

Controlling the Cascade: Kinematic Planning for N -ball Toss Juggling

Kai Ploeger and Jan Peters

Abstract—Dynamic movements are ubiquitous in human motor behavior as they tend to be more efficient and can solve a broader range of skill domains than their quasi-static counterparts. For decades, robotic juggling tasks have been among the most frequently studied dynamic manipulation problems since the required dynamic dexterity can be scaled to arbitrarily high difficulty. However, successful approaches have been limited to basic juggling skills, indicating a lack of understanding of the required constraints for dexterous toss juggling. We present a detailed analysis of the toss juggling task, identifying the key challenges and formalizing it as a trajectory optimization problem. Building on our state-of-the-art, real-world toss juggling platform, we reach the theoretical limits of toss juggling in simulation, evaluate a resulting real-time controller in environments of varying difficulty and achieve robust toss juggling of up to 17 balls on two anthropomorphic manipulators.

I. INTRODUCTION

The task of juggling pushes the boundary of control performance of robotic manipulators, testing both their speed and dexterity. In this paper, we consider the task of toss juggling to investigate the feasibility of achieving N -ball juggling in the intricate cascade patterns achieved by skilled humans. By pushing the boundary of dynamic robotic manipulation with juggling, we aim to stimulate research for faster and more dexterous robot control, as juggling requires consideration of both mechanical design and control strategy. By opening up the solution space of possible control strategies away from popular static approaches to more dynamic alternatives [1], [2], [3], we gain potential benefits in overcoming torque constraints and under-actuation, energy efficiency, speed, simpler movements, and simpler hardware requirements.

A. Problem Statement

The definition of juggling is controversial both in the juggling and in the scientific community. We restrict our study to single person toss juggling of balls by four assumptions.

Definition I.1 (Toss Juggling).

- Balls are held in hand for a non-zero duration;
- Balls are thrown to the same height;
- Balls are thrown from and aimed at defined locations;
- Balls are thrown and caught to a constant rhythm.

While Definition I.1 is a simplification compared to most definitions of toss juggling, it does not make the task easy. On the contrary, instead of just requiring no ball to drop to the ground, it is our goal to closely track ball trajectories that are well defined through take-off and touch-down location as well as flight time.




Fig. 1. Stable juggling of seven and more balls in a cascade pattern on an anthropomorphic two-arm robotic setup. Trajectories are planned in real-time and adapted to disturbances.

Given Definition I.1, two distinct juggling patterns arise: For odd numbers of balls, they always cross from one hand to the other following the cascade pattern, and for even numbers of balls, they always return to the same hand following the fountains pattern illustrated in figure 2. Juggling a cascade pattern with an even number of balls or a fountain pattern with an odd number of balls is not possible without breaking the symmetry of the movement. As recognizable in Figure 2, in both cases, balls have to be caught at and thrown from sufficiently distant locations to avoid ball collisions.

Juggling is an infinite horizon problem but can be divided into a sequence of finite-horizon problems alternating between catching and throwing. The main problem in catching is to decide where to catch the object. In juggling, the search space is reduced by aiming balls to intercept a desired touch-down location. Also, the difficulty of the throwing task is reduced by aiming each throw at the same target. Why is juggling more challenging than either of its parts? Because each sub-task in the sequence complicates the ones before and after, leading to quickly compounding errors. The catching phase can only start after clearing the previous ball from the hand, leaving just the last fraction of a ball's flight time to execute the catch. The limited time a ball spends in-hand leaves little time to recover from the distribution of catching positions and requires a distribution of throwing movements. As these time constraints can be changed arbitrarily by adjusting the targeted pattern, toss juggling is an excellent subject to explore dynamic manipulation.




Fig. 2. (Left) Crossing throws in a five ball *cascade* (Middle) Non-crossing throws in a four-ball *fountains* pattern (Right) The handcycle with dwell time T_d and vacant time T_v and carry distance d_d .

