Fast Contact-Implicit Model-Predictive Control

Simon Le Cleac'h[1*], Taylor A. Howell[1*], Mac Schwager[2], and Zachary Manchester[3]

Abstract—**We present a general approach for controlling robotic systems that make and break contact with their environments. Contact-implicit model-predictive control (CI-MPC) generalizes linear MPC to contact-rich settings by relying on linear complementarity problems (LCP) computed using strategic Taylor approximations about a reference trajectory and retaining non-smooth impact and friction dynamics, allowing the policy to not only reason about contact forces and timing, but also generate entirely new contact mode sequences online. To achieve reliable and fast numerical convergence, we devise a structure-exploiting, interior-point solver for the LCP contact dynamics and a custom trajectory optimizer for trajectory-tracking MPC problems. We demonstrate CI-MPC at real-time rates in simulation, and show that it is robust to model mismatch and can respond to disturbances by discovering and exploiting new contact modes across a variety of robotic systems, including a pushbot, hopper, and planar quadruped and biped.**

Index Terms—**Model-Predictive Control, Legged Robots, Contact Modeling, Optimization and Optimal Control.**

I. Introduction

Controlling systems that make and break contact with their environments is one of the grand challenges in robotics. Numerous approaches have been employed for controlling such systems, ranging from hybrid-zero dynamics [13, 1, 10], to complementarity controllers [2], to neural-network policies [4, 5], and model-predictive control (MPC) [14, 11]. There have also been numerous successes deploying such approaches on complex systems in recent years: direct trajectory optimization and LQR on Atlas [8], smooth-contact models and differential dynamic programming on HRP-2 [12, 7], zero-moment point and feedback linearization on ASIMO [6], and MPC with simplified dynamics models on Cheetah [3] and ANYmal [9]. However, reliable *general-purpose* control techniques that can reason about contact events and can be applied across a wide range of robotic systems without requiring application-specific model simplifications, gait-generation heuristics, or extensive parameter tuning remain elusive.

In this work, we focus on the problem of local tracking control for systems that experience contact interactions with their environments. Our approach combines a differentiable "hard-contact" rigid-body dynamics formulation with strategic linearizations, exploitation of the trajectory-optimization problem structure, and specialized numerical optimization techniques. The result is a model-predictive-control algorithm that can effectively reason about contact changes in the presence of large disturbances while remaining fast enough for real-time execution on modest computing hardware.

We formulate dynamics with contact as a complementarity problem that simultaneously satisfies impact and friction

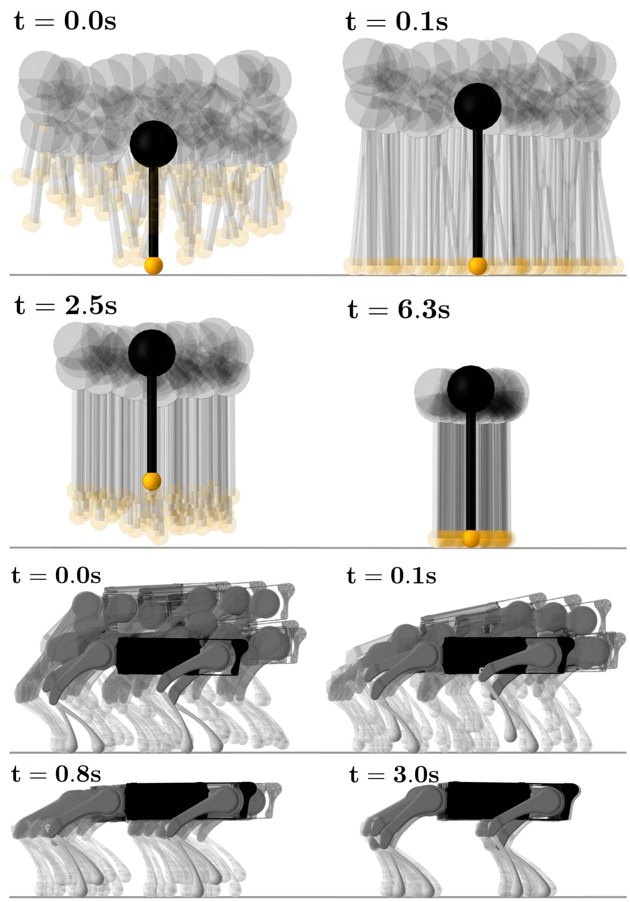

Fig. 1: Monte Carlo simulations of initial conditions for systems tracking a reference trajectory. 100 initial configurations are randomly sampled for a hopper (top) and quadruped (bottom). Perturbations from the reference initial configuration include large translations, tilts, and joint-angle offsets. The policy successfully converges to the reference gait for all initial conditions on both systems.

