

Feedback Control for Autonomous Riding of Hovershoes by a Cassie Bipedal Robot

Shuxiao Chen*, Jonathan Rogers*, Bike Zhang*, and Koushil Sreenath

Abstract—Motivated towards achieving multi-modal locomotion, in this paper, we develop a framework for a bipedal robot to dynamically ride a pair of Hovershoes over various terrain. Our developed control strategy enables the Cassie bipedal robot to interact with the Hovershoes to balance, regulate forward and rotational velocities, achieve fast turns, and move over flat terrain, slopes, stairs, and rough outdoor terrain. Our sensor suite comprising of tracking and depth cameras for visual SLAM as well as our Dijkstra-based global planner and timed elastic band-based local planning framework enables us to achieve autonomous riding on the Hovershoes while navigating an obstacle course. We present numerical and experimental validations of our work.

I. INTRODUCTION

While locomotion using legs is efficient when traveling over rough and discrete terrain, wheeled locomotion is more efficient when traveling over flat continuous terrain [7]. Humans are able to optimize locomotion efficiency by using multiple locomotion modalities that comprise of not only being able to walk and run, but also being able to ride various micro-mobility platforms, such as Segways and Hovershoes. Enabling legged robots to autonomously ride on various personal mobility platforms will offer multi-modal locomotion capabilities, improving the efficiency of locomotion over various terrains.

Autonomous robots with multi-modal locomotion capabilities can have a big impact in the real-world from package delivery to security and surveillance to search and rescue missions [16]. In order to address this problem, we use the Cassie bipedal robot developed by Agility Robotics to autonomously ride Hovershoes, see Fig. 1.

A. Challenges

There are multiple challenges in developing an autonomous framework for Cassie to ride Hovershoes. First, Cassie is a complex, underactuated robot with 20 degrees-of-freedom (DOFs), making it a difficult system to control. Second, the Hovershoes are a sensitive and highly dynamic platform that serve as two decoupled moving platforms under each of Cassie’s feet and are hard to coordinate. Third, the algorithms we can develop neither have access to the internal states of the Hovershoes nor can they directly specify



Fig. 1: Snapshot of an experiment of the bipedal robot Cassie autonomously riding on Hovershoes. Experimental videos are at <https://youtu.be/b2fKBb.0iTo>.

the torque inputs on the Hovershoes, since Cassie can only indirectly interact with the Hovershoes through the contact forces. Fourth, the Cassie-Hovershoes system needs to be robust to kinematic and dynamic variances arising from (i) manual initialization of the Hovershoes that result in an initial relative translation and orientation between the Hovershoes; (ii) manual placement of Cassie’s feet on the Hovershoes resulting in off-center foot placement; and (iii) differences in the dynamic characteristics between the Hovershoes due to manufacturing. Fifth, the Cassie-Hovershoes system needs to be able to autonomously avoid obstacles, resulting in fast perception and planning demands.

We solve these challenges by designing controllers that assume the dynamics are decoupled along various degrees-of-freedom, strategically choosing a minimalist sensor suite for state estimation, and assuming kinematics is sufficient for path planning. This motivates the simplest approach to solving the problem - using PD controllers on decoupled dynamics, using a VIO and depth sensor for state estimation and object detection, and simple global and local kinematic planners. In the future, we will look at more complex controllers, estimators and planners.

B. Related Work

1) *Multi-Modal Locomotion*: There are some examples of multi-modal locomotion. For instance, DRC-HUBO+ is a bipedal robot that can have wheeled attachments at the knee joint of each leg [12], enabling both legged and wheeled locomotion. Boston Dynamics’ Handle robot—designed for box handling in warehouses—is another legged robot with wheels for feet [1]. ANYmal is yet other wheeled-legged

* These authors contributed equally to this work.

This work was supported by National Science Foundation Grant IIS-1834557 and Berkeley Deep Drive.

S. Chen, J. Rogers, B. Zhang and K. Sreenath are with the Department of Mechanical Engineering, University of California, Berkeley, CA, 94720, USA, {shuxiao.chen, jr1397, bikezhang, koushils}@berkeley.edu



Fig. 2: Segway, Hoverboard, and Hovershoes are shown from left to right. The complexity in terms of user control increases along the arrow head direction.

robot used for traveling on irregular terrain [2]. While these robots can offer multi-modal locomotion, the integrated leg-wheel design is neither optimized for legged nor for wheeled mobility, resulting in complex and heavy feet.

Micro-mobility platforms—such as Segways and Hoverboards—allow legged robots without wheeled feet to also have multi-modal locomotion capabilities. This is advantageous since it keeps the legs lighter for faster motions during legged movement, while also having the option of wheeled mobility by riding on these platforms when needed. Segways and Hoverboards are relatively easy to ride since the platforms for either foot are connected for increased stability and turning is accomplished through simply leaning on a bar or differentially actuating the two connected platforms. However, these micro-mobility systems have limited abilities since they have a single platform for the user’s feet, as discussed next.