B. Juggling Theory

Each hands cycle time $T_c = T_v + T_d$ is split into a vacant time T_v that it spends empty and a dwell time T_d holding a ball. The catching movement is executed during T_v and the throwing movement during T_d . Both can be related to the flight time T_f that each ball spends in the air by Claude Shannon's juggling theorem [4], [5]

$$\frac{T_f + T_d}{N_b} = \frac{T_v + T_d}{N_h},$$

where N_b is the number of juggled balls, and N_h is the number of hands. The dwell ratio $R = T_d/T_c$, $R \in (0, 1)$ describes how long balls stay in a hand, with typical values for humans around $R \approx 2/3$. Increasing the dwell ratio reduces the average amount of balls in the air per hand

$$W = \frac{T_f}{T_c} = \frac{N_b}{N_h} - R.$$

Since the horizontal velocity of all balls can be assumed constant and identical without air friction, we can relate the horizontal distance between balls

$$d_b = \frac{d_f}{W} - 2r_b$$

to the distance d_f between take-off and touch-down locations and the ball radii r_b . The horizontal distance between balls is also their minimal distance reached near the top of the pattern. Requiring the minimal distance d_b to be positive results in a kinematic upper bound on N_b that is independent of the throw height.

$$N_b < \left\lfloor \frac{d_f}{r_b} + 2R \right\rfloor \quad (1)$$

As d_f only depends on the carry distance d_d for fountain patterns and can not be increased by moving the hands further apart from each other, fountain patterns are generally considered more difficult than cascade patterns.

Contributions: We propose a novel decomposition of the infinite-horizon toss juggling movement into a sequential short-horizon trajectory optimization problem, identifying the critical constraints necessary for dexterous dynamic manipulation with switching contacts. To the best of our knowledge, this formulation is the first to demonstrate stable juggling of five and more balls with anthropomorphic manipulators.

II. RELATED WORK

From a mathematic perspective, toss juggling has been studied exhaustively [4], [6], [7], [5], [8], [9], while roboticists focused mostly on the related task of paddle juggling that has no dwell time. The robot-ball interactions can then be modeled as impulse exchanges [10], [11] leading to a class of mirror algorithms [12], [13] that were extended to two balls on a spatial robot [14], [15]. Open-loop stability of paddle juggling has been achieved by decelerating during the upward stroke [16], [17] and a parabolic paddle [18].

The two components of toss juggling, throwing and catching, have been studied individually. While throwing, dynamic effects such as dynamic closure [1] and whipping in the kinematic chain [19] are at play. Most approaches capable of reliably hitting targets have been data-driven, by iterative adapting the aiming [20] or by optimizing a movement primitive through trial and error [21].

The necessary complexity of approaches for the ball-catching task depends on the mechanical setup. In case of contacts with low restitution—balls not bouncing off the hand—it is sufficient to intercept the ball trajectory. Online trajectory optimization was successfully applied minimizing the mechanical energy [22] or various kinematic cost functions [23]. A common simplification to the catching problem is the introduction of a catching plane [21]. In case of higher restitution, it is necessary to reduce impact velocities to prevent the ball from bouncing off. Additionally, matching the accelerations [24], [25] widens the time interval of low relative velocity, increasing robustness to disturbances. To solve the more complex task of catching non-spherical objects, distributions over grasp positions and catching behaviors were learned from demonstrations [26].

Claude Shannon did the first investigations in robotic toss juggling in the form of juggling mechanisms solving a particular case of toss juggling, where balls are bounced off a drum on the floor. While mechanisms of this kind [16], [27] implemented the carry distance d_d necessary to avoid collisions by rolling the balls in the hand, later juggling mechanisms [28] also implemented an active carry phase. All juggling mechanisms relied on funnel-shaped hands dissipating excess energy for passive stability.

Most throwing strategies in robotic juggling are open-loop as well. An elliptic trajectory [29] was executed on a robot with two degrees of freedom (DoF). Initializing and improving the trajectory parameterization through a ballistic model led to stable one-handed two-ball juggling. The addition of a hand-tuned throwing motion to a previous catching approach [21] resulted in a shared tree ball cascade with a human partner. Another approach also used a hand-tuned throwing motion [30], but in this case, the ball was caught by the same articulated hand-arm system that threw the ball to juggle two balls in one hand. In our previous work [31] used model-free episodic reinforcement learning on the physical system to optimize a parameterized movement primitive and achieved stable one-handed two-ball toss juggling for up to 33 minutes on an anthropomorphic manipulator.

III. JUGGLING IN TASK SPACE

Dexterous toss juggling can be formulated as a trajectory optimization problem. Since hands do not interact directly, the movements of each hand can be planned individually but need to be conditioned on the states of the balls.

