Ball Juggling on the Bipedal Robot Cassie

Katherine L. Poggensee*, Albert H. Li*, Daniel Sotsaikich*,
Bike Zhang, Prasanth Kotaru, Mark Mueller, Koushil Sreenath

***Abstract*— The increasing integration of robots in daily life necessitates research in multitasking strategies. The act of juggling offers a simple platform to test techniques which may be generalizable to more complex tasks and systems. This paper presents both analytical and empirical results for successful ball-juggling on the bipedal robotic research platform Cassie. A control strategy inspired by mirror law algorithms was simulated on a simple paddle-ball system and then extended to the Cassie-ball system in simulations and experiments, using two low-level control schemes. A Poincaré analysis demonstrated stability for both controllers. Both simulated and experimental results show that the proposed strategy is robust to a wide range of physical parameters and that the act of juggling while balancing is achievable through multiple methods.**

I. Introduction

A. Motivation

As autonomous systems are integrated into daily life, the ability to safely and accurately perform complex dynamical tasks in uncertain environments becomes increasingly essential. The difficulty in simultaneously completing these tasks while also maintaining the safety of robots, humans, and the surrounding environment has acted as a barrier to the wider adoption of robotic technology, especially for mobile robots with the potential to cause injury or damage.

In light of these motivations, juggling can be a useful task to study. Unlike prehensile catching, juggling requires interaction with an object in free fall and repeated redirection of its trajectory to maintain a periodic orbit. The repeated interactions between the ball and robot can also affect the stability of the robot itself. Thus, juggling may provide useful analogues for more generalizable methods for dynamical multitasking in collaborative or populated environments.

B. Related Work

The simplest juggling system, a ball and a fixed planar paddle, has been used to develop controllers that often exhibit juggling motions similar to those observed in skilled human jugglers [2], [14]. These techniques have been further extended to 3-dimensional juggling [12], juggling of multiple balls [11], and open-loop juggling through clever design of a paddle with curved geometry [10].

* These authors contributed equally to this work.
K. L. Poggensee, A. H. Li, and D. Sotsaikich are with the Department of Mechanical Engineering, Stanford University, CA, 94305, USA, {ktpognc, ahli, dsot}@stanford.edu.
B. Zhang, P. Kotaru, M. Mueller, and K. Sreenath are with the Department of Mechanical Engineering, University of California, Berkeley, CA, 94720, USA, {bikezhang, prasanth.kotaru, mwm, koushils}@berkeley.edu.

Fig. 1. Snapshots of experiments on the Cassie-paddle-ball system. Two juggles, shown here, take 0.6 sec per bounce on average. Experimental videos are accessible at https://hybrid-robotics.berkeley.edu/cassie-juggling.

While the paddle-ball system is effective for studying the dynamics of ball bouncing, it is limited to the task of juggling. Recent work with robotic juggling has incorporated robots not designed specifically for this task. Quadrotors [7], [4], one degree-of-freedom robots [15], and robotic arms [13] have been used to juggle independently as well as cooperatively. Unlike the fixed-base robot in [13], quadrotor juggling control has the added complexity of maintaining its height while bouncing the ball along the desired trajectory.

C. Contributions

The focus of this paper is the implementation of juggling control strategies on the underactuated, bipedal robotic research platform Cassie, built by Agility Robotics. Cassie has been shown to be capable of performing highly dynamical tasks, e.g., walking on different terrains [5] and moving on hovershoes [3]. For this problem, Cassie must maintain balance while juggling using the same set of actuators. The contributions of this paper with respect to prior work are:

- the characterization of the system governing the interactions between the ball and robot-paddle system,
- stability analysis validating the controller designs,
- simulation of the robot-ball system using optimization-based and PD-based approaches to simultaneously balance on two feet while juggling, and
- experiments demonstrating the validity of the controller.

D. Organization

The paper is organized as follows. Section II introduces the dynamical models for a simple paddle-ball system which is then extended to the robot-ball system. Section III presents the controller design and describes two methods for maintaining balance while juggling, followed by an analysis of the controller stability in Section IV. Section V summarizes simulation results for both balance methods. Section VI details experiments on Cassie using one of the presented juggling controllers and analyzes deviations from the modeled behavior. The final section (VII) lists shortcomings with the current implementation, conclusions, and future work.

II. Hybrid Dynamical Model for Juggling

In this section, the dynamics for two systems – a simple paddle-ball and the Cassie-ball system – are described. The paddle-ball system was used to test the efficacy of the controller on a more intuitive platform, while the Cassie-ball system provided a realistic model for experiments.