Hovershoes offer more versatile movements since each of the user’s feet can move independently. However, Hovershoes increase the riding complexity as they are significantly more sensitive and have no internal controller for turning. Though the Hovershoes are the most complex to ride, they offer the greatest versatility as each foot can move independently to avoid obstacles and go on uneven terrain, which is why we selected this platform for our research. We illustrate the degree of complexity of riding these products in Fig. 2.

2) *Controls*: There has been research on developing controls for balancing on unstable environments such as seesaws and Bongo Boards with robust lateral stabilization [14]. However, these environments do not translate; but rather, they only tilt. There has been recent work on balancing on platforms that both tilts and translates, a Hoverboard, through sequential online learning control [10]. Additionally, a balancing controller for Cassie worked in riding a Segway by regulating the center-of-mass (COM) position [5]. However, these control strategies, though, are only for rigidly connected wheeled platforms and do not work on Hovershoes, which are not connected and add additional degrees-of-freedom for more versatility. Our research aims to develop a control algorithm for a bipedal robot to balance on a highly dynamic and decoupled platform under each foot.

3) *Planning*: Generic planning algorithms are primarily categorized into sampling-based and search-based methods. Sampling-based planning methods are computationally economical and offer high update rates. Nevertheless, feasible

trajectories are not always guaranteed and there is some stochasticity in the choice of the plan with the same environmental layout potentially resulting in different plans [9]. In contrast, search-based methods produce deterministic solutions, however, they may incur redundant computation due to their iterative searches [4].

4) *Vision*: An accurate odometry source is of vital importance to robot movements and velocity tracking. As a benchmark study reveals, VINS-Mono, a monocular VIO algorithm, with loop-closure detection offers decent accuracy and robustness and so does OKVIS, a stereo algorithm [3], [15]. Besides, there are also mature back-end mapping solutions to collaborate with the front-end VIO, e.g. ORB-SLAM2 which introduces a lightweight framework and RTAB-Map that uses OctoMap to achieve good memory management [13], [11].

For walking tasks on legged robots, there has been research on utilizing the contact between the feet and the static environment to optimize for odometry [6]. However, since in our case the feet are always in contact with the Hovershoes and the motion is relatively smooth, some standard robust VIO and SLAM solutions are sufficient for our scenario [8].

C. Contribution

The contributions of our work thus are:

1) *Dynamics*: We developed a dynamical model of the integrated Cassie-Hovershoes system for numerical validation.

2) *Controls*: We designed a control strategy for bipedal robots to robustly balance on independently mobile and decoupled wheeled platforms.

3) *Autonomy*: We developed a framework comprising of a vision system for odometry and obstacle detection, a path planner for online trajectory generation, and a control strategy for the Cassie bipedal robot to autonomously ride the Hovershoes and avoid obstacles.

4) *Real-world Experiments*: We demonstrated the robustness of our controller for Cassie to robustly ride on Hovershoes subject to external perturbations on real-world terrain (flat and rough ground as well as stairs), track commanded translational and rotational velocities, perform turning maneuvers, and ride the Hovershoes in a wave pattern.

5) *Simplicity*: We presented a solution that uses PD controllers designed assuming decoupled dynamics, an almost minimal sensor-suite for pose estimation and object detection, and kinematic planners to successfully solve the complex problem of making a bipedal robot autonomously ride a pair of Hovershoes.

D. Organization

The rest of the paper is organized as follows. Section II presents the dynamical model of the Cassie-Hovershoes system. Section III describes our proposed controller design. Section IV presents the simulation results. The perception and planning are described in Section V. Section VI demonstrates experimental results with discussion. Section VII discusses shortcomings of our work. Finally, Section VIII summarizes the work and provides thoughts on future work.

II. DYNAMICAL MODEL OF CASSIE ON HOVERSHOES

Having established the need for multi-modal locomotion, described the challenges, and outlined our solution, we now present a dynamical model of Cassie and Hovershoes that also captures the interaction between each other.

A. Dynamical Model of Cassie

Cassie is a highly dynamic, under-actuated bipedal robot. Cassie has twenty DOFs as listed in (1):

$$\begin{aligned} q = [q_x, q_y, q_z, q_{yaw}, q_{pitch}, q_{roll}, \\ q_{1L}, q_{2L}, q_{3L}, q_{4L}, q_{5L}, q_{6L}, q_{7L}, \\ q_{1R}, q_{2R}, q_{3R}, q_{4R}, q_{5R}, q_{6R}, q_{7R}]^T, \end{aligned} \quad (1)$$

where, (q_x, q_y, q_z) and $(q_{yaw}, q_{pitch}, q_{roll})$ are the Cartesian coordinates of the pelvis and the Euler Angles in the Z-Y-X order, and (q_{1L}, \dots, q_{7L}) , (q_{1R}, \dots, q_{7R}) are the generalized coordinates of the left and right legs, respectively. These correspond to the DOFs for each leg and are defined in (2).