A. Throwing

For precise throwing, a set of take-off constraints has to hold. The take-off hand position $\mathbf{x}(t_{\text{TO}})$ and hand velocity $\dot{\mathbf{x}}(t_{\text{TO}})$ at t_{TO} have to result in a target touch-down ball location $\mathbf{b}(t_{\text{TD}+1}) = \mathbf{b}_{\text{TD,des}}$ after the flight time $T_f = t_{\text{TD}+1} - t_{\text{TO}}$. According to the parabolic ballistic model $\mathbf{b}(t) = \mathbf{b}_{\text{TO}}(t) + \dot{\mathbf{b}}_{\text{TO}}(t)t + \frac{1}{2}\mathbf{g}t^2$ we define our take-off constraints similar to [1] as

$$\mathbf{0} = \mathbf{x}(t_{\text{TO}}) - \mathbf{b}_{\text{TO,des}}, \quad (2)$$

$$\mathbf{0} = \mathbf{x}(t_{\text{TO}}) + \alpha \dot{\mathbf{x}}(t_{\text{TO}})T_f + \frac{1}{2}\mathbf{g}T_f^2 - \mathbf{b}_{\text{TD,des}}, \quad (3)$$

$$\mathbf{0} = \ddot{\mathbf{x}}(t_{\text{TO}}) - \mathbf{g}, \quad (4)$$

assuming $\mathbf{b}_{\text{TO}} = \mathbf{x}_{\text{TO}}$ and $\dot{\mathbf{b}}_{\text{TO}} = \alpha \dot{\mathbf{x}}_{\text{TO}}$. The moment of take-off is defined as the point at which the contact forces vanish and constraint (4) holds. The scalar $\alpha \approx 1$ compensates for the small amount of potential energy that is stored in the ball-hand contacts and leads to slight excess ball take-off velocity $\|\dot{\mathbf{b}}_{\text{TO}}\| > \|\dot{\mathbf{x}}_{\text{TO}}\|$. To prevent post-take-off ball contact with parts of the hand by premature lateral movement, we require the relative acceleration between ball and hand to be collinear with the hand normal vector \mathbf{e}_h on a short interval $\tau \in (0, T_{\text{PTO}}]$ after take-off by

$$\mathbf{0} = (\ddot{\mathbf{x}}(\tau) - \mathbf{g}) \times \mathbf{e}_h. \quad (5)$$

Since planned cycles end at t_{TO} , this constraint acts on the initial interval of the planned movement.

B. Catching

To catch a ball, we predict the time \tilde{t}_{TD} and location $\tilde{\mathbf{b}}_{\text{TD}}$ of touch-down on a fixed horizontal plane and match the hand position $\mathbf{x}(\tilde{t}_{\text{TD}})$ accordingly

$$\mathbf{0} = \mathbf{x}(\tilde{t}_{\text{TD}}) - \tilde{\mathbf{b}}_{\text{TD}}. \quad (6)$$

To prevent pre-touch-down contact between the ball and the outer side of the hand, the relative velocity between ball and hand should be collinear with the hand normal vector on a short interval $\tau \in [0, T_{\text{PTO}}]$ before touch-down by

$$\mathbf{0} = (\dot{\mathbf{x}}(\tilde{t}_{\text{TD}} - \tau) - \dot{\tilde{\mathbf{b}}}(\tilde{t}_{\text{TD}} - \tau)) \times \mathbf{e}_h,$$

analogous to constraint (4) but in velocities, not in accelerations. Since we found this constraint challenging to optimize, we approximate it efficiently through

$$\mathbf{0} = \dot{\mathbf{x}}(\tilde{t}_{\text{TD}} - \tau) \times \dot{\tilde{\mathbf{b}}}(\tilde{t}_{\text{TD}} - \tau), \quad (7)$$

choosing a hand normal \mathbf{e}_h near collinear to the ball's predicted touch down velocity $\dot{\tilde{\mathbf{b}}}_{\text{TD}}$.




Fig. 3. **Pre-touch-down and post-take-off constraints:** (Left) After take-off, the relative hand acceleration $\ddot{\mathbf{x}} - \mathbf{g}$ and the hand normal \mathbf{e}_h are constrained to be collinear. (Right) Before touch-down, the hand velocity $\dot{\mathbf{x}}$ and the ball velocity $\dot{\mathbf{b}}$ are constrained to be collinear.

