Full Dynamics LQR Control With Multi Contact Phases For Bipedal Walking

Sean Mason[1], Stefan Schaal[1,2] and Ludovic Righetti[2]
[1] University of Southern California, USA
sean.mason@usc.edu
[2] Max Planck Institute, Tuebingen, Germany
ludovic.righetti@tuebingen.mpg.de , *sschaal@usc.edu*

I. Introduction

Humanoid robots are expected to both locomote and interact with objects within unstructured environments. A key component to realizing this goal is the ability for our algorithms to handle multiple contact switching scenarios within a whole-body control framework. To complement the many advances in trajectory optimization we must recognize which tools work best to robustly track these trajectories while providing a well-defined disturbance rejection behavior. Over recent years, optimal control strategies have showed very promising results in simulation and on real systems for torque controlled humanoids that use operational-space techniques to achieve whole-body manipulation and locomotion. Previous work, [5], [12], [1], [11], [2], has utilized Quadratic Programs (QPs) that optimize over a variety of constraints (e.g. dynamic consistency, joint tracking, friction cones, etc.) in order to compute joint torques. Trajectories are often planned using operational-space techniques and then converted to joint torques using QPs, achieving whole-body manipulation. QPs can further be organized into hierarchies to solve whole-body optimal control problems such that there is a set priority in goals that the robot should achieve and tasks of higher priorities will always be achieved first [4]. Unfortunately, coupled with the growing flexibility of these methods there is also the added computational overhead, complexity in tuning, and a lack of theoretical disturbance rejection metrics, such as the gain and phase margin seen in classical control, that prevent these algorithms from being compared with one another.

Currently, it remains unclear what level of complexity in controllers is needed on real systems. Real systems have problems such as model error, sensor noise, actuator saturations, backlash, stiction, and imperfect state estimation. Because of theses issues, the advances seen by algorithms in simulation do not always transfer to real systems. To address this unknown, we propose testing simpler optimal control strategies that still offer whole-body control and can handle multiple contact constraints. We are interested in studying to what extent contact-consistent linearized dynamic models can be used for the control of whole-body behaviors on real robots. In particular, we are interested in understanding how the performance compares with more advanced operational-space control approaches on real robots. Previously in [7],

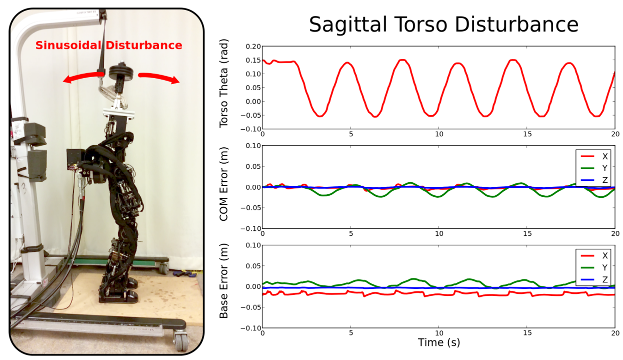

Fig. 1. Balancing experiment conducted on the hydraulic torque controlled Sarcos humanoid. The weighted torso tracked a sine trajectory while the lower body balanced using a LQR controller with multiple foot constraints.

we derived a linear quadratic regulator (LQR) making use of the linearized full robot dynamics which are consistent with the contact constraints. It is well-known that LQR controllers provide very good gain and phase margins and by extension a certain tolerance to non-linearities. With more complex optimization-based algorithms, it becomes difficult to provide similar stability analyses. The advantage of the constrained LQR controller is that it explicitly takes into account the coupling between the different joints to create optimal feedback controllers. To address the complexity of tuning this high-dimensional controller in joint space, we have shown in [7] how this framework also allows costs on operational-space quantities (i.e. end effector tracking, center of mass (CoM) tracking, and momentum) to be embedded into the total cost function. We do this through a change of variable in states and rewrite the LQR cost function to be in terms of these new variables. Experiments on a torque-controlled humanoid robot (Fig. 1) have demonstrated that our computationally light weight control policy had push-recovery and tracking performances competitive with more sophisticated balance controllers [11], [2], [8], [6] rejecting impulses up to 11.7 Ns with peak forces of 650 N. We have extended the original approach [7] to include contact switching and re-linearization of the dynamics for different contact situations. Our recent results show that our approach can be extended to more complex scenarios such as walking as detailed in the following section.

II. Infinite Horizion LQR using a single linearization

Previously in [7], we conducted balancing experiments by using a single linearization and static gains provided by an infinite horizon control problem. We derived a contact-consistent LQR formulation using the linearized model of the full dynamics, similar to [13]. The robot was pushed with a rigid rod instrumented with an ATI force/torque sensor, from which the maximum perturbation impulse was computed. We have since extended this approach to a decoupled task situation. It is often the case that the upper body of the robot is interested in a task that can be considered decoupled from the lower body (i.e. working at a table) and the lower body should stabilize the full system. In one task, the CoM is stabilized to keep balance and the other task consists in tracking a fast torso motion. These two operational-space tasks are embedded in the cost used to derive the LQR controller. Fig. 1 shows the lower body balancing while the torso (equipped with weights to simulate the upper body mass) tracked sinusoidal trajectories. The resulting upper body motion simulated forces which may be experienced during a whole-body manipulation task. With just the simple LQR feedback the robot was able to balance while the weighted torso moved at frequencies between 0-1.5 Hz. Note that just performing the same task using PD control for each joint would lead to a fall.