A. Simple Paddle-Ball Model

The paddle was modeled as a floating rectangular rigid body in $SE(3)$ and the ball as a point mass in $\mathbb{R}^3$. The ball height and velocity relative to the paddle frame, $\zeta_b, \dot{\zeta}_b \in \mathbb{R}$, are

$$\zeta_b = (\mathbf{x_b} - \mathbf{x_p})^\top \mathbf{R} \mathbf{e_3} \tag{1}$$

$$\dot{\zeta}_b = \left((\dot{\mathbf{x}}_\mathbf{b} - \dot{\mathbf{x}}_\mathbf{p})^\top \mathbf{R} + (\mathbf{x_b} - \mathbf{x_p})^\top \mathbf{R} \hat{\mathbf{\Omega}} \right) \mathbf{e_3}, \tag{2}$$

where $\mathbf{x_b}, \mathbf{x_p} \in \mathbb{R}^3$ are the Cartesian positions of the ball and paddle respectively, $\mathbf{R} \in SO(3)$ is the rotation matrix from the paddle to the world frame, $\mathbf{\Omega} \in \mathbb{R}^3$ is the paddle angular velocity vector, $\hat{\cdot} : \mathbb{R}^3 \to SO(3)$ is the skew-symmetric operator, and $\mathbf{e_3} \in \mathbb{R}^3$ is the third canonical basis vector.

Contact occurs at the surface where $\zeta_p \equiv 0$ and can either be modeled as an instantaneous, rigid contact or as a compliant-contact model. While instantaneous contact models motivated the control architecture used in this paper, compliant contact better represents experimental conditions.

The compliant contact between the paddle and ball was modeled as a mass-spring-damper system. If $\zeta_p > 0$, no contact occurs, and the ball is only subject to gravitational forces. Adapting the methods in [1], the force on the ball in the paddle frame is given by

$$\mathbf{f_{n,b}} = \sqrt{\max\{-\zeta_p, 0\}} \left(-G_k \zeta_p - G_b \dot{\zeta}_p \right) \mathbf{e_3}, \tag{3}$$

where G_k, G_b are the paddle stiffness and damping coefficients. Contact is assumed to be frictionless because the expected lateral speed and rotation of the paddle were limited by the controller. The contact also results in an equal and opposite force, $-\mathbf{f_{n,b}}$, on the paddle as well as the induced moment, $-\mathbf{R}^\top(\mathbf{x_b} - \mathbf{x_p}) \times \mathbf{f_{n,b}}$.

Thus, for paddle inputs $\mathbf{f}, \mathbf{M} \in \mathbb{R}^3$ and the contact model from (3), the paddle-ball model is compactly written as

$$\Sigma_{pb} : \begin{cases} m_b \ddot{\mathbf{x}}_\mathbf{b} = \mathbf{R} \mathbf{f_{n,b}} - m_b g \mathbf{e_3} \\ m_p \ddot{\mathbf{x}}_\mathbf{p} = -\mathbf{R} \mathbf{f_{n,b}} + \mathbf{f} - m_p g \mathbf{e_3} \\ \mathbf{J_p} \dot{\Omega} = -\mathbf{R}^\top (\mathbf{x_b} - \mathbf{x_p}) \times \mathbf{f_{n,b}} \\ \qquad\quad + \mathbf{M} - \mathbf{\Omega} \times \mathbf{J_p} \mathbf{\Omega}, \end{cases} \tag{4}$$

where $\mathbf{J_p}$ is the paddle inertia matrix, m_p, m_b the paddle and ball masses, and g the gravitational acceleration constant.

B. Full Robot-Ball Model

The dynamical model of Cassie is described in detail in [5]. The generalized coordinates contain the position and orientation of the pelvis and the angular positions of the seven joints for the left and right legs, respectively, written as

$$\begin{aligned} \mathbf{q_{robot}} := [&q_x, q_y, q_z, q_{yaw}, q_{pitch}, q_{roll}, \\ &q_{1L}, q_{2L}, q_{3L}, q_{4L}, q_{5L}, q_{6L}, q_{7L}, \\ &q_{1R}, q_{2R}, q_{3R}, q_{4R}, q_{5R}, q_{6R}, q_{7R}]^\top. \end{aligned} \tag{5}$$

The seven joint coordinates refer to the hip roll, hip yaw, hip pitch, knee pitch, shin pitch, tarsus pitch, and toe pitch respectively. Only q_1, q_2, q_3, q_4, and q_7 are actuated, while q_5 and q_6 are passive, moving subject to springs in the leg.

Cassie is modeled with two contact points per foot, where the force at each point is modeled like the stick/slip model in [1]. The normal ground deformation during contact, ζ_g, is defined similarly to ζ_p in the paddle-ball model. The magnitude of the foot-ground contact force, $f_{n,g}$, is likewise computed similarly to $\|\mathbf{f_{n,b}}\|$.