C. Trajectory Optimization

Given a desired take-off location $\mathbf{b}_{\text{TO,des}}$, position, velocity, and acceleration are fully constrained such that subsequent cycles of the juggling motion are independent of each other. Therefore, we formulate the mathematical program as a discrete-time trajectory optimization problem with piecewise constant jerk $\ddot{\mathbf{x}}_k \in \{\ddot{\mathbf{x}}_0, \dots, \ddot{\mathbf{x}}_{K_c-1}\}$ over one cycle, starting at the moment of take-off $t_{\text{TO}-1}$ of the previous cycle, while complying with constraints \mathbf{h} (2-7).

$$\min_{\{\ddot{\mathbf{x}}_0, \dots, \ddot{\mathbf{x}}_{K_c-1}\}} \sum_{k=0}^{K_c} \ddot{\mathbf{x}}_k^T \ddot{\mathbf{x}}_k \text{ s.t.} \quad (8)$$

$$(\mathbf{x}_0, \dot{\mathbf{x}}_0, \ddot{\mathbf{x}}_0) = (\mathbf{x}_{\text{TO}-1}, \dot{\mathbf{x}}_{\text{TO}-1}, \ddot{\mathbf{x}}_{\text{TO}-1})$$

$$\mathbf{h}(\mathbf{x}_k, \dot{\mathbf{x}}_k, \ddot{\mathbf{x}}_k, \mathbf{x}_k) = \mathbf{0}$$

$$-\ddot{x}_{i,\max} \leq \ddot{x}_{k,i} \leq \ddot{x}_{i,\max}$$

To avoid unattainable peak accelerations, we penalize the sum of squared accelerations and box-constrain the jerk $\ddot{\mathbf{x}}_k$ at each support point k and each dimension i . Planning in piece-wise constant jerk instead of piece-wise constant accelerations results in piece-wise linear accelerations. This allows us to define a precise time for the acceleration constraint (4) to hold despite discretization. Additionally, it allows us to increase the discretization time step Δt , reducing the number of decision variables. We solve the problem in CasADi [32] by direct shooting with interior point methods and the explicit Euler integration scheme

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k + \dot{\mathbf{x}}_k \Delta t + \frac{1}{2} \ddot{\mathbf{x}}_k \Delta t^2 + \frac{1}{6} \ddot{\mathbf{x}}_k \Delta t^3, \\ \dot{\mathbf{x}}_{k+1} &= \dot{\mathbf{x}}_k + \ddot{\mathbf{x}}_k \Delta t + \frac{1}{2} \ddot{\mathbf{x}}_k \Delta t^2, \\ \ddot{\mathbf{x}}_{k+1} &= \ddot{\mathbf{x}}_k + \ddot{\mathbf{x}}_k \Delta t. \end{aligned} \quad (9)$$

Since we integrate the kinematics, not the system dynamics, symplectic integration is unnecessary. The resulting piece-wise linear acceleration allows us to plan smooth trajectories despite a low number of support points $K_c < 30$. Additionally—in contrast to piece-wise constant accelerations—the take-off acceleration constraint (4) is only fulfilled exactly at t_{TO} , cleanly defining from which position \mathbf{x}_{TO} , the ball takes off and with which velocity $\dot{\mathbf{x}}_{\text{TO}}$.

IV. JUGGLING IN JOINT SPACE

A. Planning

To plan juggling movements in joint space, we formulate the trajectory optimization problem similar as in task space (8) and solve for a cycle of piece-wise constant joint space jerks $\ddot{\mathbf{q}}_k \in \{\ddot{\mathbf{q}}_0, \dots, \ddot{\mathbf{q}}_{K_c-1}\}$, integrating analogous to (9).

$$\begin{aligned} \min_{\{\ddot{\mathbf{q}}_0, \dots, \ddot{\mathbf{q}}_{K_c-1}\}} \quad & \sum_{k=0}^{K_c} \ddot{\mathbf{q}}_k^T \ddot{\mathbf{q}}_k \quad \text{s.t.} \\ (\mathbf{q}_0, \dot{\mathbf{q}}_0, \ddot{\mathbf{q}}_0) = & (\mathbf{q}_{\text{TO}-1}, \dot{\mathbf{q}}_{\text{TO}-1}, \ddot{\mathbf{q}}_{\text{TO}-1}) \\ \mathbf{h}(\mathbf{q}_k, \dot{\mathbf{q}}_k, \ddot{\mathbf{q}}_k, \mathbf{q}_k) = & 0 \\ -\ddot{q}_{i,\max} \leq \ddot{q}_{k,i} \leq & \ddot{q}_{i,\max} \end{aligned}$$

Before, the robot and ball states were defined in a shared euclidean space. Now the robot state is defined in the joint space. Therefore we need to reformulate constraints \mathbf{h} (2-7), which condition the robot's movement on the predicted and desired ball trajectories. First, we match the balls predicted touch-down location $\mathbf{x}_{\text{TD}} = \mathbf{b}_{\text{TD}}$ and desired take-off location $\mathbf{x}_{\text{TO}} = \mathbf{b}_{\text{TO}}$ as well as compute a desired take-off hand velocity by solving equation (3) for

$$\dot{\mathbf{x}}_{\text{TO}} = \frac{\mathbf{b}_{\text{TD}+1,\text{des}} - \mathbf{b}_{\text{TO}}}{\alpha T_f} - \frac{\mathbf{g} T_f}{2\alpha}.$$