$$\begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \\ q_5 \\ q_6 \\ q_7 \end{bmatrix} = \begin{bmatrix} \text{hip roll} \\ \text{hip yaw} \\ \text{hip pitch} \\ \text{knee pitch} \\ \text{shin pitch} \\ \text{tarsus pitch} \\ \text{toe pitch} \end{bmatrix}. \quad (2)$$

Fig. 3a shows the generalized coordinates of Cassie's pelvis and right leg. The generalized coordinates of Cassie's left leg are similar to the right leg states. Each of Cassie's legs has seven DOFs with five of them being actuated: q_1, q_2, q_3, q_4 , and q_7 . The corresponding motor torques are u_1, u_2, u_3, u_4 , and u_5 . The other two DOFs, q_5 and q_6 , are passive, corresponding to stiff springs.

The dynamics of Cassie can then be expressed in the following Euler-Lagrange dynamics:

$$D(q)\ddot{q} + H(q, \dot{q}) = Bu + J_s^T(q)\tau_s + J_c^T(q)F_c, \quad (3)$$

where q is the generalized coordinate vector as defined in (1), $D(q)$ is the mass matrix, $H(q, \dot{q})$ contains the centripetal, Coriolis, and gravitation terms, B is the motor torque matrix, u is the motor torque vector of dimension 10 corresponding to the actuators on the two legs, $J_s(q)$ is the Jacobian for the spring torques, τ_s is the spring torque vector, $J_c(q)$ is the Jacobian for the ground contact forces, and F_c is the ground contact force vector.

B. Dynamical Model of Hovershoes

The Hovershoes is a highly sensitive wheeled platform. It has an internal controller that regulates the pitch to zero when there is no external force. The internal parameters and states of the hovershoe are unknown. Here, we develop a closed-loop dynamical model of the Hovershoe. We consider $\theta, \psi, u_\theta, u_\psi$ to be the pitch and yaw angles and torques in the body X-Y-Z frame, and x, y, v to be the x, y position and

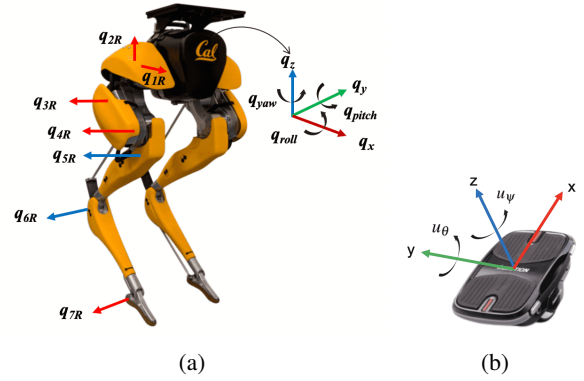


Fig. 3: (a) Kinematic model of Cassie showing the robot's generalized coordinates in the body frame. (b) Hovershoe model where u_θ and u_ψ change the pitch and yaw of the Hovershoe in the X-Y-Z body frame, respectively.

scalar speed in the global X-Y frame. Our dynamical model then is,

$$J_\theta \ddot{\theta} = -c_1 \theta - c_2 \dot{\theta} + u_\theta, \quad (4)$$

$$J_\psi \ddot{\psi} = -c_3 \psi + u_\psi, \quad (5)$$

$$\dot{x} = v \cos(\psi), \quad (6)$$

$$\dot{y} = v \sin(\psi), \quad (7)$$

$$m\dot{v} = c_4 \theta. \quad (8)$$

Here, the parameters (c_1, c_2, c_3) correspond to the Hovershoe internal controller, c_4 corresponds to the contact model between the Hovershoe and the ground. The same dynamics is used for the left and right Hovershoes.

The Hovershoe pitch dynamics is given by (4), where J_θ is the moment of inertia of the Hovershoe about the y-axis, c_1 is the coefficient for the stiffness term, c_2 is the coefficient for the damping term, and u_θ is the input torque from a rider's toe about the y-axis. The "feedback" with θ and $\dot{\theta}$ terms in this equation account for the Hovershoe's internal stabilization controller that drives θ to zero when u_θ is zero.

The yaw dynamics of the Hovershoe is described by (5), where J_ψ is the moment of inertia of the Hovershoe about the z-axis, c_3 is the coefficient for the ground contact damping term, and u_ψ is the input torque from the rider's toe about the z-axis. Here, we only have a $\ddot{\psi}$ term since there is no internal Hovershoe controller driving ψ to zero when u_ψ is zero. Figure 3b shows u_θ and u_ψ on the Hovershoe.

The global translational dynamics of the Hovershoe are captured by (6)-(7), where v is the speed and m the mass of the Hovershoe. Finally, the acceleration dynamics of the Hovershoe is captured by (8) as a function of Hovershoe pitch angle.