For ground stiffness and damping coefficients Γ_k, Γ_b, tangential ground deformation ξ_g, coefficient of friction μ, and planar foot velocity $\dot{\mathbf{x}}_\mathbf{foot} \in \mathbb{R}^2$, the stick/slip forces tangential to the ground are $\mathbf{f_{stick}}, \mathbf{f_{slip}} \in \mathbb{R}^2$, given by

$$\mathbf{f_{stick}} = \sqrt{\max\{-\zeta_g, 0\}} \left(-\Gamma_k \xi_g - \Gamma_b \dot{\mathbf{x}}_\mathbf{foot} \right) \tag{6}$$

$$\mathbf{f_{slip}} = \mathbf{f_{stick}} \times \frac{\mu f_{n,g}}{\|\mathbf{f_{stick}}\|}, \tag{7}$$

where $\mathbf{f_{slip}}$ is only computed for nonzero $\mathbf{f_{stick}}$. The frictional force $\mathbf{f_{friction}} = \mathbf{f_{slip}}$ when $||f_{stick}|| > \mu f_{n,g}$, and $\mathbf{f_{friction}} = \mathbf{f_{stick}}$ otherwise. The total contact force is then

$$\mathbf{f_C} = \begin{bmatrix} \mathbf{f_{friction}} \\ f_{n,g} \end{bmatrix}. \tag{8}$$

Because the paddle was mounted rigidly to Cassie's pelvis, the transformation from the robot to the paddle is known and the paddle-ball model was easily integrated with the robot model. Thus, using the same implicit switching condition as in (4), the complete Cassie-ball dynamics are

$$\Sigma_{cb} : \begin{cases} m_b \ddot{\mathbf{x}}_\mathbf{b} = \mathbf{R} \mathbf{f_{n,b}} - m_b g \mathbf{e_3} \\ \mathbf{M}(\mathbf{q}) \ddot{\mathbf{q}} + \mathbf{H}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{B}\mathbf{u} + \mathbf{J_s}^\top(\mathbf{q}) \tau_\mathbf{s} \\ \qquad\qquad + \mathbf{J_C}^\top(\mathbf{q}) \mathbf{f_C} \\ \qquad\qquad + \mathbf{J_b}^\top(\mathbf{q}) \mathbf{f_{n,b}}, \end{cases} \tag{9}$$

where $\mathbf{M}(\mathbf{q})$ is the mass matrix, $\mathbf{H}(\mathbf{q}, \dot{\mathbf{q}})$ contains the Coriolis and gravity terms, $\mathbf{B}$ is the motor torque matrix, $\mathbf{u} \in \mathbb{R}^{10}$ is the vector of the input torques for all actuated joints, $\mathbf{J_s}(\mathbf{q})$ is the spring coordinate Jacobian, $\tau_\mathbf{s}$ is the spring torque vector, $\mathbf{J_C}(\mathbf{q})$ is the ground contact position Jacobian, and $\mathbf{J_b}(\mathbf{q})$ is the ball contact position Jacobian.

III. Feedback Control for Juggling

The control strategy based on the dynamics from the previous section was inspired by *mirror algorithms*, a family of control algorithms that exhibit stable juggling of a ball to a desired apex [2].

This strategy employs a transformation between the ball trajectory and the desired paddle trajectory. In successful juggling motions, e.g., in human juggling [14], the paddle moves downward mirroring the upward motion of the ball and vice versa, with the magnitude scaled based on the error in the desired energy of the ball. This energy-based method helps to "servo energy" into the ball as described in [2] for a simple actuated paddle. It was also implemented in [9] for robotic walking. In addition to the vertical motion of the paddle, its attitude may be independently adjusted to bounce

the ball to reach desired planar positions. Hence juggling can be accomplished by constraining the desired pose of the paddle to be a function of the current and desired position of the ball. This is similar to virtual constraints [5] where a wrench on the paddle enforces the virtual constraint.

First validated on the simple paddle-ball system described in (4), this strategy was then implemented on the Cassie-ball system, described in (9), in simulation using a ground contact force optimization-based scheme presented in [8]. Another controller was developed as a benchmark for the contact force optimization controller, using the same high-level juggling controller with a different low-level motor torque controller modified from the balancing controller presented in [5].

A. Juggling Controller

The juggling controller carries out two goals: to control the vertical motion of the paddle to modulate the height of the ball and to control the attitude of the paddle to move the ball to a desired planar position. The output of the controller is a desired wrench which is applied to the paddle.

The force component is the sum of a feedforward term $\mathbf{f_{ff}} = -k_{ff,\eta}\, m_p\, g$ to compensate for the weight of the paddle and a feedback term, $\mathbf{f_{fb}}$, such that

$$f^x_{fb} = -k^x_P\ e^x_p - k^x_D\ \dot{e}^x_p \tag{10}$$

$$f^y_{fb} = -k^y_P\ e^y_p - k^y_D\ \dot{e}^y_p \tag{11}$$

$$\begin{aligned} f^z_{fb} = &-k^z_P \left(e^z_p + k_W(x^z_b - z_C)\right) \\ &- k^z_D \left(\dot{e}^z_p + k_W \dot{x}^z_b\right), \end{aligned} \tag{12}$$

where $\mathbf{e_p}$ is defined as the position error between the center of mass of the paddle and the desired contact position, z_C the desired contact height, and $\mathbf{x_b}, \mathbf{x_p}$ the current ball and paddle positions. In practice, the desired contact position was chosen to keep the paddle under the desired apex.