Since \mathbf{x}_{TO} , $\dot{\mathbf{x}}_{\text{TO}}$, \mathbf{g} , and \mathbf{x}_{TD} are independent of the decision variables, we can solve the inverse kinematics problems

$$\begin{aligned} (\mathbf{q}_{\text{TO}}, \dot{\mathbf{q}}_{\text{TO}}, \ddot{\mathbf{q}}_{\text{TO}}) &= \mathbf{f}_{\text{kin}}^{-1}(\mathbf{x}_{\text{TO}}, \dot{\mathbf{x}}_{\text{TO}}, \mathbf{g}), \\ \mathbf{q}_{\text{TD}} &= \mathbf{f}_{\text{kin}}^{-1}(\mathbf{x}_{\text{TD}}), \end{aligned}$$

a priori. Now constraints (2-4) and (7) can be replaced by

$$\begin{aligned} (\mathbf{q}_{k\text{TO}}, \dot{\mathbf{q}}_{k\text{TO}}, \ddot{\mathbf{q}}_{k\text{TO}}) &= (\mathbf{q}_{\text{TO}}, \dot{\mathbf{q}}_{\text{TO}}, \ddot{\mathbf{q}}_{\text{TO}}), \\ \mathbf{q}_{k\text{TD}} &= \mathbf{q}_{\text{TD}}. \end{aligned}$$

Finally, we extend the post-take-off constraints (4) and the pre-touch-down constraints (7) by solving the forward kinematics at optimization time,

$$0 = \dot{\mathbf{x}}_{k\text{TD}-i} \times \mathbf{b}_{k\text{TD}-i}, \quad (10)$$

$$\dot{\mathbf{x}}_{k\text{TD}-i} = \mathbf{J}(\mathbf{q}_{k\text{TD}-i}) \dot{\mathbf{q}}_{k\text{TD}-i}, \quad (11)$$

$$0 = (\ddot{\mathbf{x}}_j - \mathbf{g}) \times \mathbf{e}_{\text{hand}}(\mathbf{q}_j), \quad (12)$$

$$\ddot{\mathbf{x}}_j = \mathbf{J}(\mathbf{q}_j) \ddot{\mathbf{q}}_j + \dot{\mathbf{J}}(\mathbf{q}_j, \dot{\mathbf{q}}_j) \dot{\mathbf{q}}_j, \quad (13)$$

with hand Jacobian \mathbf{J} at support points $i \in \{1, \dots, N_{\text{TD}}\}$ before touch-down and support points $j \in \{1, \dots, N_{\text{TO}}\}$ after the previous take-off. The previous take-off is at $k=0$, so the post-take-off constraints apply to the start of the movement.

B. Control

To execute the planned trajectories, a tracking controller is necessary. We consider a simple proportional derivative (PD) controller (14) with the options to add either compensate gravity (PD+G) (14) or add a feed-forward torque (PD+FF) (17) that is pre-computed from the planned trajectory through the inverse dynamics model $\mathbf{f}_{\text{dyn}}^{-1}$. Lastly, we test a joint space inverse dynamics controller (ID) (17) that also uses the same inverse dynamics model for error correction through a corrected reference acceleration.

$$\tau_{\text{PD}} = \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}) + \mathbf{K}_D(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) \quad (14)$$

$$\tau_{\text{PDG}} = \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}) + \mathbf{K}_D(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{g}(\mathbf{q}) \quad (15)$$

$$\tau_{\text{PDFF}} = \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}) + \mathbf{K}_D(\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{f}_{\text{dyn}}^{-1}(\mathbf{q}_d, \dot{\mathbf{q}}_d, \ddot{\mathbf{q}}_d) \quad (16)$$

$$\tau_{\text{ID}} = \mathbf{f}_{\text{dyn}}^{-1}(\mathbf{q}, \dot{\mathbf{q}}, \mathbf{K}_P(\mathbf{q}_d - \mathbf{q}) + \mathbf{K}_D(\dot{\mathbf{q}}_d - \dot{\mathbf{q}})) \quad (17)$$

