Running into a Trap: Numerical Design of Task-Optimal Preflex Behaviors for Delayed Disturbance Responses

Johnathan Van Why[1], Christian Hubicki[1], Mikhail Jones[1], Monica Daley[2] and Jonathan Hurst[1]

***Abstract*— Legged robots enjoy kilohertz control rates but are still making incremental gains towards becoming as nimble as animals. In contrast, bipedal animals are amazingly robust runners despite lagged state feedback from protracted neuromechanical delays. Based on evidence from biological experiments, we posit that much of disturbance rejection can be offloaded from feedback control and encoded into feedforward pre-reflexive behaviors called preflexes. We present a framework for the offline numerical generation of preflex behaviors to optimally stabilize legged locomotion tasks in the presence of response delays. By coupling directly collocated trajectory optimizations, we optimize the preflexive motion of a simple bipedal running model to recover from uncertain terrain geometry using minimal actuator work. In simulation, the optimized preflex maneuver showed 30-77% economy improvements over a level-ground strategy when responding to terrain deviating just 2-4cm from the nominal condition. We claim this "preflex-and-replan" framework for designing efficient and robust gaits is amenable to a variety of robots and extensible to arbitrary locomotion tasks.**

I. INTRODUCTION

Legged control approaches exist on a broad spectrum in their use of sensory feedback. While some machines can rapidly switch between high-level controllers with kilohertz precision, others manage to stabilize their motions entirely with feedforward signals. Animals have large neuromechanical delays preventing fast state feedback [1], yet still run with enviable robustness to disturbances [2]. This demonstrates that much of the task of stabilizing a gait can be designed without the need for fast feedback and may be encoded into pre-reflex, or *preflex*, behaviors [3]. Preflex responses arise from a combination of passive mechanical dynamics and feedforward actuation [4], the latter of which we can optimize beforehand with optimal control tools.

The potential benefits of preflex behaviors can extend to legged robots as well. While robots do not have an animal's fundamental sensing delays, phenomenon like sensor noise and uncertain ground properties impose delays in model estimation. Additionally, some controls techniques, such as model-predictive control, take time to re-optimize after disturbances. A task-stabilizing preflex can buy a controller precious time to react and re-plan.

We present a method for systematically generating preflex responses for defined locomotor tasks. Given a sampling of possible disturbances, we optimize a feedforward preflex trajectory which is active from the beginning of the task, through each disturbance, and through a defined response delay. Using trajectory optimization, the preflex trajectory is designed to allow a post-delay optimal recovery maneuver to successfully complete the locomotion task. This optimization minimizes the maneuver's energy consumption, rendering the preflex efficient in addition to robust.

This work was supported by the Human Frontier Science Program (grant no. RGY0062/2010) and the Defense Advanced Research Projects Agency (grant no. W31P4Q-13-C-0099). Email at: vanwhyj@onid.orst.edu

[1] These authors are with the School of Mechanical, Industrial, & Manufacturing Engineering at Oregon State University, Corvallis, OR, 97331

[2] Monica Daley is with the Structure & Motion Laboratory at The Royal Veterinary College, Hatfield, Hertfordshire, United Kingdom

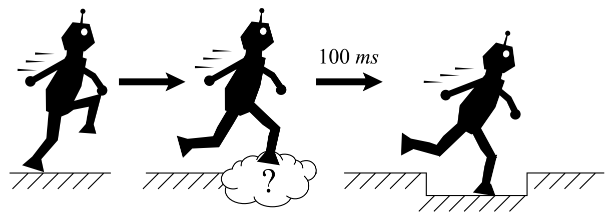

Fig. 1. Humans run with large feedback delays; as a result of these delays, disturbances can have disastrous effects on state trajectories during dynamic locomotion.

To demonstrate the approach, we generate a preflex behavior to efficiently stabilize a limit-cycle gait in the presence of uncertain ground geometry with human-like delays (100ms). An optimal level-ground strategy was generated as a baseline to compare the preflex behavior against. As a test of efficacy, we compared the preflex optimized for several ground heights to this level-ground strategy. After the 100ms delay, the tuned preflex enabled recovery using significantly less energy than the baseline strategy for disturbances with magnitudes of 2 to 4 centimeters.