The gain $k_W = k_{ff,\eta} - k_{P,\eta}\ (\eta_{des} - \eta)$ is also the sum of a feedforward term, $k_{ff,\eta}$, the nominal relation between the ball and paddle energies, and a feedback term, $-k_{P,\eta}\ (\eta_{des} - \eta)$, which changes proportionally to the error in the energy of the ball. Only the vertical mechanical energy was used in the calculation of η, η_{des}, which are then written as

$$\eta_{des} = -m_b\, g\, x^z_{apex} \tag{13}$$

$$\eta = -m_b\, g\, x^z_b + \frac{1}{2} m_b \left(\dot{x}^z_b\right)^2. \tag{14}$$

The moment to be applied to the paddle drives the paddle orientation $\mathbf{R}$, with respect to the inertial frame, to the desired orientation $\mathbf{R_{des}}$, and is likewise the sum of a feedforward term, $\mathbf{M_{ff}}$, and a feedback term, $\mathbf{M_{fb}}$, where

$$\begin{aligned} \mathbf{M_{ff}} = &\boldsymbol{\Omega} \times \mathbf{J_p}\boldsymbol{\Omega} \\ &- \mathbf{J_p}(\hat{\boldsymbol{\Omega}}\mathbf{R}^\top \mathbf{R_{des}}\boldsymbol{\Omega_{des}} - \mathbf{R}^\top \mathbf{R_{des}}\dot{\boldsymbol{\Omega}}_{\mathbf{des}}) \end{aligned} \tag{15}$$

$$\mathbf{M_{fb}} = -k_R\, \mathbf{e_R} - k_\Omega\, \mathbf{e_\Omega}. \tag{16}$$

Since the robot-paddle transformation is a translation, the pelvis and paddle orientation are equivalent. The desired pitch and roll are computed by PD for ball error $\mathbf{e_b}$ between the ball position $\mathbf{x_b}$ and the desired apex position $\mathbf{x}_{apex}$ as

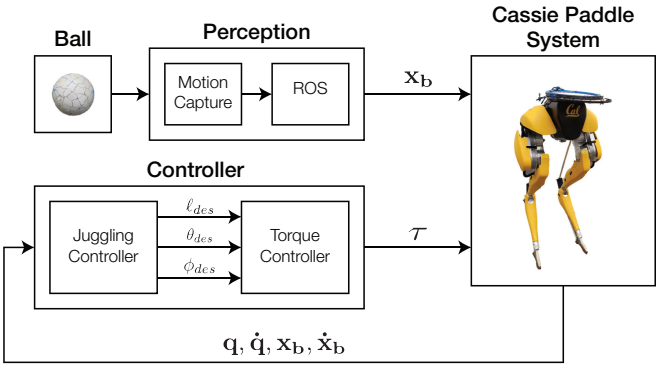

Fig. 2. Control framework relating the high-level juggling controller on the paddle, the low-level torque controller on Cassie, and the perception via motion capture.

$$\phi_{des} = k_{P,\phi}\ e^y_b + k_{D,\phi}\ \dot{e}^y_b \tag{17}$$

$$\theta_{des} = -(k_{P,\theta}\ e^x_b + k_{D,\theta}\ \dot{e}^x_b). \tag{18}$$

The desired rotation matrix is then defined as

$$\mathbf{R_{des}} = \mathbf{R_y}(-\phi_{des})\ \mathbf{R_x}(\theta_{des}), \tag{19}$$

where $\mathbf{R_x}(\cdot)$ and $\mathbf{R_y}(\cdot)$ are rotation matrices corresponding to rotations about the x and y axes, respectively.

The moment feedback in (16) is derived from a feedback law presented in [6], where $\boldsymbol{\Omega}$ is the paddle angular velocity vector and the attitude errors are

$$\mathbf{e_R} = \frac{1}{2}(\mathbf{R_{des}}^\top \mathbf{R} - \mathbf{R}^\top \mathbf{R_{des}}) \tag{20}$$

$$\mathbf{e_\Omega} = \boldsymbol{\Omega} - \mathbf{R}^\top \mathbf{R_{des}}\boldsymbol{\Omega_{des}}. \tag{21}$$

This high-level control architecture for regulating the pose of the paddle is implemented using two different low-level controllers on Cassie suitable for multitasking juggling and balancing, described in the following sections.

B. Contact Force Optimization Controller

For the desired wrench on the paddle described in (10-12) and (15-16), juggling while balancing was achieved by optimizing ground contact forces. Specific details of the optimization-based controller are found in [8], but the salient points and modifications are summarized in this section.