V. EXPERIMENTS

As in previous work, we use low restitution balls and cone-shaped unactuated hands, with all contacts modeled as spring-damper systems in MuJoCo [33]. While this choice of a cone-shaped hand with 0.17m diameter provides advantageous passive stability [31], an actuated hand that fully opens at take-off can start moving sideways towards the following touch-down location earlier. This decreases the required vacant time T_v and average number of balls in air per hand W . In a simplified setup, we assume complete state control over two free-floating hands to verify the planned trajectories in task space independent of additional sources of errors like the inverse kinematics and tracking controllers. In the setup depicted in figure 1, the hands are attached to 4-DoF Barrett WAM arms which are torque controlled at 500Hz using the inverse dynamics controller (17) unless stated otherwise. The robot bases are mounted jointly, shaping a torso with a shoulder width of 0.76m. This configuration allows the elbows to drop down naturally and juggle in a human-typical pose, which is especially suitable for juggling since adaptations of the hand position in the horizontal catching plane only cause slight hand rotation.

While our method generalizes to fountain juggling, and we achieve stable juggling of six balls, the described unactuated hands are not suited for juggling even ball numbers requiring large carry distances d_d . Therefore we restrict the presented experiments to cascade juggling patterns with odd numbers of balls.

In the following sections, we conduct several experiments to measure the performance of our approach and test the relevancy of individual components of planner and controller. Runs are started in running patterns for each experiment, with most balls already in the air. Our primary metric to evaluate the stability of a juggling patten are the number of catches before the first ball drops, counting only balls that were thrown by the robot. As we do not count balls that were initialized in air, this metric is comparable to human juggling, typically starting with all balls held in hands.

A. Verifying Theoretical Limits

Even with total dexterity, the maximum amount of balls that can be juggled is limited by the distance between take-off and touch-down locations d_f through the kinematic upper bound (1). We validate this upper bound empirically in cascade patterns at dwell ratio $R = 0.5$ using the simplified setup with floating hands. As shown in figure 4, we achieve stable juggling for high amounts of balls up to 19 balls at the kinematic upper bound, defining stable as reaching 500 catches. For 21 and more balls at throw heights greater




Fig. 4. **Theoretical limit vs. achieved patterns:** The maximum number of balls is limited by the horizontal distance they travel. Stable juggling close to the upper bound was achieved in task space for more than 500 catches.




Fig. 5. **Robustness to disturbances:** Ball trajectories in cascade patterns were disturbed by isotropic white noise on the take-off velocity. Catches counted until 100 averaged over 50 runs, with and without ball collisions.

than 15m, throwing precision decreases due to unmodeled contact dynamics, and extra lateral travel d_f is required. On the anthropomorphic two-arm setup, we maximally achieve 50.000 catches with the maximum possible 17 balls at $d_f = 0.62\text{m}$, requiring balls to pass each other within a distance of 2.5mm.

B. Robustness to Disturbances

We disturbed each ball trajectory immediately after take-off by adding isotropic white noise to the velocity to evaluate the robustness of cascade patterns achieved with the two-arm setup. For each data point in figure 5, we counted the number of catches up to a maximum of 100 and averaged over 50 trials. As expected, higher numbers of balls are more susceptible to disturbances. To identify what causes balls to drop, we reran the entire experiment with all inter ball collisions turned off—balls tunneling through each other—indicated by the dashed lines. It is apparent that the importance of ball collisions increases with the number of balls, but their overall contribution to the degradation of robustness is only minor compared to the compounding turbulence caused by insufficiently dissipated kinetic energy.

C. Necessity of Trajectory Constraints

The additional constraints (10-13) are enforced on short intervals after take-off and before touch-down to ensure the successful execution of throws and catches. In the time-discrete case, these constraints are only implemented on




Fig. 6. **Necessity of trajectory constraints:** Alternating the number of constrained support points N_{TO} and N_{TD} for three different carry distances d_c in a five-ball cascade pattern shows the indispensability of the pre-touch-down and post-take-off constraints. The achieved catches were capped at 100 averaged over 50 runs with take-off ball disturbance of $\sigma = 0.01\text{m/s}$.




Fig. 7. **Evaluating tracking performance:** Evaluating the average error of touch-down locations in undisturbed five-ball cascade patterns (Left) Comparison of different tracking controllers show the need for inverse dynamics control. (Right) While ID control tracks closer, PD+FF control is less susceptible to varying uncertainties in the dynamics model.

the N_{TO} and N_{TD} support points within these intervals. To evaluate if these additional constraints are necessary and on which interval they need to hold, we again count catches until a maximum of 100, average over 50 trials, and disturb ball velocities by the isotropic white noise $\sigma = 0.01\text{m/s}$ immediately after take-off. We vary N_{TO} and N_{TD} between 0 and 2 while keeping the other constant at 2. Figure 6 shows that juggling a five-ball cascade pattern is infeasible without these additional constraints and that they become more critical at more considerable carry distances. Juggling with minimal carry distances, to avoid the additional constraints, would render the pattern susceptible to disturbances.