Note that we did not conduct system identification to determine the parameters of the model. Since computing the tilt angle and applied torque on an accelerating Hovershoe was hard, we instead focused on capturing the structure of the Hovershoe model for use in our control design.

C. Contact Model between Cassie and the Hovershoes

The Cassie robot interacts with the Hovershoes through feet contacts. The contact force is $F_c = [F_1, \dots, F_4]^T$ in (3)




Fig. 4: Autonomy framework where $\dot{\psi}_{est}$ is yaw rate of change estimated from VIO, v_{est} is the estimated velocity from VIO, $\dot{\psi}_d$ is the desired yaw rate of change given from the path planner, and v_d is the desired velocity given from the path planner. The user provides a goal location as input.

corresponding to the four contacts (front and back contact locations for either foot) with each F_i having an x, y, and z components in the local Hovershoe frame, applied on Cassie.

The two torques u_θ and u_ψ , with respect to y-axis and z-axis, respectively are computed as,

$$\begin{bmatrix} u_\theta \\ u_\psi \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \sum_{i=1}^4 r_i \times -F_i, \quad (9)$$

where, r_i is the relative position of the i^{th} contact locations with respect to each Hovershoe position, and $i = \{1, 2\}$ for left hovershoe and $i = \{3, 4\}$ for right hovershoe. The contact forces F_i are themselves determined through a compliant ground model with stick-slip friction and are then used in F_c in (3), while u_θ, u_ψ from (9) are used in (4)-(5).

III. CONTROL DESIGN

Having presented the dynamical model of the Cassie with Hovershoes system in Section II, we now proceed to discuss our control strategy that will be implemented on Cassie in order to achieve our goal of driving the Hovershoes to desired locations or at desired speeds while Cassie still balances on them. It must be noted that we can neither directly measure the Hovershoes internal states nor can we directly command input torques on the Hovershoe. Our only means of interaction between the Cassie and Hovershoe are through the contact forces. Fig. 4 illustrates our autonomy framework detailing the interactions between the vision, planner, controller, and the Cassie with Hovershoes system.

Our proposed control strategy includes a Hovershoe X-Y controller, a velocity controller, and a turning controller, all integrated with a balancing controller. The X-Y controller regulates the relative position of each Hovershoe with respect to Cassie. The velocity controller tracks a desired translational velocity and the turning controller maintains a desired rotational velocity. The balancing controller incorporates the output of the three controllers and regulates the position of Cassie's COM. Note that the gains of our proposed controller are empirically determined.

A. Hovershoe X-Y Controller

Each Hovershoe serves as a highly dynamic, moving platform for each foot of Cassie. A low-level controller that regulates the relative x-y position of each Hovershoe was designed for Cassie to keep the Hovershoes together. In the

absence of this controller, the Hovershoes drift apart resulting in Cassie falling off of the moving Hovershoes.

The Hovershoe X-Y controller is also critical since it makes the system robust to variances that occur when the Hovershoes are initialized with a non-zero relative translation or orientation, any off-center placement of Cassie's feet on the Hovershoes, and any dynamic variances between the Hovershoes.

The x-axis controller keeps the relative x-coordinate of the Hovershoes at zero in the local Cassie frame. If one of the Hovershoes is ahead of the other, then Cassie needs to bring one Hovershoe back and the other forward by pitching the Hovershoe accordingly with its toe (recall from (8) that increasing the pitch of the Hovershoe results in forward motion). This is accomplished through an inner-outer loop controller. The outer loop computes the desired difference in the pitch of either toe, see (10)-(11), while the inner loop computes the desired toe pitch torque to realize the toe pitch difference, see (12)-(13). Here $x_{foot L}$ and q_{7L} are the left foot x-position and toe pitch angles respectively. A similar notation is used for the right leg. The computed toe pitch torque gets added to the left and subtracted from the right toe torques that will be specified by the nominal balancing controller.

$$q_{\Delta toe}^{des} = -K_p^x (x_{foot L} - x_{foot R}), \quad (10)$$

$$q_{\Delta toe} = q_{7L} - q_{7R}, \quad (11)$$

$$e_{\Delta toe} = q_{\Delta toe} - q_{\Delta toe}^{des}, \quad (12)$$

$$u_5^{\Delta toe} = -K_p^{\Delta toe} e_{\Delta toe} - K_d^{\Delta toe} \dot{e}_{\Delta toe}. \quad (13)$$

The y-axis controller is designed to keep the y-coordinate of the left and right toe at y_{offset} and $-y_{offset}$, respectively. Due to the Hovershoes' nonholonomic constraint, the Hovershoe can not move sideways directly along the y-axis. The y-coordinate can only be changed when the Hovershoes yaw and move forward, so that the legs can come closer or move apart depending on the yaw angle of each Hovershoe. The y-axis controller has an inner-outer loop as well. The outer loop determines each hip's desired yaw angle through the difference between each toe's y-coordinate and desired y_{offset} , see (14)-(15). The hip yaw is used since the toe does not have an individual motor for toe yaw. In the inner loop, the desired hip yaw angle is converted into hip yaw motor torque, see (16).

$$e_y = y_{foot} \pm y_{offset}, \quad (14)$$

$$q_2^{des} = -K_p^y e_y, \quad (15)$$

$$u_2^y = -K_p^{q_2} (q_2 - q_2^{des}). \quad (16)$$

In the above equations, \pm refers to addition to one leg and subtraction for the other. To achieve a wave pattern with the Hovershoes, we can specify a time-varying $y_{offset}(t)$.

