Reliable Dynamic Motions for a Stiff Quadruped

Katie Byl, Alec Shkolnik, Sam Prentice, Nick Roy and Russ Tedrake

Computer Science and Artificial Intelligence Lab, MIT, Cambridge, MA 02139
katiebyl@alum.mit.edu, {shkolnik, prentice, nickroy, russt}@mit.edu

Summary. We present a kinodynamic planning methodology for a high-impedance quadruped robot to negotiate a wide variety of terrain types with high reliability. We achieve motion types ranging from dynamic, double-support lunges for efficient locomotion over extreme obstacles to careful, deliberate foothold and body pose selections which allow for precise foothold placement on rough or intermittent terrain.

1 Introduction

Legged robots such as RHex [1] can handle many types of extreme terrain with impressive robustness, but an approach that is blind to the upcoming terrain is inherently incapable of precise foothold selection. On the other hand, position-controlled robots which are capable of precise foot placements have thus far been able to traverse significantly rough terrain only by using relatively slow, crawling gaits [10, 8, 6], typically enabled by decoupling body and leg movements, e.g., [11]. Robots such as the Raibert hopper [9] demonstrate that simple control principles can stabilize a dynamic legged machine and allow it to select foot placement on mildly rough and intermittent terrain [3]. This has been generalized for a dynamic quadruped [2] as well. However, the control problem of how to make the gaits of such robots adaptable to varying terrain has not yet been addressed adequately in the literature. There has been surprisingly little work on kinodynamic planning of underactuated gaits with the careful foot placement necessary for rough terrain on real robots.

In our approach to the foothold selection and dynamic motion planning tasks for legged locomotion, we first design a family of simple, low-level motion strategies which can achieve either slow, careful footsteps or fast, dynamic motions, as required by the upcoming terrain. A high-level planner can then select from among these low-level options to obtain an appropriate trade-off between speed and accuracy to travel across rough terrain.

The platform we use to demonstrate this approach is a small, point-footed quadruped robot called LittleDog, which is made by Boston Dynamics and is

currently used as part of the DARPA-sponsored Learning Locomotion project. The goal of the project is to develop a walking controller which enables the robot to traverse a variety of rough terrain both quickly and robustly. The LittleDog robot has 12 motors with large gear ratios, two at each hip and one at each knee. The high gear ratios (85:1) make the joints so stiff that the robot can hold its own weight in an upright stance even when the power is switched off. The body has a mass of 1.8 kg, which is significantly greater than the mass of each leg (at 0.25 kg). The robot has a variety of sensors which include (most notably) encoders at each of the 12 joints, an IMU, 3-axis force sensors at each foot, and current, voltage and temperature sensors. Additionally, it operates within a Vicon motion capture (mocap) environment which provides 6-DOF position information for the robot body and for the terrain with a 50 ms latency; a sub-millimeter accuracy terrain map is provided. The robot has a 1 kHz internal PD control loop on desired joint trajectories. All other control is done by an external computer in a 100 Hz loop through a wireless connection.

Fig. 1. LittleDog walks carefully on pegs.

When exact foot placement is critical, we can use the stiff joints of the robot to position the body and legs carefully, as shown in Figure 1. For more dynamic motions, the velocity and power limits do not allow for a truly ballistic, airborne "jump". However, by carefully planning the desired ground reaction forces, we can achieve repeatable, bipedal "lunging" motions which allow the front legs to clear gaps or vertical obstacles or to initiate a climb where there is a step in terrain height, as depicted in Figures 2, 5 and 8.

Fig. 2. LittleDog climbs a step dynamically.

2 Technical Approach

Planning ground reaction forces is central to our methodology. For slow motions, such as the peg-walking in Figure 1, we use the standard, quasi-static crawling gait [7] in which the projection of the center of mass (COM) of the robot is always within the support polygon. Slow walking is performed sequentially by first moving the body into a 3-legged support polygon and then moving the remaining "swing" leg.