Grasping methods consider the net wrench produced by a set of contact points on an object. Each of these forces must satisfy the positivity restriction, i.e., the contact points push but do not pull. Contact is also governed by a Coulombic friction model. The resulting set of allowable contact forces at point i creates a friction cone centered about the axis along the surface normal:

$$\mathscr{F}_i = \left\{ \mathbf{f}_i \in \mathbb{R}^3 \,\middle|\, \sqrt{(f^x_i)^2 + (f^y_i)^2} \le \mu f^z_i, f^z_i \ge 0 \right\}. \tag{22}$$

A grasp map $\mathbf{G}$ is computed to transform the vector of contact forces on the object $\mathbf{f_C}$ to a resultant wrench $\mathbf{f}$. Each column of $\mathbf{G}$ is a contact map that transforms the contact

force at a single point to a wrench in the object coordinate frame [8]. The total wrench on the object $\mathbf{f}$ is thus $\mathbf{G}\mathbf{f}_\mathbf{C}$.

The contact forces $\mathbf{f}_\mathbf{C}$ can be computed as the product of the pseudoinverse of the grasp map $\mathbf{G}^\dagger$ and the desired wrench $\mathbf{f}$. In general, this calculation does not guarantee the aforementioned constraints are satisfied, so a multi-objective optimization problem is solved subject to these constraints. The first two terms (J_1, J_2) ensure forces and moments generated on the paddle from the contact forces at the feet are close to the desired values computed by the high-level controller, while the third (J_3) minimizes the magnitude of the contact force, and are written as

$$J_1 = \left\| \begin{bmatrix} \mathbf{I}_3 & \mathbf{0}_{3\times 3} \end{bmatrix} (\mathbf{f} - \mathbf{G}\mathbf{f}_\mathbf{C}) \right\|_2^2 \tag{23}$$

$$J_2 = \left\| \begin{bmatrix} \mathbf{0}_{3\times 3} & \mathbf{I}_3 \end{bmatrix} (\mathbf{f} - \mathbf{G}\mathbf{f}_\mathbf{C}) \right\|_2^2 \tag{24}$$

$$J_3 = {\mathbf{f}_\mathbf{C}}^\top \mathbf{f}_\mathbf{C}. \tag{25}$$

For weights $\alpha_i > 0$, a single objective function is written as

$$J = \alpha_1 J_1 + \alpha_2 J_2 + \alpha_3 J_3. \tag{26}$$

To prioritize balance, $\alpha_1 \gg \alpha_2 \gg \alpha_3$.

Finally, contact wrenches $\mathbf{f}_k$ at each contact point k are constructed using these forces, which then allow the computation of the joint torques $\boldsymbol{\tau}$ required to achieve the desired contact forces as

$$\boldsymbol{\tau} = {\mathbf{J}_\mathbf{C}}^\top \mathbf{f}_\mathbf{C} \tag{27}$$

C. PD Torque Controller

While the presented contact force optimization controller works in simulation, a simpler joint PD controller from [5] was also explored for this task. As in the optimization-based controller, the attitude of the paddle locates the ball in the plane while the vertical motion of the paddle is used to regulate the energy of the ball. The desired attitude is controlled according to (17), (18).

In contrast to the optimization-based controller, the ground contact forces are computed to modulate the length ℓ of the virtual leg, defined from Cassie's hip to its toe [5], to control the vertical height of the ball. For a desired contact height z_C as in (12) and vertical offset between the paddle and hip h_{off}, the desired leg length is

$$\ell_{des} = -k_W(x_b^z - z_C) - h_{off} + z_C. \tag{28}$$

The motor torques can then be computed using the balancing controller presented in [5] with the modification that the control law for both hip pitch motors, q_3, is

$$q_3 = -k_{P,3}(\phi - \phi_{des}) - k_{D,3}\dot{\phi}. \tag{29}$$

Balance was maintained by setting the virtual leg angle, i.e., the angle between the vertical axis and the virtual leg, to zero. Thus, the center of mass of the robot was constrained over the support polygon formed by the feet placed squarely on the floor.

The stability of these controllers is detailed in the next section, followed by simulation and experimental results.

IV. Stability Analysis

Stability of the controller described in the previous section was determined by Poincaré analysis of the Cassie-ball system. A Poincaré section $\mathcal{S}$ was defined at the apex of the ball trajectory, i.e., $\dot{x}_b^z \equiv 0$. The return map $P : \mathcal{S} \to \mathcal{S}$ is defined by the continuous dynamics described in (9).

By choice of the Poincaré section, the system passes through three phases: free-fall of the ball until impact with the paddle, collision dynamics between the ball and the paddle, and the ball leaving the surface of the paddle until it reaches its apex. The complexity of modeling the system dynamics of the robot and the effects of repeated interactions between the robot and ball required simulating the Cassie-ball dynamics to identify a periodic orbit for a chosen x_{apex}^z.