B. Velocity Controller

Once Cassie can balance on the Hovershoes and maintain a relative position between the Hovershoes while they are moving, we can then design higher-level controllers for

Cassie to perform speed regulation and turning maneuvers. These high-level controllers are required for interfacing with the path planner in the autonomous system.

The velocity controller maintains the velocity of Cassie on the Hovershoes at a specified set point given either by an operator or by the path planner. Velocity control is accomplished through regulating the x -COM position of Cassie, where the desired x -COM position is proportional to the error in velocity:

$$COM_x^{\text{des}} = -K_p^{\text{vel}} (v_{\text{est}} - v_d) \quad (17)$$

Here, v_{est} is the velocity of Cassie estimated from our vision system and v_d is the desired velocity.

C. Turning Controller

The turning controller maintains a desired yaw rate for Cassie's pelvis that is specified either by an operator or by the path planner. For the two Hovershoes to turn along a curve, the Hovershoes not only need to constantly yaw but the outer Hovershoe needs to go faster than the inner Hovershoe. Furthermore, for fast turns, Cassie will also need to bank by leaning into the turn. Leaning into a turn ensures Cassie remains balanced on the Hovershoes during fast turning maneuvers and does not tip over. This is accomplished by a three-part controller. The Hovershoes are turned by yawing Cassie's hips proportional to the yaw rate error, see (18)-(19). The outer Hovershoe is made to go faster than the inner Hovershoe by differentially actuating Cassie's toes, see (20) where the computed toe difference torque will be added to one toe and subtracted from the other in the balancing controller. Finally, Cassie is made to lean into the turn by changing the desired COM position along the y -axis when there is a large angular and/or tangential velocity, see (21)-(22).

$$e_{\dot{\psi}} = \dot{\psi}_{\text{est}} - \dot{\psi}_d, \quad (18)$$

$$u_2^{\text{turn}} = -K_p^{\text{yaw}} e_{\dot{\psi}}, \quad (19)$$

$$u_5^{\text{turn}} = -K_p^{\text{pitch}} e_{\dot{\psi}}, \quad (20)$$

$$\theta_{\text{tilt}} = \arctan \frac{\dot{\psi}_{\text{est}} v_{\text{est}}}{9.81} \quad (21)$$

$$COM_y^{\text{des}} = -K_p^{\text{shift}} (L \theta_{\text{tilt}}). \quad (22)$$

Here, y_{vision} is the actual velocity of Cassie computed from our vision system and L is the distance of Cassie's COM to the Hovershoes platform plane.

D. Integrating with a Nominal Balancing Controller

Our proposed controller is based on a nominal balancing controller for Cassie from [5]. The nominal balancing controller works by regulating the COM position error and the torso orientation. In order to move the COM along the x -axis, the controller modules the toe pitch torque. In order to move the COM along the y -axis, the controller changes the leg length difference between the left and right legs. We modify the nominal torque \tilde{u}_i from the nominal balancing controller as follows:




Fig. 5: Simulation visualization with Cassie on the Hovershoes. This is a static balancing task demonstration.