As the speed of walking increases, we account for accelerations of the body to ensure the zero-moment point (ZMP) remains in the support polygon at all times. One strategy to achieve this is a preview control method introduced in [5] and later developed for bipedal robot walking in [4]. If we assume the center of mass has a constant height and neglect rotational accelerations, we can decouple the x and y motions and can derive the simple relationship between the motion of the COM, x_m (or y_m), and the corresponding ZMP, x_{zmp} (or y_{zmp}). Referring to the left-hand images in Figure 3, the ZMP location must satisfy $(x_m - x_{zmp})(\ddot{z}_m + g) = z_m \ddot{x}_m$. When the COM remains at a constant height, $z = z_m$, $\ddot{z}_m = 0$, and this simplifies to the following:

$$x_{zmp} = \frac{x_m g - z_m \ddot{x}_m}{g} \tag{1}$$

To model dynamic lunging motions where the body intentionally rotates – for instance, to clear an obstacle or to complete a step of a pacing gait – we must enhance the planar model from [4] to include both the mass *and* rotational inertia of the robot body. Because the legs have considerably less inertia than the body, we model them as massless. With this simplification, our dynamic model of the robot is essentially a "brick". We ignore the orientations of the individual joints and model only the relative location of the ground contact point with respect to the body, as illustrated at right in Figure 3.
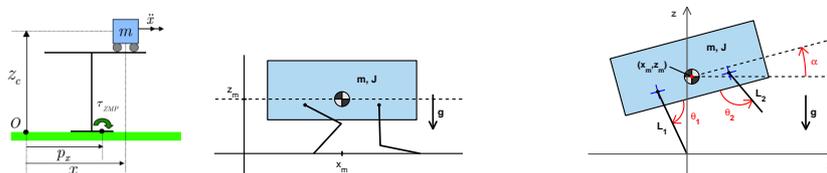

Fig. 3. 2D models of robot during laterally-symmetric dynamic motions. Leftmost image is the "cart-table model" from Figure 3 of [4] (reproduced with permission of the first author), which corresponds to fully-actuated motions at constant height (center). In underactuated lunging (at right), the modeled ground contact is a point.

Note that if we are given a prescribed (and twice differentiable) $x_m(t)$ trajectory over time, we can use Equation 1 to solve (trivially) for ZMP location. This is the **forward problem** for ZMP planning. The **inverse problem** involves first specifying a desired x_{zmp} trajectory and then solving for

a particular trajectory of mass, $x_m(t)$, such that we minimize error in ZMP tracking over time. In the inverse problem, positions of the mass, x_m, at the start and at the end of a desired $x_{zmp}(t)$ trajectory typically coincide with the end conditions on the locations of the ZMP, so the robot is in a stable, zero-acceleration configuration both before and after the particular planned motion trajectory. During the intervening trajectory, the ZMP location is dominated in the short term by the $\ddot{x}$ term (which can vary rapidly) and in the long term by the x position itself. Because of the sign difference in the influence of x and $\ddot{x}$ (i.e., in $\frac{x_m g - z_m \ddot{x}_m}{g}$), these short-term and long-term goals are at odds; we will generally have to shift the ZMP in the *opposite* direction from our eventual goal. There are a variety of ways to solve control problems of this type; the *preview control* solution [5, 4] is simply the one we prefer to use.

The preview control method augments this inherently second-order system (i.e., a simple mass and inertia being pushed around) by including acceleration as an additional state. By using acceleration as a state variable, we can create a state-space model which has the ZMP itself as a desired output, and we can then use a linear-quadratic cost – mostly on the desired output, with a small additional cost on actuation effort – to solve for the required x trajectory over time which will minimize our cost function. Performing this for both (x, z) and (y, z) relationships independently results in a way to generate $x(t)$ and $y(t)$ trajectories automatically for whatever overall speed we deem prudent, based on how rough the terrain is.