[1] Simon Le Cleac'h and Taylor A. Howell are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA. {simonlc, thowell}@stanford.edu

[2] Mac Schwager is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, USA. schwager@stanford.edu

[3] Zachary Manchester is with The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. zacm@cmu.edu

(Corresponding author: Taylor A. Howell)

* These authors contributed equally to this work.

constraints. By employing a interior-point method to optimize this problem, we naturally and reliably converge from "soft" to "hard" contact as the central-path parameter is decreased to zero. At a solution point, the implicit-function theorem is then utilized to efficiently compute dynamics derivatives for use in the policy. To enable real-time performance for model-predictive control, we pre-compute linearizations of the system's dynamics, signed-distance functions, and friction cones about a reference trajectory, while explicitly retaining complementarity constraints that encode contact switching behavior, resulting in a sequence of lower-level time-varying linear-complementarity problems (LCP). An upper-level trajectory-optimization problem is then optimized using an efficient structure-exploiting solver. We refer to this algorithm as *contact-implicit model-predictive control* (CI-MPC).

Finally, we demonstrate that our CI-MPC policy can generate new contact sequences online and reliably track reference trajectories while subject to significant model mismatch and large disturbances for a number of qualitatively different robotic systems, including a pushbot, hopper, and planar quadruped and biped.

Our specific contributions are:

- A contact-dynamics formulation that can be reliably evaluated and efficiently differentiated with a custom interior-point solver
- Fast, structure-exploiting solvers for the contact-dynamics and trajectory-optimization problems
- A model-predictive-control framework for robotic systems with contact dynamics
- A collection of simulation examples demonstrating the performance of the CI-MPC algorithm on a variety of robotic systems across a range of highly dynamic balancing and locomotion tasks

A pre-print of our paper is available on arXiv: `arxiv.org/abs/2107.05616`. An open-source implementation of the algorithm, `ContactImplicitMPC.jl`, along with a set of experiments on various dynamical systems are available at: `github.com/dojo-sim/ContactImplicitMPC.jl`.

Acknowledgments

This work was supported in part by Frontier Robotics, Innovative Research Excellence, Honda R&D Co., Ltd, ONR award N00014-18-1-2830, NSF NRI award 1830402, and DARPA YFA award D18AP00064. Toyota Research Institute ("TRI") provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

References

[1] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle. Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics. *IEEE Transactions on Automatic Control*, 59(4):876–891, 2014.

[2] A. Aydinoglu, V. M. Preciado, and M. Posa. Contact-aware controller design for complementarity systems. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1525–1531. IEEE, 2020.

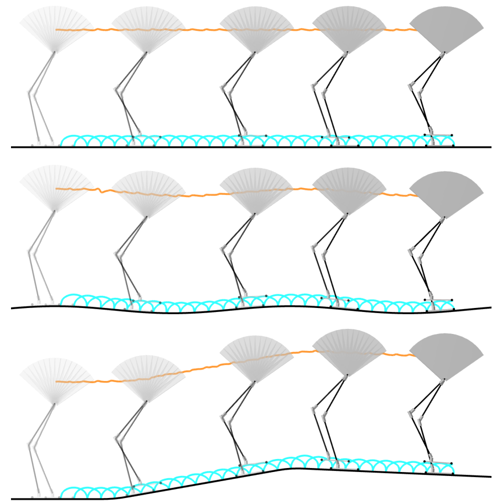

Fig. 2: Biped walking from left to right across flat (top), sinusoidal (middle), and piecewise linear (bottom) terrain using the same policy.

[3] G. Bledt. *Regularized predictive control framework for robust dynamic legged locomotion*. PhD thesis, Massachusetts Institute of Technology, 2020.

[4] N. Heess, D. TB, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. M. Eslami, et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.

[5] E. Heiden, D. Millard, E. Coumans, Y. Sheng, and G. S. Sukhatme. Neuralsim: Augmenting differentiable simulators with neural networks. *arXiv preprint arXiv:2011.04217*, 2020.

[6] K. Hirai, M. Hirose, Y. Haikawa, and T. Takenaka. The development of Honda humanoid robot. In *Proceedings. 1998 IEEE International Conference on Robotics and Automation*, volume 2, pages 1321–1326. IEEE, 1998.

[7] J. Koenemann, A. Del Prete, Y. Tassa, E. Todorov, O. Stasse, M. Bennewitz, and N. Mansard. Whole-body model-predictive control applied to the HRP-2 humanoid. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3346–3351. IEEE, 2015.

[8] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake. Optimization-based locomotion planning, estimation, and control design for the Atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, 2016.

[9] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter. Learning quadrupedal locomotion over challenging terrain. *Science Robotics*, 5 (47), 2020.

[10] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. *arXiv preprint arXiv:2103.14295*, 2021.

[11] J.-P. Sleiman, F. Farshidian, M. V. Minniti, and M. Hutter. A unified mpc framework for whole-body dynamic locomotion and manipulation. *IEEE Robotics and Automation Letters*, 6(3):4688–4695, 2021.

[12] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

[13] E. R. Westervelt, J. W. Grizzle, and D. E. Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1):42–56, 2003.

[14] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli. Gait and Trajectory Optimization for Legged Systems Through Phase-Based End-Effector Parameterization. *IEEE Robotics and Automation Letters*, 3(3): 1560–1567, July 2018.