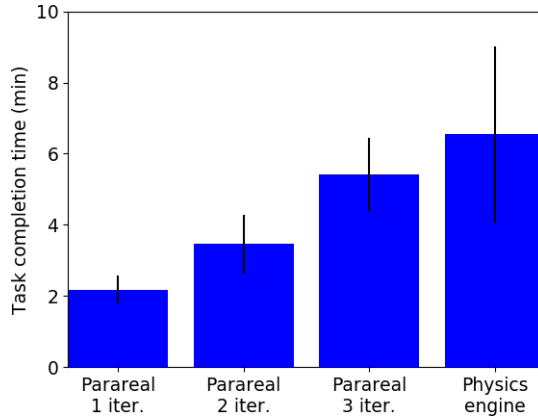


Fig. 7: Total task completion time (within 95 percent confidence interval of the mean) for push planning with obstacle avoidance using different physics models for 100 randomly sampled initial states.

## 5 Discussion and Future Work

The method we presented here opens up important and exciting questions for manipulation planning and control. We discuss some of them briefly below.

*Different coarse models:* We proposed a cheap, general and accurate enough coarse pushing model. This can be replaced with other models.

We would like to explore learned pushing/poking models [3, 11, 10, 33, 18, 40, 4] as coarse models. Devising/learning coarse models for other more complex planning environments (e.g. for manipulation in clutter) is another important problem.

*Degree of parallelization:* In this paper, we used 4 cores for parallelization. With a higher number of parallel slices, it may be possible to get higher speedups.

*Task adaptivity:* Different tasks require different degree of accuracy [1]. For example, think of searching for a sock in the sock drawer, versus searching for a wine glass in the glass cabinet. It is okay for a robot’s physics predictions to be coarse in the former example, which is not the case in the latter. Parareal can be used to explore this spectrum, and generate coarse predictions when it is sufficient for the task, and more accurate predictions as the task requires.

## 6 Acknowledgements

This project was funded from the European Unions’s Horizon 2020 programme under the Marie Skłodowska Curie grant No. 746143, and from the UK EPSRC under grants EP/P019560/1, EP/R031193/1, and studentship 1879668.

## References

- [1] W. C. Agboh and M. R. Dogar. Pushing fast and slow: Task-adaptive planning for non-prehensile manipulation under uncertainty. In *WAFR*, 2018.

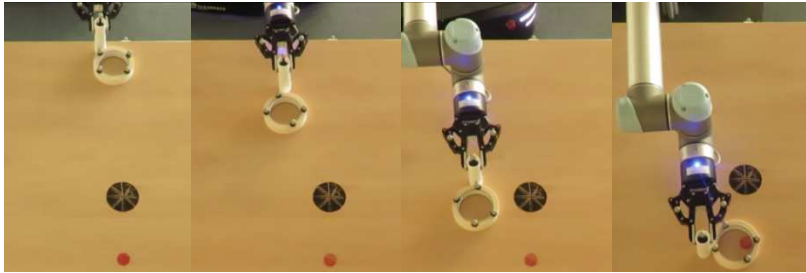


Fig. 8: Pushing with a hybrid physics model. We complete the push planning task about **four times faster** than a physics engine.

- [2] W. C. Agboh and M. R. Dogar. Real-time online re-planning for grasping under clutter and uncertainty. In *IEEE-RAS Humanoids*, 2018.
- [3] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine. Learning to poke by poking: Experiential learning of intuitive physics. In *NeurIPS*, 2016.
- [4] A. Ajay, J. Wu, N. Fazeli, M. Bauza, L. P. Kaelbling, J. B. Tenenbaum, and A. Rodriguez. Augmenting physical simulators with stochastic neural networks: Case study of planar pushing and bouncing. In *IROS*, 2018.
- [5] E. Arruda, M. J. Mathew, M. Kopicki, M. Mistry, M. Azad, and J. L. Wyatt. Uncertainty averse pushing with model predictive path integral control. In *Humanoids*, 2017.
- [6] M. Bauza and A. Rodriguez. A probabilistic data-driven model for planar pushing. In *ICRA*, 2017.
- [7] F. Ebert, S. Dasari, A. X. Lee, S. Levine, and C. Finn. Robustness via retrying: Closed-loop robotic manipulation with self-supervised learning. In *CoRL*, 2018.
- [8] T. Erez, Y. Tassa, and E. Todorov. Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx. In *ICRA*, 2015.
- [9] T. Fan, J. Schultz, and T. Murphey. Efficient computation of higher-order variational integrators in robotic simulation and trajectory optimization. In *WAFR*, 2018.
- [10] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *ICRA*, 2017.
- [11] C. Finn, I. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction. In *NeurIPS*, 2016.
- [12] M. Gifthalder, M. Neunert, M. Stäuble, J. Buchli, and M. Diehl. A family of iterative gauss-newton shooting methods for nonlinear optimal control. *IROS*, 2018.
- [13] S. Goyal, A. Ruina, and J. Papadopoulos. Planar sliding with dry friction part 1. limit surface and moment function. *Wear*, 143(2):307 – 330, 1991.
- [14] J. A. Haustein, J. King, S. S. Srinivasa, and T. Asfour. Kinodynamic randomized rearrangement planning via dynamic transitions between statically stable states. In *ICRA*, 2015.
- [15] F. R. Hogan and A. Rodriguez. Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. *WAFR*, 2016.
- [16] A. M. Johnson, J. King, and S. Srinivasa. Convergent planning. *RA-L*, 2016.
- [17] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. Stomp: Stochastic trajectory optimization for motion planning. In *ICRA*, 2011.