D. Tracking Control

We compared the tracking controllers (14-17) by evaluating them on the throwing accuracy, averaging the distance between planned and achieved touch-down locations. For the task of juggling five balls in an undisturbed cascade pattern, this average touch-down error is shown for each controller on the left side of figure 7. We show that the simple PD and PD+G controllers are at a significant disadvantage by not taking the robot's inertia into account. Juggling more than five balls with either of them was infeasible. While adding a feed-forward torque, the PD+FF controller halves the touch-down error, the ID controller nearly eliminates it.

Since the ground truth dynamics model of a physical system is never known, we tested the touch-down error of the PD+FF and the ID controller with respect to model errors juggling five and seven balls. We averaged the touch-down error of the first 20 catches over 50 different models with masses varied by standard deviations relative to the ground truth masses. The right side of figure 7 shows that the overall average touch-down error of ID-controlled throws is substantially lower but significantly more susceptible to model errors.

VI. CONCLUSIONS

We presented a novel decomposition of the infinite-horizon toss juggling movement into a sequential short-horizon trajectory optimization problem. Moreover, we identified several critical constraints necessary for dexterous dynamic manipulation with switching contacts. To the best of our knowledge, this formulation is the first to demonstrate stable juggling for five and more balls with anthropomorphic manipulators, reaching a kinematic upper bound of 17 balls.

We plan to evaluate our approach on a real-world physical robot. In prior work on a physical platform, we reached stable juggling of two balls in one hand, equivalent to juggling four balls in a fountains pattern. Since we did not reach the robot’s torque limit, a five-ball cascade pattern appears to be a realistic goal given the current hardware constraints. Our current choice of end-effector design restricts possible lateral movements with an object, and the possible types of object. Therefore, it may also be worthwhile to investigate different hand choices further. While we focused on the stationary infinite-horizon control problem of running a juggling pattern, experiments on a physical system will additionally require a solution to the transient short-horizon control problem of starting a juggling pattern.

Our approach is limited by the assumptions in Definition I.1. The restriction to non-zero dwell times in a) separates toss juggling from the different skill of paddle juggling. However, the three other assumptions lead to exciting directions for future work. Firstly, relaxing assumption b) and accounting for throws of different heights grants access to various juggling patterns mixing crossing and non-crossing throws. Secondly, by relaxing assumption c) and adapting take-off and touch-down positions, online collision avoidance between ball trajectories can be taken into account in real-time. Finally, breaking the constant juggling rhythm required in assumption d) and adapting the phase of the juggling pattern to one or more partners allows for interactive juggling—either by sharing a joint pattern with skilled partners or by having a novice partner add a ball to a running pattern by throwing it in, or removing a ball by snatching it out.