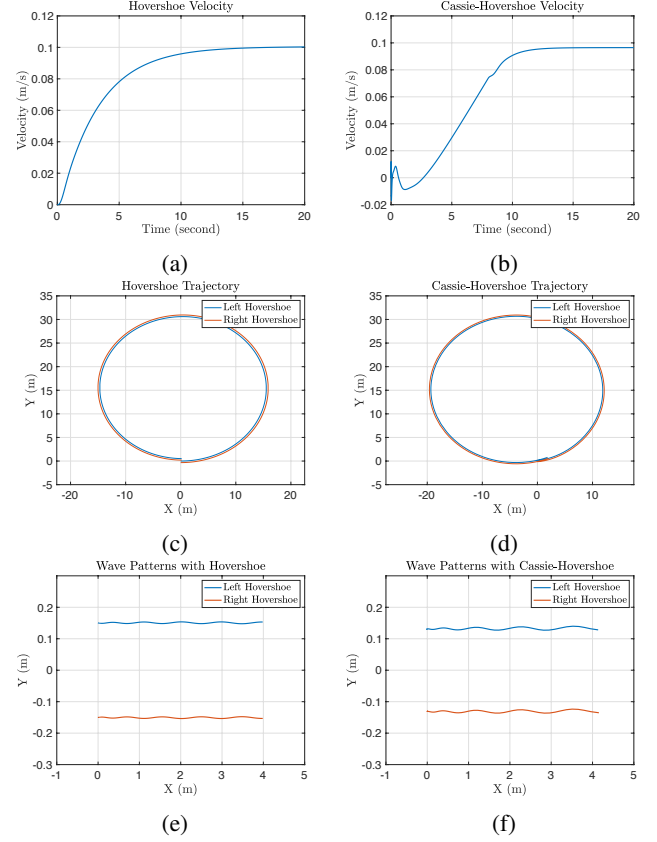


Fig. 6: Simulation results for the Hovershoes (left column) and Cassie-Hovershoe system (right column) for (a,b) forward velocity tracking, (c,d) going around in a circle, and (e,f) following a wave pattern.

$$u_1 = \tilde{u}_1, \quad (23)$$

$$u_2 = \tilde{u}_2 + u_2^y + u_2^{\text{turn}}, \quad (24)$$

$$u_3 = \tilde{u}_3 (COM_y^{\text{des}}), \quad (25)$$

$$u_4 = \tilde{u}_4 (COM_y^{\text{des}}), \quad (26)$$

$$u_5 = \tilde{u}_5 (COM_x^{\text{des}}) \pm u_5^{\Delta \text{toe}} \pm u_5^{\text{turn}} + u_5^{\text{damping}}. \quad (27)$$

In particular, the hip yaw torque u_2 in (24) comprises of the yaw torque from the nominal balancing controller \tilde{u}_2 , as well as torques from the y -axis controller u_2^y , and turning controller u_2^{turn} . The toe torque u_5 in (27) comprises of the




Fig. 7: The left picture shows the real scene and the right demonstrates the segmented ground and obstacles.

nominal balancing controller toe torque \tilde{u}_5 as well as the toe pitch difference torque $u_5^{\Delta\text{toe}}$ from the x-axis controller, the turning controller torque u_5^{turn} , and a damping controller u_5^{damping} that suppresses the oscillation of the toe pitch joint. As mentioned earlier, \pm refers to a torque that is added to one leg and subtracted from the other.

IV. SIMULATION RESULTS

In order to verify the proposed controller design presented in Section III on the dynamical model from Section II, we ran two numerical simulations. First, we considered the Hovershoes (without Cassie) and provided the controllers full state knowledge of the Hovershoes and allowed the controllers to directly actuate the Hovershoes. In the next simulation, we considered the Cassie-Hovershoes integrated system where the controller did not have access to the Hovershoe state and could only indirectly actuate the Hovershoes through Cassie. We discuss these simulation results in the following subsections. The simulations ran in MATLAB 2018a installed on a laptop (Intel Core i5, RAM 8GB).

A. Hovershoes Simulation

We modeled each Hovershoe in Matlab with the dynamics specified in (4)-(8). In the simulation, we assumed perfect state knowledge of the Hovershoes, and the controllers would provide the input torques: u_θ and u_ψ .

We verified the velocity controller by setting a desired velocity of 0.1 m/s and verifying the tracking as shown in Fig. 6a. We verified the turning controller by setting a desired rotational velocity and ensuring the Hovershoes turned in a circle as shown in Fig. 6c. Finally, we verified the wave pattern controller by setting the y_{offset} to a sinusoidal wave, resulting in Fig. 6e.

B. Cassie-Hovershoes System Simulation

We next simulated the Cassie-Hovershoes integrated system dynamics (3)-(8), with our designed controller. Fig. 5 shows the visualization of Cassie and the Hovershoes balancing in the simulation. We used the same set points as the previous Hovershoes simulation to test our velocity, turning, and wave pattern controller and obtained a similar performance, see Figs. 6b, 6d, and 6f.

V. PERCEPTION AND PLANNING

Having numerically validated our control strategy, we will next pave the path for experiments by presenting our perception and path planning strategies and also detailing the communication between perception, planning, and control.




Fig. 8: Hardware infrastructure diagram for the autonomous Cassie system. The vision and path planning computer and the controller computer are onboard Cassie, while the monitoring computer is connected wirelessly and used to see signals and tune parameters.

A. Perception

In order to modularize our set-up to be simple and easy to manage, and at the same time decrease the computational load, we opted for an integrated VIO sensor, the Intel RealSense T265, which is comprised of a pair of fish-eye lenses is combined with an on-board implementation and extension of an algorithm utilizing stereo VIO with loop-closure detection to produce reliable odometry [8]. Additionally, a rigidly attached depth camera, Intel RealSense D435i, was used for the obstacle mapping. The D435i camera processed and matched the stereo images using dedicated on-board hardware. By designing the vision set-up as above, we are able to incorporate the entire perception, planning, and control pipeline using the CPUs on Cassie, and avoid the need of power-hungry GPUs.

