Predictive Planning Based on Reactive Control

Patrick J. Clary* and Jonathan W. Hurst*
*Oregon State University, Corvallis, USA
claryp@oregonstate.edu

Summary

Cassie, the Dynamic Robotics Lab's newest bipedal robot, currently walks using controllers that react to disturbances instead of preemptively planning around obstacles. This gives it unprecedented speed and robustness when guided by a human operator, but leaves it unable to operate autonomously when planning is needed, such as on staircases and in crowds of people. Experiments in simulation show that these reactive controllers can result in sufficiently predictable behavior to use them as a basis for model-predictive planning for obstacles that cannot be handled reactively. This movement control approach will be further developed to allow robots like Cassie to quickly traverse complex environments while retaining the robust disturbance recovery characteristics of the underlying reactive controller.

1 Introduction

The bipedal robot Cassie (Figure 1) is currently able to stand, walk, and run using controllers that do not attempt to plan movement in advance, but combine feedforward and feedback elements to remain upright and regulate velocity. The reactive controllers used on Cassie drive the robot's stepping with a clock and use the robot's center-of-mass velocity to control foot placement. With different controllers or parameter settings, Cassie can be made to stand, walk, run, and jump in many different ways. These controllers give the robot inherent robustness to disturbances like pushes, slips, and drop steps, and they can be designed to use only proprioceptive sensor data from joint encoders and an inertial measurement unit.

Some situations cannot be handled by reacting to disturbances as they occur. Large enough obstacles will destabilize the robot regardless of what the controller does to react. Stairs significantly constrain foot placement and require a structured pattern of movement to negotiate. A robot operating in a crowd of humans must be careful to avoid crashing into people, even when a collision would not harm the robot itself.

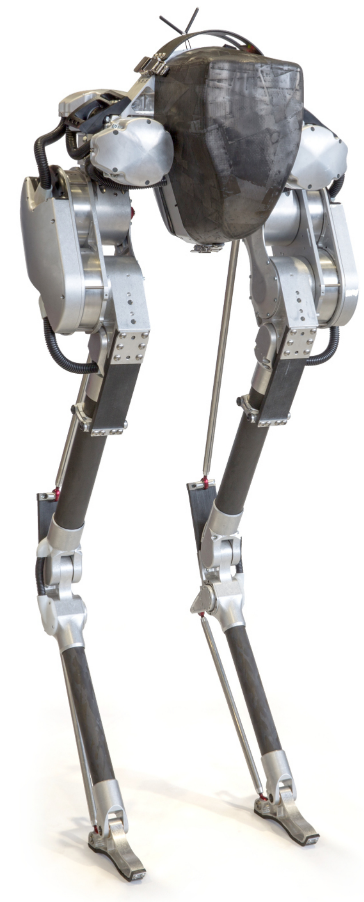

Figure 1: Cassie

Bipeds like Cassie will need to use some kind of preemptive planning to operate in these environments. The usual approach, finding motion plans in a kinematic space and executing these plans with a full-body trajectory follower, is much less robust to typical locomotion disturbances than reactive control, so we are now investigating other ways of implementing planning on Cassie.

2 Reactive Controllers as Actions

We are casting the planning problem as choosing a sequence of actions with a high probability of satisfying movement objectives. The characteristics of the action space are critical to making such an approach feasible for real-time control. Rather than instantaneous joint torques, we propose using the execution of a reactive controller with fixed parameters for a short amount of time (0.1–1 second) as an action. For example, one action could be taking a step forward using a basic walking controller, while another action could cause the robot to jump or take a high step that clears an obstacle. This action space is compact, implicitly excluding "white noise"-type torque profiles that are not useful for locomotion, and expressive, allowing meaningful behavior to be described with relatively few parameters.

With careful controller design, these actions can be chained together to produce complex motion plans. Using model-predictive methods, the robot can choose a plan that is likely to accomplish its objectives while negotiating obstacles, but that falls back on robust reactive behavior when the prediction is incorrect due to imperfect sensing or unpredictable external disturbances.

3 Experiments

A proof of concept has been implemented on a planar walking simulation with physical parameters resembling Cassie. Using a short-horizon Monte-Carlo planner to choose actions, the robot model is capable of negotiating a variety of terrain types (Figure 2). The simulator includes some stochasticity and the planner uses state abstraction techniques to ensure that its plans are robust.

The simulator used here is simpler than the multibody simulation used for controller development on Cassie, but captures much more of the dynamics than a SLIP model. The robot model includes damping, body inertia, toe mass, motor dynamics, and communication delays. As a result, the controller and planner should not be able to exploit patently nonphysical effects and this type of control should be feasible on the physical robot.

This proof of concept implementation is capable of safely and reliably guiding the robot over terrain filled with gaps, drop steps, and staircases. When the robot approaches particularly difficult obstacles, it stops short of them and waits until it finds a suitable plan, if such a plan exists. When subjected to large center-of-mass velocity disturbances at random intervals, the robot responds immediately due to the reactive controller and is usually able to stabilize itself and continue.

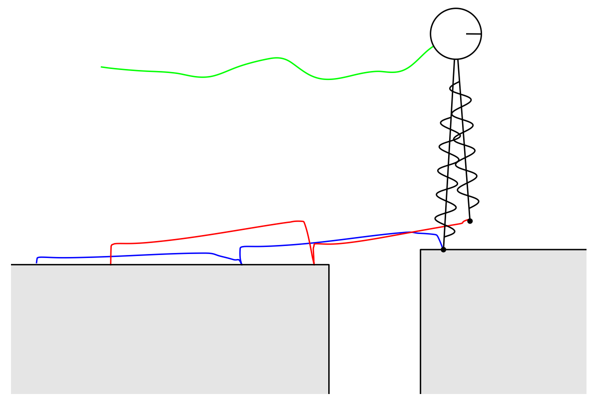

Figure 2: Simulated robot after executing a planned maneuver. Traces show center-of-mass and foot movement.

4 Future Work

Future work will be focused towards enabling a demonstration of Cassie negotiating large and complex obstacles in real-time. This will require additional controller development on Cassie, additional modeling and verification of the robot's dynamics, and improvements to the planner that allow it to make good choices faster.

To overcome processing speed limitations, we are investigating using learned heuristics to suggest better candidate actions. A human directing the search process would be able to find a viable plan with far fewer action-outcome predictions than the existing random sampler. Deep networks that take the robot's state and a convolutional view of the terrain may be able to map situations to candidate actions well enough to make real-time operation feasible.