REFERENCES

- [1] M. T. Mason and K. M. Lynch, “Dynamic manipulation,” in *IOS 1993*, vol. 1. IEEE, 1993, pp. 152–159.
- [2] N. Furukawa, A. Namiki, S. Taku, and M. Ishikawa, “Dynamic regrasping using a high-speed multifingered hand and a high-speed vision system,” in *ICRA 2006*. IEEE, 2006, pp. 181–187.
- [3] H. Ha and S. Song, “Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding,” in *CoRL 2022*. PMLR, 2022, pp. 24–33.
- [4] C. E. Shannon, “Scientific aspects of juggling,” p. 924, 1993.
- [5] Y. Yam and J. Song, “Extending shannon’s theorem to a general juggling pattern,” *Studies in Applied Mathematics*, vol. 100, no. 1, pp. 53–66, 1998.
- [6] J. Buhler, D. Eisenbud, R. Graham, and C. Wright, “Juggling drops and descents,” *The American Mathematical Monthly*, vol. 101, no. 6, pp. 507–519, 1994.
- [7] P. J. Beek and A. Lewbel, “The science of juggling,” *Scientific American*, vol. 273, no. 5, pp. 92–97, 1995.
- [8] B. Polster, *The mathematics of juggling*. Springer, 2003.
- [9] A. Mays, “Combinatorial aspects of juggling,” *Praca magisterska, University of Melbourne*, 2006.
- [10] M. Buehler and D. E. Koditschek, “Robotics in an intermittent dynamical environment: A prelude to juggling,” *Departmental Papers*, 1987.
- [11] E. W. Aboaf, C. G. Atkeson, and D. J. Reinkensmeyer, “Task-level robot learning: Ball throwing,” *AI Memo 1006*, 1987.
- [12] M. Buehler, D. E. Koditschek, and P. J. Kindlmann, “A family of robot control strategies for intermittent dynamical environments,” *Control Systems Magazine*, vol. 10, no. 2, pp. 16–22, 1990.
- [13] M. Buehler, D. E. Koditschek, and P. J. Kindlmann, “Planning and control of robotic juggling and catching tasks,” *IJRR*, vol. 13, no. 2, pp. 101–118, 1994.
- [14] A. A. Rizzi and D. E. Koditschek, “Progress in spatial robot juggling,” in *ICRA 1992*. IEEE, 1992, pp. 775–780.
- [15] —, “Further progress in robot juggling: The spatial two-juggle,” in *ICRA 1993*. IEEE, 1993, pp. 919–924.
- [16] S. Schaal and C. G. Atkeson, “Open loop stable control strategies for robot juggling,” in *ICRA 1993*. IEEE, 1993, pp. 913–918.
- [17] S. Schaal, C. G. Atkeson, and D. Sternad, “One-handed juggling: A dynamical approach to a rhythmic movement task,” *Journal of Motor Behavior*, vol. 28, no. 2, pp. 165–183, 1996.
- [18] P. Reist and R. D’Andrea, “Design and analysis of a blind juggling robot,” *Transactions on Robotics*, vol. 28, no. 6, pp. 1228–1243, 2012.
- [19] T. Senoo, A. Namiki, and M. Ishikawa, “High-speed throwing motion based on kinetic chain approach,” in *IROS 2008*. IEEE, 2008, pp. 3206–3211.
- [20] E. W. Aboaf, C. G. Atkeson, and D. J. Reinkensmeyer, “Task-level robot learning,” in *ICRA 1988*. IEEE, 1988, pp. 1309–1310.
- [21] J. Kober, M. Glisson, and M. Mistry, “Playing catch and juggling with a humanoid robot,” in *Humanoids 2012*. IEEE, 2012, pp. 875–881.
- [22] R. Lampariello, D. Nguyen-Tuong, C. Castellini, G. Hirzinger, and J. Peters, “Trajectory planning for optimal robot catching in real-time,” in *ICRA 2011*. IEEE, 2011, pp. 3719–3726.
- [23] B. Bäuml, T. Wimböck, and G. Hirzinger, “Kinematically optimal catching a flying ball with a hand-arm-system,” in *IROS 2010*. IEEE, 2010, pp. 2592–2599.
- [24] B. Hove and J.-J. E. Slotine, “Experiments in robotic catching,” in *ACC 1991*. IEEE, 1991, pp. 380–386.
- [25] A. Namiki and N. Itoi, “Ball catching in kendama game by estimating grasp conditions based on a high-speed vision system and tactile sensors,” in *Humanoids 2014*. IEEE, 2014, pp. 634–639.
- [26] S. Kim, A. Shukla, and A. Billard, “Catching objects in flight,” *Transactions on Robotics*, vol. 30, no. 5, pp. 1049–1065, 2014.
- [27] C. Atkeson, “Shannon (robot) juggling: 3 and 5 balls,” *YouTube*: <https://youtu.be/dyc5bgpY86c>, 2017.
- [28] machines100, “Juggling machine,” *YouTube*: <https://youtu.be/kj36Z5ZIC6Q>, 2007.
- [29] T. Sakaguchi, Y. Masutani, and F. Miyazaki, “A study on juggling tasks,” in *IROS 1991*. IEEE, 1991, pp. 1418–1423.
- [30] T. Kizaki and A. Namiki, “Two ball juggling with high-speed hand-arm and high-speed vision system,” in *ICRA 2012*. IEEE, 2012, pp. 1372–1377.
- [31] K. Ploeger, M. Lutter, and J. Peters, “High acceleration reinforcement learning for real-world juggling with binary rewards,” in *CoRL 2020*. PMLR, 2021, pp. 642–653.
- [32] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “Casadi: A software framework for nonlinear optimization and optimal control,” *MPC*, vol. 11, no. 1, pp. 1–36, 2019.
- [33] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *IROS 2012*. IEEE, 2012, pp. 5026–5033.