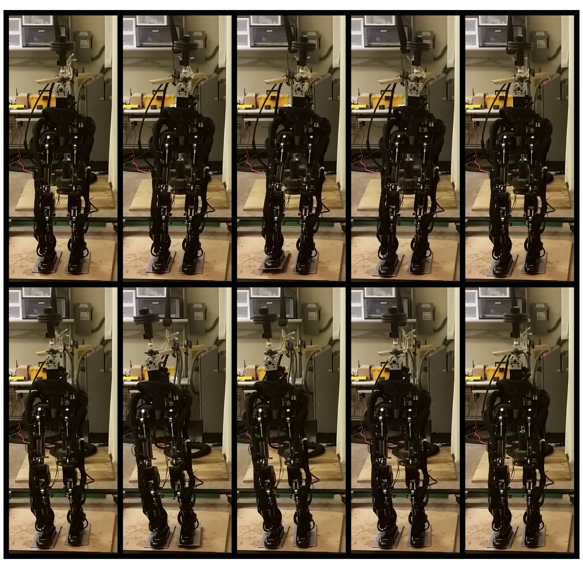

Fig. 2. Snap shots of the robot walking using 5 key linearization poses and changing contact constraints.

III. Time Varying Dynamics and Constraints

We have extended our approach to more complex scenarios such as walking. Rather than using a single linearization as in the balancing experiments, we linearize the full robot dynamics around key contact poses such as those shown in Fig. 5. One of our assumptions is that only a small set of linear models is necessary to control many interesting tasks. By using a small set of linearized models we are able to compute efficient controllers using LQR design

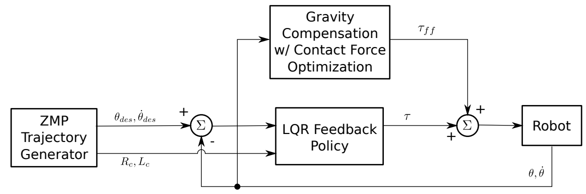

Fig. 3. Block diagram depicting the control architecture used in the walking experiments. θ is the joint and base state, τ_{ff} is the feed-forward torque, and R_c , L_c denote the time varying contact trajectories for the right and left foot.

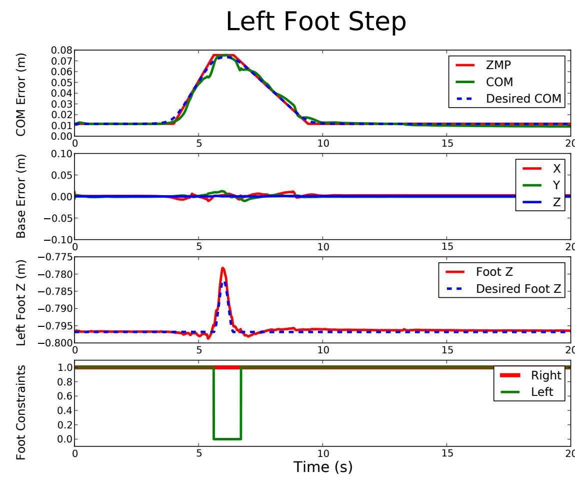

Fig. 4. Plot of the robot taking a single step. In this plot there is multiple LQR gain transitions and contact switching. Tracking of the CoM is shown along the frontal plane with the ZMP plan showed for reference. Additionally, the base error, the foot step z direction (upwards), and the contact switching plan are shown.

methods. In Fig. 5 we show the different feedback gain matrices resulting from linearizations around different poses and contact configurations. This plot highlights the synergies among the joints that are used to regulate the CoM motion for different contact situations. The plot also provides insight into the tracking strategies used by the LQR controller (i.e. which joints react to error in which states), an analysis tool that is not available by many other optimization methods. While any walking trajectory generation method would be appropriate, we started by considering the popular zero moment point (ZMP) walk modeled after Kajita's ZMP walk with preview control [3]. Using the resulting CoM trajectory with predefined foot step trajectories, we generate desired joint space trajectories using inverse kinematics and feed-forward torques provided by gravity compensation. While in the double support phase torque redundancy can be exploited to optimize contact forces. The gravity compensation term uses this redundancy to generate task-consistent interaction forces [10]. We found that this step is crucial in order to more closely realize the desired ZMP trajectory. An overview of the control scheme is shown in Fig. 3.

