Fast Linearized Model-Predictive Control for Legged Robots

Simon Le Cleac'h, Taylor Howell, and Mac Schwager
Stanford University, California, USA
simonlc, thowell, schwager@stanford.edu

Zachary Manchester
Carnegie Mellon University, Pennsylvania, USA
zacm@cmu.edu

I. Introduction

We present a general approach for controlling robotic systems that make and break contact with their environments. Our linearized contact-implicit model-predictive control framework is a natural extension of LQR or linear model-predictive control to the contact setting. A differentiable contact simulator, formulated as an interior-point solver that reliably converges to hard-contact physics, is employed as a lower-level implicit constraint in a bi-level model-predictive control formulation. This policy leverages precomputed linearizations about a reference trajectory to simplify the problem and speedup online computation. This approach is efficient and reliable while retaining the binary contact switches encoded by complementarity constraints, enabling policies to be robust to contact timing. In simulation, we demonstrate the robustness of the linearized model-predictive control policy to varying sampling rates, disturbances, and unmodeled environments for a collection of robotic systems, including: a hopper, planar quadruped, and planar biped.

II. Approach

A. Differentiable Physics Simulator

A configuration-based time-stepping scheme is used to simulate rigid-body systems that experience contact. At each time step, a feasibility problem, formulated as a Nash Equilibrium between impact and friction optimization problems [2] is solved. An interior point method successively reduces the complementary slackness which directly corresponds to converging from "soft" to "hard" contact. We present an approach for differentiating through this equilibrium point [1] in order to compute the gradients of the current configuration with respect to previous configurations and system parameters.

B. Linearized Model-Predictive Control

For online control, we efficiently solve a linearized tracking problem. The differentiable simulator is utilized as a lower-level constraint which encodes the dynamics in an upper-level motion-planning problem. To reduce computational complexity, nonlinear terms are linearized about a reference trajectory, with the exception of complementarity constraints. In this way, we simplify the contact dynamics while retaining their crucial component: contact switches.

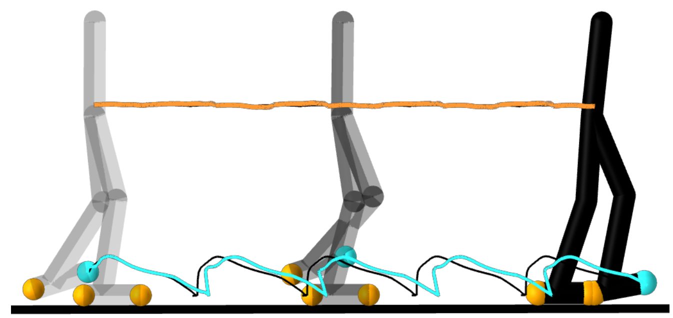

Fig. 1. Simulation results for linearized model-predictive control applied to a planar biped. The system tracks the center of mass (orange) and foot position (blue) of the reference gait.

III. Results

The linearized model-predictive control policy is demonstrated by tracking a reference trajectory in the presence of varying sample rates, disturbances, and uneven terrain for three systems in simulation. The policy is re-optimized faster than real time at each time step for a hopper, planar quadruped, and planar biped. We compare our policy against system-specific baselines: the Raibert heuristic [4] for the hopper and a spring flamingo policy [3] to compare against our biped.

References

[1] Brandon Amos and J. Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pages 136–145. PMLR, 2017.

[2] Zachary Manchester and Scott Kuindersma. Variational contact-implicit trajectory optimization. In *Robotics Research*, pages 985–1000. Springer, 2020.

[3] Jerry E. Pratt. *Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots*. PhD thesis, Massachusetts Institute of Technology, May 2000.

[4] Marc H. Raibert, H. Benjamin Brown Jr., Michael Chepponis, Jeff Koechling, Jessica K. Hodgins, Diane Dustman, W. Kevin Brennan, David S. Barrett, Clay M. Thompson, John Daniell Hebert, Woojin Lee, and Borvansky Lance. Dynamically stable legged locomotion. Technical report, Massachusetts Institute of Technology Cambridge Artificial Intelligence Lab, 1989.