What is exciting and somewhat unexpected is that we can also plan *double-support* motions reliably using this same, basic preview control strategy. For a dynamic lunging motion, there is no longer a support polygon, since the robot balances precariously on just two feet during a motion. Our planning now requires a ZMP which is *precisely* on the line connecting these two feet for a short duration (about 0.3 sec). We will simultaneously open the inter-leg angle between the support legs and the body quickly during our ZMP-balancing time window. Figure 3 (right) illustrates such a motion with a 2D planar model. Although rotational accelerations now clearly violate our assumptions in Equation 1, pinning the gross motion of the center of mass to the idealized preview control solution in fact goes a long way toward decoupling the calculations for the x motion of the body versus its orientation angle, α: the overall motion of the mass in x essentially balances an inverted pendulum system. Commanding an additional hip angle on top of the nominal motions obtained from the inverse kinematics solution for the x motion generates an additional, rotational momentum. This motion acts to destabilize this planned equilibrium, but if these motions are commanded to happen rapidly enough, then 1) the body successfully does pitch upward and 2) the additional moments caused by pitching the body shift both the COM and the "fictitious ZMP" forward of the line of double-support stance, stabilizing the total motion to ensure the robot lands by falling forward rather than backward.

Similar tactics are employed to achieve a fast, stable dynamic walking gait. When walking, we utilize a "diagonal" gait, in which a single body movement

accompanies two leg steps: a front leg step, followed by a back leg on the diagonally-opposing side. As the body accelerates from rest to the front swing foot, the ZMP moves in the direction *opposite* the acceleration. By design, the ZMP still remains within the support triangle, however. As the body decelerates with the front leg still in the air, the robot enters a brief double-support phase during which it is rocked onto the front foot which is about to touch down. The entire motion produces a fast gait and ensures that feet are unloaded before they are lifted off the ground, enhancing the dynamic stability of the system and allowing for dynamic walking over rough terrain.
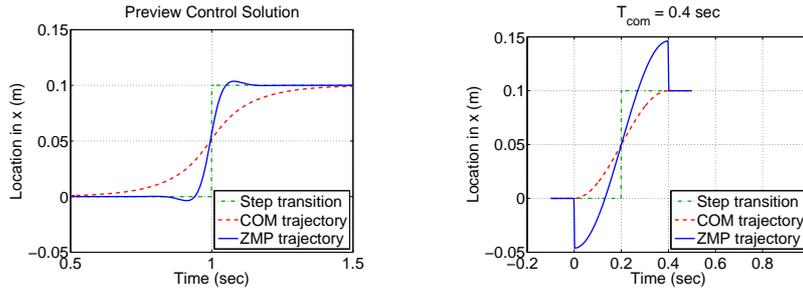

Fig. 4. The **inverse** (left) and **forward** (right) problems ZMP trajectory.

At left in Figure 4 is an example of a solution to the inverse problem for the 1D motion of a point mass which is at a height of $z_m = 0.135$ m; the preview control model used is described in more detail in [4]. The desired ZMP trajectory in this example is a step function, going from 0 to 0.1 m at $t = 1$ second. The center of mass moves monotonically from the 0 to 0.1 m while the ZMP exhibits some overshoot, going somewhat outside of this 0-0.1 m range. The magnitude of the overshoot in ZMP is only about 5 mm while the distance to the edge of the support polygon is typically on the order of 5 cm (10x further). However, there are two significant penalties to be paid for such an overly-cautious solution. First, calculating the inverse solution in real-time is computationally expensive, which would result in significant pauses between steps on the real robot. Second, the actual time required to move the mass is significantly slower (approximately 1.0 seconds) than it needs to be. To optimize our speed in fast walking, we deliberately plan for more aggressive excursions of the ZMP to achieve motions of the mass which are approximately twice as fast as the inverse solution shown in Figure 4. The solution we use is a smooth half-cosine wave shape motion for the mass, timed such that the ZMP moves nearly as cleanly from one support polygon to the next as for the inverse solution while overshoot (i.e. extra push-off and braking by the stance legs) is allowed for additional speed. At right in Figure 4 is a forward solution of the following form:

$$x_m = x_0 + (x_f - x_0) * 1/2 * (1 - cos(\pi/T_{cos})) \tag{2}$$

where x_0 and x_f are the initial and final positions of the mass, respectively. Note that both the trajectories of both x_m and x_{zmp} are each half-cosine waves. From Equation 1, we find that the peak overshoot in ZMP occurs for the mass trajectory in Equation 2 comes at both $t = 0$ and at $t = T_{cos}$. Solving for the discursion of the ZMP at $t = 0$, we find:

$$x_{zmp}(0) = \frac{-z_m}{g}\ddot{x}_m = \frac{-z_m}{g}\frac{(x_f - x_0)}{2}\left(\frac{\pi}{T_{cos}}\right)^2 \qquad (3)$$

3 Results

We apply the dynamic locomotive behaviors described in the previous section to negotiate a wide variety of terrain type. Some examples are shown in Figure 5. We calculated the average speed of crossing the terrains as the robot moves from start to goal over a distance of 1.2 meters. The speed ranged from 5.6 cm/sec (for very rocky terrain) to 12.5 cm/sec when walking on flat terrain. For a sense of scale, the leg length is approximately 16 cm when fully extended and about 12 cm in a nominal stance pose with a slight bend at the knee. The body is approximately 20 cm in length.

There are two classes of terrain in which our method yields the most significant speed improvement over a crawl gait solution. One is for intermittent obstacles, including the thin, vertical ("jersey") barriers and the gap obstacles, as illustrated in Figure 5. Here, swinging both front legs over the obstacle in a double-support lunge allows us to perform at over twice speed of a statically-stable crawl. The second type of terrain where dynamic planning greatly improves speed is for rough terrain where a) footholds are flat enough to avoid significant slippage and b) the path between footholds do not present significant vertical obstacles. Here, there is little risk of collisions with terrain features which would knock us off course during an open-loop trajectory; we cross such rough terrain at over 80% of our speed on flat terrain.

Fig. 5. LittleDog traversing various terrain dynamically.

In addition to the speed benefits of planning dynamic, double-support motions on the terrain types mentioned, we also report significant reliability. For the terrain types where we plan for the fastest speeds (gaps, jersey barrier and modular boards), we observe success rates at or above 95% at top speeds. For such terrain, it is worth noting that success requires appropriate clearance of the upcoming obstacle and joint trajectories that are not near the saturation

limits (primarily in joint velocity). For example, although the peak angle of pitch during a double support lunge had some variability, as shown in Figure 6, the duration of the motion was repeatable enough to ensure we attained our desired end pose as the robot rotated back onto all four feet after clearing an obstacle. The problems of finding mid-air trajectories for the non-support feet which would avoid collisions and would end in an appropriate configuration at landing were relatively simple, as compared with planning the dynamic motions of robot body and the corresponding joint trajectories for the two legs used in double support. For fast, double-support walking, repeatability also depended critically on avoiding unexpected collisions with the terrain or other situations (such as steep terrain footholds) which would unexpectedly knock the robot body and/or stance feet off the planned path. Despite the open-loop control strategy we employ within a particular diagonal gait step, the dynamic motions at the heart of our strategy were quite repeatable.
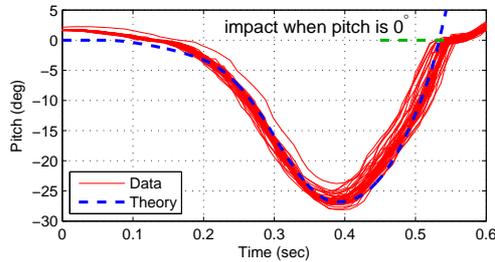

Fig. 6. Data showing pitch for 40 consecutive lunge trials.