We claim that this method is unique among delay-considerate control approaches in multiple ways. First, it generates feedforward motions in concert with the best post-delay recovery maneuvers, fully embracing the coupling between feedback and feedforward control. Second, these behaviors are crafted using large-scale nonlinear programming methods using sparse collocation formulations, thereby certifying a locally optimal solution. We believe this approach is applicable to a wide variety of systems and a wide variety of disturbances, allowing for the creation of robust robotic systems without extensive feedback requirements. This capability promises the creation of more reliable robotic systems with less need for sophisticated feedback infrastructure.

II. BACKGROUND

The challenge of designing gaits that are robust to uncertainties has been approached in a number of ways. These methods often fall on the extremes of the feedback spectrum, either employing high-gain feedback control or eschewing feedback in near entirety. However, biological studies demonstrate that animals make use of both preflexive stabilization [4] and planned maneuvers [5] to regulate gaits and tasks [6].

On the feedback side of the spectrum, numerous control frameworks have been employed to recover from unexpected disturbances. The bipedal robot MABEL used a series of switching controllers designed and optimized with Hybrid Zero Dynamics to rapidly recover from walking off large and unexpected drops [7]. Rough terrain walking has also been enabled using transverse linearization stabilized using Lyapunov's direct method [8], as well as L_2-gain optimization [9]. Model-predictive control approaches can re-plan in real-time after a disturbance and a short planning delay [10]. Within a metastable locomotion framework, dynamic programming policies have been designed to minimize the likelihood of falling on randomly rough terrain [11]. Further embracing a stochastic view of locomotion, this problem has even been recast as an exercise in Bayesian inference [12].

Stable control has also been generated in a fully feed-forward manner and can be accomplished via trajectory optimization [13] or the use of offline optimized central pattern generators [14]. Further, stabilizing feed-forward leg motions can be designed by inspecting a reduced-order running model [15], with and without knees [16]. These insight-driven policies can be further informed by incorporating biological insights [17], [18]. These simple models have also been stabilized by clock-driven torques [19] and applied to the hexapedal robot RHex [20].

Toward blending these two extremes, central pattern generators have been augmented with neural networks to process and incorporate feedback signals into gait control [21]. Recent work has also used evolutionary algorithms to stabilize motions to regulate target poses for animated bipeds with biological delays [22]. Further, morphology-specific methods have been employed to handle the challenge of swing-leg control in humanoid legs using a neuromuscular control model [23].

The presented approach is a numerical method for simultaneously optimizing preflexes and their subsequent post-delay responses to facilitate energy-optimal recovery. By formulating the problem with trajectory optimization, the representation of the preflex and post-delay input signals is not limited to purely oscillatory frameworks nor feedback control formulations. We use large-scale constrained optimization tools so our results are both computationally expedited and certified to be local minima. This is in notable contrast to many methods, which employ stochastic/evolutionary optimization algorithms to handle the high dimensionality of optimal control problems.

III. MODEL

The model utilized is a fully-actuated variant of the Spring Loaded Inverted Pendulum (SLIP) model, with actuation and morphology designed to represent the ATRIAS bipedal robot [24]. This model contains two actuators, one in the leg length direction and one in the leg angle direction, and a point mass body. Each actuator is in series with a combination spring/damper, introducing inherent energy losses to the SLIP dynamics.

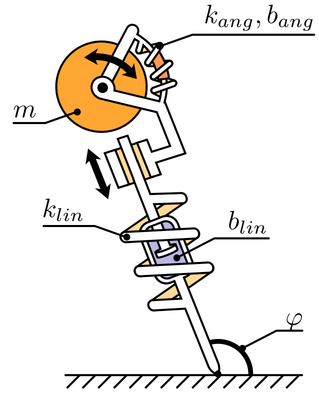

Fig. 2. Actuated damped spring-mass model used for this paper, showing the angular coordinate, point mass location, actuation, and spring/damper locations

As shown in Fig. 2, the model's position is parameterized by two generalized coordinates, radius r and leg angle φ. The equations of motion for the model, in terms of lengthwise force (f) and torque about the torso (τ), are therefore:

$$\ddot{r} = r\dot{\varphi}^2 + \frac{f}{m} - g\sin(\varphi) \tag{1}$$

$$\ddot{\varphi} = \frac{\tau}{mr^2} - \frac{2\dot{r}\dot{\varphi}}{r} - \frac{g\cos(\varphi)}{r} \tag{2}$$

with g the gravitational acceleration and m the torso mass.