Kinematic constraints on the system resulted in a lower-dimensional return map. Individual states were first perturbed and then the constraints were enforced before running the forward simulation. Of the 54 states, 40 describe Cassie, 6 describe the motion of the ball, and 8 determine deformation of the ground and paddle. Since a state on the Poincaré section must satisfy several constraints, such as the ball being at its apex, feet being in contact with the ground, and the paddle not being in contact with the ball, the number of perturbable states was reduced to 17. Therefore, the Poincaré map P is reduced to

$$x[k+1] = P(x[k]), \tag{30}$$

where $x \in \mathbb{R}^{17}$ is the vector of perturbable states. This map was computed using the finite differences method about the chosen fixed point for both juggling controllers. The largest eigenvalue magnitudes for the optimization controller and for the PD controller were 0.65 and 0.84, respectively. Thus both controllers are locally exponentially stable.

V. Simulations

To evaluate these controllers, juggling was simulated on the robot-ball system. All simulation results presented here were performed in `Matlab R2019b` on a Lenovo Thinkpad X1 Carbon laptop (Intel Core i7 CPU, 16 GB RAM).

For stationary juggling, only the position states of the ball relative to the position of the paddle were varied. For an inertial frame set on the ground below the paddle center, the chosen ball initial conditions in meters were $(0.05, 0.075, 1.3)$, $(0.05, 0.075, 1.0)$, $(0.05, 0.075, 1.6)$. Of the many initial conditions tested, these were deemed sufficient for verifying the robustness of the juggling controller. The desired apex was $(0, 0, 1.3)$, so the proposed controllers were tested for initial configurations when the ball was displaced from the apex position along all three axes.

For both controllers, the resulting trajectories converged to a stable juggling pattern on the order of 10 bounces (Fig. 3). Over a four-second period, all trajectories generated with the optimization controller converged to periodic bouncing within $8\,\mathrm{cm}$ of the desired apex height and within $2\,\mathrm{cm}$ of the desired planar x and y positions. Over the same period, all trajectories generated using the PD controller converged

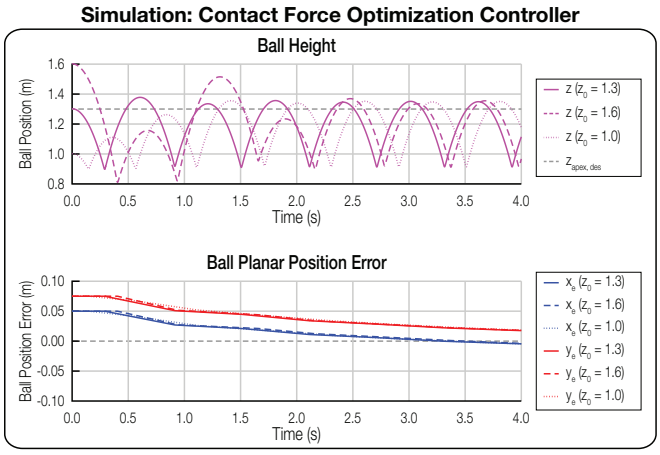

Fig. 3. Simulation of the Cassie-ball system with the contact force optimization controller. The ball is dropped from three different initial heights ($z_0 = 1.0, 1.3, 1.6$m). In all three cases, the desired apex height is $1.3\,\mathrm{m}$, the ball apex settles to within $0.08\,\mathrm{m}$ of the apex, and the planar ball position converges to zero. The simulation with the PD controller shows similar results. The subscript "e" denotes error.

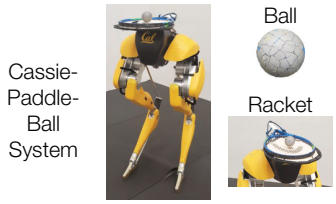

Fig. 4. Experimental set-up for the Cassie-paddle-ball system. The Cassie robot is the bipedal robotic research platform developed by Agility Robotics. The racket is a standard tennis racket. The ball is a standard racquetball covered in reflective tape.

to periodic bouncing within $2\,\mathrm{cm}$ of the desired apex height and within $3\,\mathrm{cm}$ of the desired planar positions.

For both controllers, periodic juggling was also achieved when the mass of the ball was varied from $0.04\,\mathrm{kg}$, the mass of a racquetball, to $0.8\,\mathrm{kg}$, twice the mass of a soccer ball. Other physical parameters including the stiffness and damping coefficients for contact and the friction coefficient for robot-ground contact were evaluated.

Surprisingly, the contact force optimization controller produced similar or less well-controlled trajectories than the PD controller. This may be due to difficulties interpreting the contact force optimization controller gains, which include all parameters related to the optimization algorithm and gains governing the feedback forces and moments along all axes. In contrast, the PD controller was governed by fewer gains, so brief tuning produced desirable results. Thus, the PD controller was used in experiments. However, performance for both controllers was still sensitive to gain variations, and the large dimensionality of the parameter spaces prevented an exhaustive gain optimization.