Data presented in Figure 6 come from executing a lunging behavior repeatedly. Executing identical open-loop joint trajectory plans experimentally yields highly-repeatable results. *For walking*, ZMP preview control provides an automatic mechanism for generating the (x, y) trajectory of the COM over a wide range of speeds of travel. More surprisingly, we have successfully extended this approach beyond its intended domain to aid in decoupling the linear and rotational motions of the dog during *double-support dynamic lunges*. Our experimental repeatability in controlling motions which are inherently dynamic and underactuated is a particularly exciting and enticing result toward the development of field-ready legged robotics.

Our results for the dynamic, diagonal leg-support walking gaits were also quite repeatable. By reasoning about the physics of the robot mass and inertia during double support, we were successful in scaling the speeds of motions such that accelerations remained low enough to maintain a ZMP within the support polygon during the triple-support phase at the start and end of motion while also moving the body rapidly enough to avoid accumulating excessive pitch during the double-support transient.

There are two key ideas we employed in successfully scaling the speed of this walk for a variety of different terrains. First, swing leg trajectories were slowed down whenever the terrain connecting two sequential footholds presented vertical obstacles which might cause collisions. The curving path one must plan to avoid an obstacle is inherently longer than the near-straight start-to-goal path used for the feet when walking on flat terrain. Also, joint trajectories were intentionally executed more slowly on more extreme terrain, to ensure that joint velocities were not saturated, which would cause potentially dangerous deviations from the path foot trajectory on bumpy or stepwise terrain.

The second key idea was maintaining the same characteristic double-support transition time even as each particular leg trajectory slowed down. The effect of the mass and inertia of the leg was not significant in our experiments, so that the overall body motions and allowable double-support transients remained essentially constant (at about 0.25 sec in duration) over the various terrain boards we tested. Our most significant speed improvements from using this dynamic double-support gait came when the quarter-second transition was a significant fraction of a two-legged motion, but the same reasoning allowed for steady, repeatable walking over a wide range of overall speeds on rough terrain.
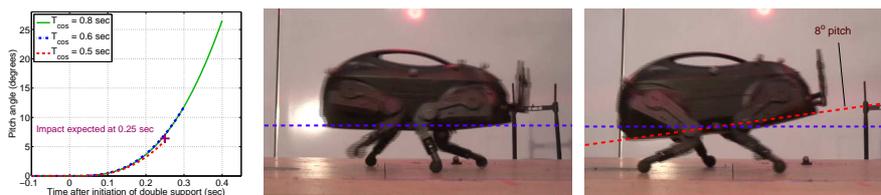

Fig. 7. Predicted pitching (nose-down) motion during double support phase.

4 Discussion

Many of the most successful bipedal robots to date (e.g., Honda's ASIMO, Kawada's HRP-2, and KAIST's Hubo, to name a few) are based on motion planning and feedback stabilization strategies which regulate the zero-moment point (ZMP) of the ground reaction forces to be safely inside the supporting polygon. This regulation is accomplished with high-gain servo motors, but is aided by direct measurements of the ground reaction forces through load cells in the feet. The methods presented here can be interpreted as extending the ZMP-style motion planning into dynamic regimes where the supporting polygon is a line, and the distance between the ZMP and the edge of the support polygon (aka the ZMP-margin) is at best identically zero.

Perhaps the most important lesson here is that this reasoning about ground reaction forces proved to be relevant and powerful despite the fact that our

robot was seemingly missing all of the prerequisites for force control. Specifically, recall that commands here were delivered over a low-bandwidth (100 Hz) command interface to highly-geared (stiff) position-controlled joints on the robot. Furthermore, we did not have an accurate dynamic model of the robot, nor accurate force sensors in the feet; no load cell feedback was used in generating these motions. In fact, it is the gross motion of the center of mass that dominates the production of ground reaction forces, and desired center of mass trajectories can be tracked nicely with stiff position control. This idea is presumably well-understood by the ZMP-walking community, but is perhaps an under-appreciated aspect of the success of their robots.