The transformation between the two cameras was known and published internally via ROS. The depth camera was slightly tilted towards the ground to cover more ground space near Cassie. The depth maps were converted to point clouds, which were then filtered and voxelized to a horizontal resolution of 5cm and vertical resolution of 10 cm to reduce the memory usage and accelerate the computation. After transforming the point clouds to the base frame of the Hovershoes, a ground segmentation with a tuned tolerance to noise was applied to remove the ground plane and register the obstacles onto an global occupancy map. Fig. 7 shows an example of our segmentation. The map registration was achieved by using the RTAB-Map ROS package, but the built-in appearance-based loop closure detector was turned off. In order to increase the refresh rate for the local costmap, the segmented 3-D obstacles were further converted and merged to a pseudo-laser scan to decrease the data rate. Based on the above perception infrastructure, we next detail how we can achieve real-time planning in order to circumvent obstacles.

B. Motion Planning

To create real-time motion plans for the Cassie-Hovershoes system, we require the planner to be reliable and kinematically feasible. For the global planner, we employed



Fig. 9: Cassie riding the Hovershoes on (a) different heights, (b) rough terrain, (c) up a 7° incline, and (d) down a 13° incline.




Fig. 10: Cassie robustly descending three stairs with heights 79%, 93%, 93% of the wheel radius respectively.

a search-based method, Dijkstra’s algorithm, for its guarantee of a solution. The costmap decay rate was tuned to confine the Dijkstra’s behavior to follow a path that stays in the middle of two potentially incoming obstacles.

For the local planner, the optimization-based Timed Elastic Band (TEB) planner was applied on top of the Dijkstra’s algorithm as the TEB planner looks 1.75 m ahead of the current position and generates a short-term, kinematically feasible path for Cassie. In order to emulate the forward non-holonomic property of the Hovershoes, customized car-like optimization weights were used to encourage forward motion. The optimization process was adjusted to adapt the computing power of Cassie and the resulting planner frequency could reach up to slightly less than 50 Hz. However, 10 Hz proved to be sufficient for Cassie to autonomously navigate through a narrow corridor with obstacles.

C. Communication between Perception, Planning & Control

Fig. 8 shows a diagram of the communication setup and hardware infrastructure. Please note that the router can be replaced with a Wi-Fi hotspot. In addition, between the real-time controller and vision system, we use the UDP to communicate to each other. Finally, the vision and path planning have a separate computer other than the controller because only the controller requires a real-time computer. We use a gantry during experiments for safety purposes only.

VI. EXPERIMENTAL RESULTS

After developing our control strategy, path planning algorithm, and perception system, we now present the results of our autonomous framework. First, we demonstrate the robustness of our controller. Then, we show the performance of our velocity and turning maneuver. Last, we demonstrate that Cassie successfully navigating an obstacle course autonomously while riding the Hovershoes.

A. Robustness

Once Cassie could balance and perform the tasks in simulation, we transferred the controller to Cassie and checked the robustness of the controller subject to disturbances that could




Fig. 11: Top: velocity response to three successive step inputs; Bottom: yaw rate of change response to two successive step inputs.

cause Cassie to fall. Robustness was checked by kicking the Hovershoes and pushing Cassie’s pelvis and tarsus while it was riding the Hovershoes. In addition, Cassie successfully traversed over rough terrain with cracks and bumps, as well as riding on up and down slopes with 7° and 13° inclines respectively, and riding with Hovershoes at different heights (see Fig. 9). A challenge task was to go down a variety of steps. Cassie successfully went down single stairs that were 79% and 171% of the Hovershoe wheel radius and a sequence of three stairs that were 79%, 93%, 93% of the Hovershoe wheel radius, respectively. As can be seen in Fig. 10, Cassie remains balanced and the Hovershoes remain level after descending each step.

B. Vision-based Velocity Tracking

We tested the performance of our vision-based velocity controller by providing a step input change in desired velocity. We found that it can track a desired velocity with a small steady-state error. As can be seen in Fig. 11, the desired velocity step inputs of 0.5 m/s, 1.0 m/s, and 1.5 m/s, were tracked with small steady-state errors with the root mean square error (RMSE) being 0.1467 m/s.

C. Fast Turning

A similar step input for turning yaw rotational velocity of 0.5 rad/s and 0.0 rad/s produced the response shown in Fig. 11. Clearly, the tracking is not perfect and has a steady-state tracking error with RMSE is about 0.1446 rad/s. This can be improved with better choice of gains or incorporating an integral feedback. Furthermore, Fig. 12 shows Cassie performing a high-speed turn. It is clear from the center graphic of Fig. 12 that Cassie successfully leans into the turn by shifting its COM position along the y-axis in order to prevent from tipping over during the turn.

D. Wave Pattern with Feet for Obstacle Avoidance

Certain obstacles are short enough where Cassie does not need to circumnavigate them in order to avoid collision. With the Hovershoes, Cassie can split around the obstacle




Fig. 12: Cassie executing a fast turn by leaning, with the lean angle directly dependent on the forward and rotational velocities.