Our experiments have demonstrated very good performance in simulation and preliminary experiments on the real robot demonstrate the ability of this control architecture to handle stabilization over multiple steps. Fig. 2 shows a sequence of the robot walking and Fig. 4 show the plot

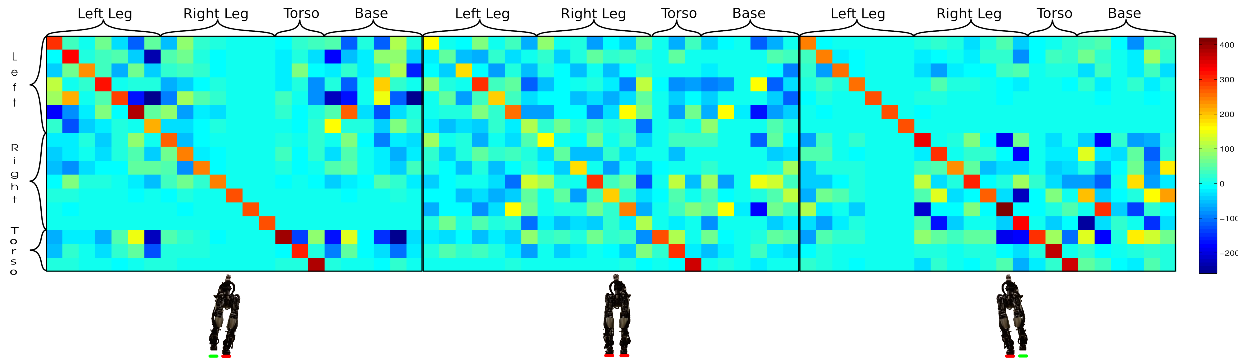

Fig. 5. Gain matrices produced from key poses and constraint conditions. From these gain matrices, one can gain intuition about the relationship between error in states and the torques that the optimal controller produces. This type of analysis is not available in other strategies that return joint torques rather that a local feedback policy.

of a single step, highlighting the tracking during a change in the contact states. However, sensor noise, stiction and backlash problems currently limit the feedback gains that can be used in our LQR designs. We are currently investigating filtering methods to reduce these issues and improve real robot performance.

IV. Conclusion and Future Work

The goal of this project is to understand how linear optimal control approaches compare to more complex whole-body approaches for multi-contact tasks. In particular we are interested to get insight on the difference of performance on real robotic systems with model inaccuracies and sensor noise. We have shown that for a bipedal system, using a simple constraint consistent LQR controller yields disturbance rejection results competitive with more complex methods. We have additionally shown that by using a small number of linear models and changing contact conditions we can track a simple walking trajectory. Our current research is related to the work of [9] and investigates how to exploit both the lightweight computational advantages of our approach while adding the flexibility that advanced operational-space methods based on QPs offer. By varying the complexity of our feedback controllers, we intend to better understand the benefits of added complexity for real robot control.

References

[1] S. Feng, X. Xinjilefu, C. Atekson, and J. Kim. Optimization based controller design and implementation for the atlas robot in the darpa robotics challenge finals. In *Humanoid Robots (Humanoids), 2015 15th IEEE-RAS International Conference on*, Nov 2015.

[2] A. Herzog, L. Righetti, F. Grimminger, P. Pastor, and S. Schaal. Balancing experiments on a torque-controlled humanoid with hierarchical inverse dynamics. In *IEEE/RSJ Intl Conf on Intelligent Robots and Systems*, 2014.

[3] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa. Resolved momentum control: humanoid motion planning based on the linear and angular momentum. In *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 2, pages 1644–1650 vol.2, Oct 2003.

[4] O. Kanoun, F. Lamiraux, and P.-B. Wieber. Kinematic control of redundant manipulators: Generalizing the task-priority framework to inequality task. *Robotics, IEEE Transactions on*, 27(4):785–792, Aug 2011.

[5] T. Koolen et al. Summary of team ihmc's virtual robotics challenge entry. In *Humanoid Robots (Humanoids), 2013 13th IEEE-RAS International Conference on*, pages 307–314, Oct 2013.

[6] S.-H. Lee and A Goswami. A momentum-based balance controller for humanoid robots on non-level and non-stationary ground. *Autonomous Robots*, 33(4):399–414, 2012.

[7] S. Mason, L. Righetti, and S. Schaal. Full dynamics lqr control of a humanoid robot: An experimental study on balancing and squatting. In *Humanoid Robots (Humanoids), 2014 14th IEEE-RAS International Conference on*, pages 374–379, Nov 2014.

[8] C. Ott, M.A. Roa, and G. Hirzinger. Posture and balance control for biped robots based on contact force optimization. In *Humanoid Robots (Humanoids), 2011 11th IEEE-RAS International Conference on*, pages 26–33, Oct 2011.

[9] M. Posa, S. Kuindersma, and Tedrake R. Optimization and stabilization of trajectories for constrained dynamical systems. 2015. Submitted paper under review. Available at `http://groups.csail.mit.edu/robotics-center/public_papers/Posa15.pdf`.

[10] L. Righetti and S. Schaal. Quadratic programming for inverse dynamics with optimal distribution of contact forces. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on*, pages 538–543, Nov 2012.

[11] B.J. Stephens and C.G. Atkeson. Dynamic balance force control for compliant humanoid robots. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1248–1255, Oct 2010.

[12] R. Tedrake et al. A summary of team mit's approach to the virtual robotics challenge. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2087–2087, May 2014.

[13] K. Yamane. Systematic derivation of simplified dynamics for humanoid robots. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on*, pages 28–35, 2012.