5 Future Work

Our future aim is to plan for arbitrary *final* footholds for the front feet and (especially) for *initial* footholds. Figure 8 shows LittleDog climbing particularly extreme terrain, using the same lunge used to climb a step (Figure 2). This sequence currently executes with approximately 70% reliability. Our current dynamic climbing motion also requires that all four feet of the dog are initially on flat, even terrain. A significant goal in generalizing foothold selection is to achieve bounding over rough terrain. This task is particularly difficult because the front and back feet have wide separation, meaning significant energy is lost in impact when the front feet of our non-compliant quadruped hit the ground.

Fig. 8. LittleDog using a dynamic lunge to get onto elevated rough terrain.

6 Conclusions

In this paper, we present a reasoned approach for planning highly-repeatable dynamic motions for a quadruped robot which lacks either passive compliance or sensory feedback bandwidth – two typical approaches for achieving dynamic gaits. Instead of using such natural or control-based compliant dynamics, we reason about a simplified physical model of the robot body to design planned ground forces and body motions which are compatible for underactuated, double-support phases in motion. If saturation limits (primarily in velocity) of the robot are carefully avoided, the high impedance of the robot ensures

that planned joint trajectories are executed with high fidelity, and planned underactuated trajectories result in good repeatability over each short (0.3 sec) double-support phase. Interspersing double-support phases with three-legged support periodically corrects for any small deviations (e.g., in pitch) to keep the overall trajectory of the robot near a higher-level set of pre-planned, nominal gait poses. This straight-forward strategy allows us to negotiate significant examples of rough terrain (gaps, hurdles, etc.) at over twice the speed of a traditional "crawl" gait. Equally significant and perhaps less intuitively, we obtain *more reliable* results by using this careful reasoning about the ground reaction forces and body motions than in using a crawl gait. Although, a crawl gait is often assumed to be a "conservative" strategy on rough terrain, violations of the static assumption and swing-leg collisions with terrain become more probable as speed increases – both resulting in unexpected moments which may topple the robot. Our approach provides a practical solution to the dual goals of increasing both speed *and* reliability of locomotion while also enabling efficient negotiation of significant terrain obstacles.

References

1. R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, H.B. Brown Jr., D. McMordie, U. Saranli, R. Full, and D.E. Koditschek. RHex: A biologically inspired hexapod runner. *Autonomous Robots*, 11(3):207–213, Nov 2001.
2. M. Buehler, R. Playter, and M. Raibert. Robots step outside. *Int. Symp. Adaptive Motion of Animals and Machines (AMAM)*, pages 1–4, 2005.
3. Jessica Hodgins and Marc Raibert. Adjusting step length for rough terrain locomotion. *IEEE Trans. on Robotics and Automation*, 7(3):289–298, 1991.
4. S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiware, K. Harada, K. Yokoi, and H. Hirukawa. Biped walking pattern generation by using preview control of zero-moment point. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1620–1626. IEEE, Sep 2003.
5. Tohru Katayama, Takahira Ohki, Toshio Inoue, and Tomoyuki Kato. Design of an optimal controller for a discrete-time system subject to previewable demand. *International Journal of Control*, 41(3):677–699, 1985.
6. J. Zico Kolter, Mike P. Rodgers, and Andrew Y. Ng. A control architecture for quadruped locomotion over rough terrain. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 811–818, 2008.
7. R.B. McGhee and A.A. Frank. On the stability properties of quadruped creeping gaits. *Mathematical Biosciences*, 3:331–351, 1968.
8. Dimitris Pongas, Michael Mistry, and Stefan Schaal. A robust quadruped walking gait for traversing rough terrain. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007.
9. Marc H. Raibert. *Legged Robots That Balance*. The MIT Press, 1986.
10. John Rebula, Peter Neuhaus, Brian Bonnlander, Matthew Johnson, and Jerry Pratt. A controller for the littledog quadruped walking on rough terrain. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007.
11. Alexander Shkolnik and Russ Tedrake. Inverse kinematics for a point-foot quadruped robot with dynamic redundancy resolution. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007.