VI. Juggling Experiments on Cassie

Having demonstrated success in simulation, the PD controller presented in Section III was implemented on Cassie.

A. Setup

A tennis racket rigidly fixtured onto the pelvis was chosen to juggle a racquetball of mass $40\,\mathrm{g}$, shown in Fig. 4. The rigidly-mounted racket demonstrated isotropic contact properties at a low weight. Several trials were also conducted using other paddles. A rectangular 3D-printed paddle was tested, but exhibited a poor coefficient of restitution between the ball and paddle and poor isotropic properties along its planar axes. Wooden blocks were used in an attempt to promote more uniform contact properties, but exhibited poor flatness while greatly increasing the weight. Despite these inconsistencies, each paddle achieved at least ten juggles, demonstrating the controller's robustness to uncertainties in physical properties.

The ball was sensed using nine OptiTrack Prime 17W motion capture cameras which published an estimated position through ROS to the computer onboard Cassie via UDP at a rate of $200\,\mathrm{Hz}$. The ball was covered in 3M 7610 reflective motion capture tape, which slightly decreased its coefficient of restitution. The control loop ran onboard Cassie in Simulink Real-Time at a rate of $2000\,\mathrm{Hz}$.

For each experiment, the desired planar ball position was set to a point above the planar center of the paddle, and juggling was initiated by releasing the ball from rest.

B. Results

The planar position of the ball was maintained in a stable region for multiple bounces, although the desired apex was typically overshot. The typical failure mode was the ball drifting outside of the planar region Cassie could reach while maintaining its balance.

Juggles of twenty, thirty, and forty bounces were achieved several times. In the most successful observed run, Cassie was able to juggle the ball 42 times (Figs. 5, 6). This run failed due to ball-racket contact on the raised edge of the racket frame, which was considered acceptable since this geometry was not modeled in simulations. Other runs of over forty juggles showed a similar failure mode.

VII. Conclusion

This paper has presented an implementation of juggling control strategies on the bipedal robot Cassie and addressed the challenge of balancing while juggling a ball to a desired periodic orbit. After guaranteeing stability and validating the presented strategy in simulation, a juggling controller was tested online, where Cassie achieved over 40 bounces.

The proposed strategy is one of the simplest for juggling. This simplicity resulted in instances of undesirable behaviors that were addressed through careful gain tuning but could be handled with a more complex control strategy. For example, the attitude of the paddle is linear with the planar position and velocity errors, which is only roughly accurate for small displacements.

Motion capture was the primary method of perception, which limited experiments to motion capture spaces. Other means of perception, such as on-board depth cameras,

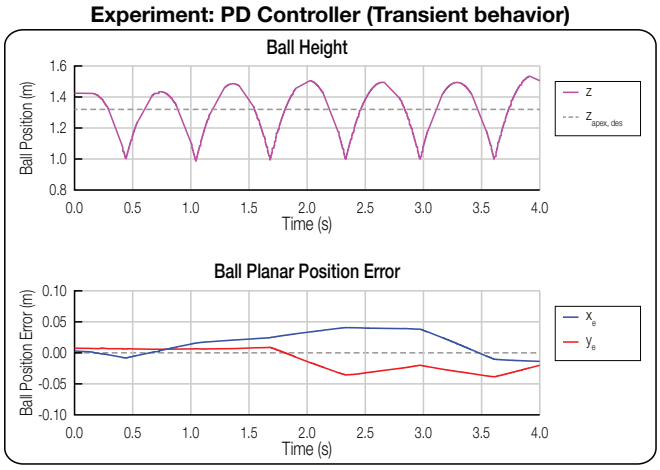

Fig. 5. Transient behavior for the PD controller in experiments. Over a 4-second window, the ball planar position errors exhibited similar convergence behavior to the simulations and consistently overshot the apex.

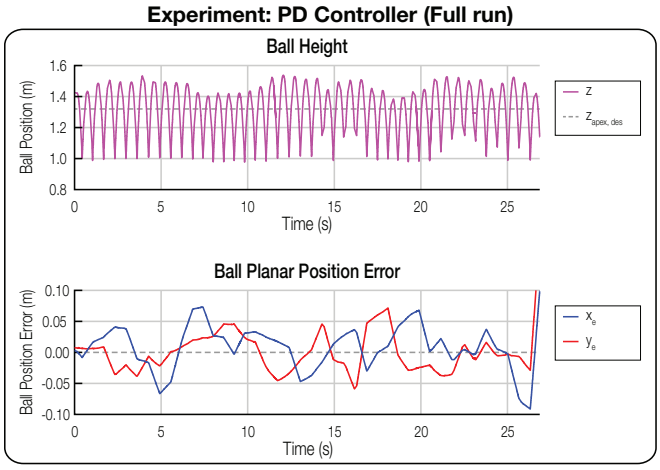

Fig. 6. Experimental results for a 42-bounce run. The peak planar position error was 7 cm, well within the confines of the racket, and the average apex overshoot was 13.8 cm. The spike in planar error corresponds to the failure mode of this run: the ball hit the frame of the racket and bounced off.