Fig. 13: Cassie showing wave pattern with feet task by avoiding an obstacle without going around it. For certain obstacles, Cassie can stay on a straight trajectory while still missing the obstacle.

using the wave pattern controller. Fig. 13 shows that Cassie successfully avoids an obstacle by doing a wave pattern maneuver with its feet. The decision making of splitting its feet was made by a human operator.

E. Obstacle Avoidance

Finally, combining our controllers with the computer vision obstacle detection system and path planner, we achieved autonomous locomotion through an obstacle course. The leftmost graphic of Fig. 14 shows the obstacle course and the rest show Cassie autonomously navigating through the obstacle course.

VII. SHORTCOMINGS OF OUR PROPOSED WORK

Our proposed method is the simplest control and autonomy package for achieving the task of riding Hovershoes. In particular, our controller decouples the system dynamics into several simple sub-components, our planner only considers a kinematical model of the system, and our vision system only uses the stereo camera and IMU without using the additional sensors on the robot. By considering a fully coupled dynamical system and incorporating the dynamics into the planner and taking all available sensing into account, we can further improve our state estimation, controller and planner performance significantly. Finally, our proposed solution assumes Cassie is somehow initialized on the Hovershoes, which is currently done manually and does not truly offer multi-modal locomotion capabilities yet.

VIII. CONCLUSIONS

In this paper, we have presented a framework for autonomous locomotion of a Cassie bipedal robot over Hovershoes. Our developed framework enables the Cassie bipedal robot to interact with the Hovershoes to balance, regulate forward and rotational velocity, achieve fast turns, and move over flat terrain, slopes, stairs, and rough outdoor terrain, as well as autonomously navigate an obstacle course.




Fig. 14: Cassie performing obstacle avoidance. The leftmost figure is a picture of the obstacle course. The rest of the pictures are sequential snapshots of the path planner and 3-D map while the robot is autonomously navigating through the course.

The future work will focus on combining Cassie walking to achieve multi-modal locomotion. Currently human operators are needed to initialize Cassie on the Hovershoes. We would like to develop a method for Cassie to step up on the Hovershoes without the need of operators.

REFERENCES

- [1] "Handle: Mobile box handling robots for logistics," <https://www.bostondynamics.com/handle>, accessed: 2019-06-06.
- [2] M. Bjelonic, C. D. Bellicoso, *et al.*, "Keep rollin' - whole-body motion control and planning for wheeled quadrupedal robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2116–2123, 2019.
- [3] J. Delmerico and D. Scaramuzza, "A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots," in *IEEE Int. Conf. on Robotics and Automation*, 2018.
- [4] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [5] Y. Gong, R. Hartley, *et al.*, "Feedback control of a cassie bipedal robot: Walking, standing, and riding a segway," in *American Control Conference*, 2019.
- [6] R. Hartley, J. Mangelson, *et al.*, "Legged robot state-estimation through combined forward kinematic and preintegrated contact factors," in *IEEE Int. Conf. on Robotics and Automation*, 2018.
- [7] K. Hashimoto, T. Hosobata, *et al.*, "Realization by biped leg-wheeled robot of biped walking and wheel-driven locomotion," in *IEEE Int. Conf. on Robotics and Automation*, 2005.
- [8] E. S. Jones and S. Soatto, "Visual-inertial navigation, mapping and localization: A scalable real-time causal approach," *The Int. J. of Robotics Research*, vol. 30, no. 4, pp. 407–430, 2011.
- [9] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The Int. J. of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [10] K. Kimura, S. Nozawa, *et al.*, "Riding and speed governing for parallel two-wheeled scooter based on sequential online learning control by humanoid robot," in *IEEE Int. Conf. on Intelligent Robots and Systems*, 2018.
- [11] M. Lobb and F. Michaud, "RTABmap as an opensource lidar and visual simultaneous localization and mapping library for largescale and longterm online operation," *Journal of Field Robotics*, vol. 36, no. 2, pp. 416–446, 2019.
- [12] J. Lim, I. Lee, *et al.*, "Robot system of DRC-HUBO+ and control strategy of team KAIST in DARPA robotics challenge finals," *Journal of Field Robotics*, vol. 34, no. 4, pp. 802–829, 2017.
- [13] R. Mur-Artal and J. D. Tardos, "ORB-SLAM2: an open-source slam system for monocular, stereo and rgb-d cameras," *IEEE Trans. on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [14] U. Nagarajan and K. Yamane, "Balancing in dynamic, unstable environments without direct feedback of environment information," *IEEE Transactions on Robotics*, vol. 30, no. 5, pp. 1234–1241, 2014.
- [15] P. L. Tong Qin and S. Shen, "VINS-Mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Trans. on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [16] S. Waharte and N. Trigoni, "Supporting search and rescue operations with uavs," in *IEEE Int. Conf. on Emerging Security Technologies*, 2010.