- [18] Alina Kloss, Stefan Schaal, and Jeannette Bohg. Combining learned and analytical models for predicting action effects. *CoRR*, abs/1710.04102, 2017.
- [19] M. Kopicki, S. Zurek, R. Stolkin, T. Moerwald, and J. L. Wyatt. Learning modular and transferable forward models of the motions of push manipulated objects. *Autonomous Robots*, 2017.
- [20] W. Li and E. Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO*, 2004.
- [21] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. Resolution d’edp par un schma en temps pararel . *Comptes Rendus de l’Acadmie des Sciences - Series I - Mathematics*, 332(7):661 – 668, 2001. ISSN 0764-4442.
- [22] K. M. Lynch, H. Maekawa, and K. Tanie. Manipulation and active sensing by pushing using tactile feedback. In *IROS*, 1992.
- [23] Yvon Maday and Gabriel Turinici. Parallel in time algorithms for quantum control: Parareal time discretization scheme. *International Journal of Quantum Chemistry*, 93(3):223–228, 2003.
- [24] Kiehl Martin. Parallel multiple shooting for the solution of initial value problems. *Parallel Computing*, 20(3):275–295, 1994.
- [25] M. Mason. Mechanics and planning of manipulator pushing operations. *IJRR*, 5(3):53–71, 1986.
- [26] Jan Matas, Stephen James, and Andrw J. Davison. Sim-to-real reinforcement learning for deformable object manipulation. In *CORL*, 2018.
- [27] V. M. Merili, T. and H.L. Akn. Push-manipulation of complex passive mobile objects using experimentally acquired motion models. *AuRo*, 38, 2015.
- [28] Michael Minion. A hybrid parareal spectral deferred corrections method. *Commun. Appl. Math. Comput. Sci.*, 5(2):265–301, 2010.
- [29] M. Neunert, M. Stuble, M. Giftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli. Whole-body nonlinear model predictive control through contacts for quadrupeds. *IEEE RA-L*, 3(3):1458–1465, 2018.
- [30] Z. Pan and D. Manocha. Time integrating articulated body dynamics using position-based collocation method. In *WAFR*, 2018.
- [31] Brian Plancher and Scott Kuindersma. A performance analysis of parallel differential dynamic programming on a gpu. In *WAFR*, 2018.
- [32] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *RSS*, 2018.
- [33] F. Ruiz-Ugalde, G. Cheng, and M. Beetz. Fast adaptation for effect-aware pushing. In *Humanoids*, 2011.
- [34] Daniel Ruprecht. Implementing parareal - openmp or mpi? *CoRR*, 2015.
- [35] Y. Tassa, T. Erez, and E. Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *IROS*, 2012.
- [36] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IROS*, 2012.
- [37] J. M. F. Trindade and J. C. F. Pereira. Parallel-in-time simulation of two-dimensional, unsteady, incompressible laminar flows. *Numerical Heat Transfer, Part B: Fundamentals*, 50(1):25–40, 2006.
- [38] G. Williams, A. Aldrich, and E. Theodorou. Model predictive path integral control using covariance variable importance sampling. *CoRR*, 2015.
- [39] K. T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez. More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing. In *IROS*, 2016.
- [40] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell. A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation. *International Journal of Robotics Research*, 2018.