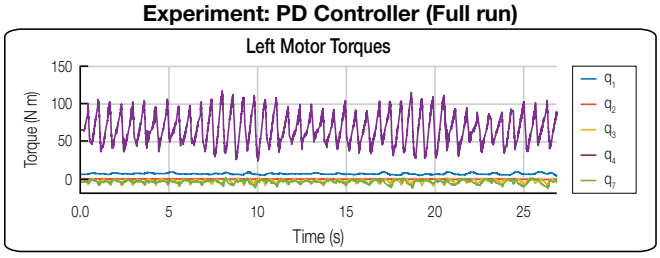

Fig. 7. Motor torque commands during a run with 42 bounces for the left leg; the right leg produced similar results. The knee flexion joint (q_4) requires the highest torque because it controls the leg length.

RGB cameras, or LiDAR, would allow juggling in any setting. Furthermore, our controller assumes the robot remains standing, limiting the range within which Cassie can juggle to the base of support. By implementing juggling while walking, Cassie could juggle more complicated trajectories and recover from more dynamic initial conditions.

Future work includes extending this control strategy to allow juggling while walking or juggling multiple balls. Collaborative juggling with moving humans or robots may also be a possible avenue for exploration and could incorporate more novel decision-making strategies.

VIII. Acknowledgments

We thank Shuxiao Chen and Philipp Wu in assisting with the experimental setup. We thank Chenxi (Toby) Tian for creating the first version of the paddle for juggling. This work was supported by National Science Foundation Grant IIS-1834557 and Graduate Research Fellowship DGE-1252522.

References

[1] M. Azad and R. Featherstone, "Modeling the contact between a rolling sphere and a compliant ground plane," *Australasian Conference on Robotics and Automation, Brisbane, Australia*, 2010.

[2] M. Bühler, D. E. Koditschek, and P. J. Kindlmann, "A family of robot control strategies for intermittent dynamical environments," *IEEE Control Systems Magazine*, vol. 10, no. 2, pp. 16–22, 1990.

[3] S. Chen, J. Rogers, B. Zhang, and K. Sreenath, "Feedback control for autonomous riding of hovershoes by a Cassie bipedal robot," in *IEEE International Conference on Humanoid Robots (Humanoids)*, Toronto, Canada, 2019, pp. 375–382.

[4] W. Dong, G.-Y. Gu, Y. Ding, X. Zhu, and H. Ding, "Ball juggling with an under-actuated flying robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Hamburg, Germany, 2015, pp. 68–73.

[5] Y. Gong, R. Hartley, X. Da, A. Hereid, O. Harib, J.-K. Huang, and J. Grizzle, "Feedback control of a Cassie bipedal robot: Walking, standing, and riding a segway," in *American Control Conference (ACC)*, Philadelphia, PA, USA, 2019, pp. 4559–4566.

[6] T. Lee, M. Leok, and N. H. McClamroch, "Control of complex maneuvers for a quadrotor UAV using geometric methods on SE(3)," *Conference on Decision and Control*, 2010.

[7] M. Müller, S. Lupashin, and R. D'Andrea, "Quadrocopter ball juggling," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 5113–5120.

[8] C. Ott, M. A. Roa, and G. Hirzinger, "Posture and balance control for biped robots based on contact force optimization," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2011, pp. 26–33.

[9] M. Raibert, *Legged Robots that Balance*. MIT Press, 1986.

[10] P. Reist and R. D'Andrea, "Design and analysis of a blind juggling robot," *IEEE Transactions on Robotics*, vol. 28, no. 6, pp. 1228–1243, 2012.

[11] A. A. Rizzi and D. E. Koditschek, "Further progress in robot juggling: The spatial two-juggle," in *IEEE International Conference on Robotics and Automation*. IEEE, 1993, pp. 919–924.

[12] A. A. Rizzi, L. L. Whitcomb, and D. E. Koditschek, "Distributed real-time control of a spatial robot juggler," *Computer*, vol. 25, no. 5, pp. 12–24, 1992.

[13] D. Serra, F. Ruggiero, V. Lippiello, and B. Siciliano, "A nonlinear least squares approach for nonprehensile dual-hand robotic ball juggling," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11 485–11 490, 2017.

[14] D. Sternad, H. Katsumata, M. Duarte, and S. Schaal, "Bouncing a ball: Tuning into dynamic stability," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 27, no. 5, pp. 1163–1184, 2001.

[15] A. Zavala-Rio and B. Brogliato, "On the control of a one degree-of-freedom juggling robot," *Dynamics and Control*, vol. 9, pp. 67–90, 1999.