The angular spring and damping constants, k_{ang} and b_{ang}, as well as the lengthwise spring and damping constants, k_{len} and b_{len}, are used to determine the lengthwise force and torque about the torso as follows:

$$f = k_{len} \cdot d_{len} + b_{len} \cdot \dot{d}_{len} \tag{3}$$

$$\tau = k_{ang} \cdot d_{ang} + b_{ang} \cdot \dot{d}_{ang} \tag{4}$$

where generalized coordinates d_{len} and d_{ang} are the lengthwise spring deflection and angular spring deflection, respectively. The deflection dynamics are derived using the lengthwise and angular actuator rates, u_{len} and u_{ang}, as well as the current torso velocity:

$$\dot{d}_{len} = u_{len} - \dot{r} \tag{5}$$

$$\dot{d}_{ang} = u_{ang} - \dot{\varphi} \tag{6}$$

The actuator rates are defined as the system control inputs.

The model parameters used were chosen to match ATRIAS and are as follows: $k_{ang} = 3200$ Nm/rad, $b_{ang} = 18.6$ Nms/rad, $k_{lin} = 6543$ N/m, $b_{lin} = 38.0$ Ns/m, $m = 59.9$ kg, and $g = 9.81$ m/s^2.

Finally, the linear actuator's acceleration was limited to a maximum magnitude of 50 m/s^2 and the angular actuator's acceleration was limited to a maximum magnitude of 50 rad/s^2.

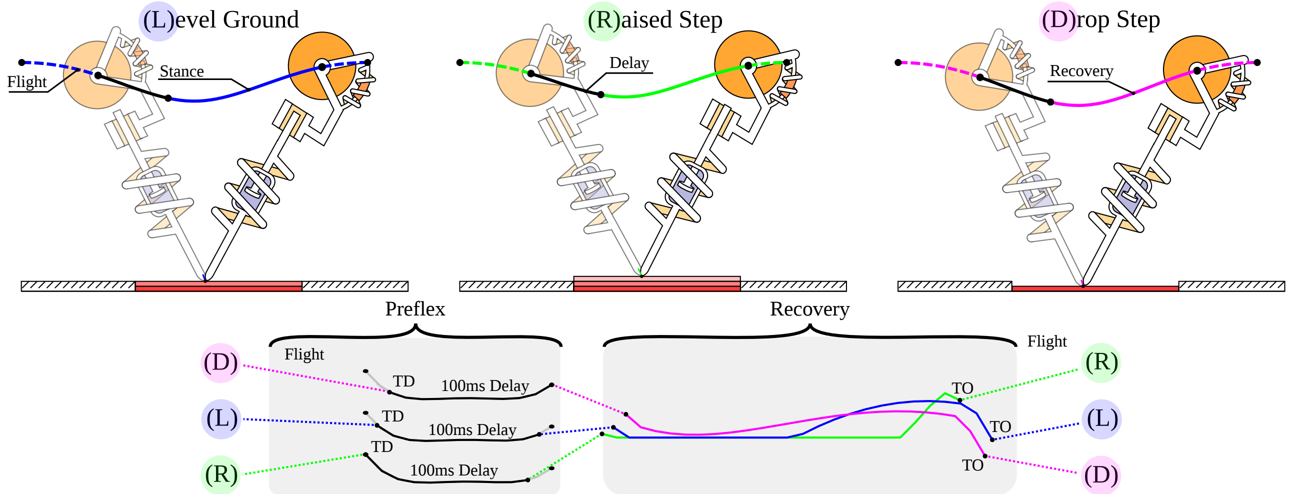

Fig. 3. Optimization problem setup, depicting the simultaneous event sequences, shared preflex inputs, and independently-optimized recovery inputs.

IV. METHODS

A. Optimization Scenario

To generate a preflex for handling a range of step disturbances, a trajectory optimization problem incorporating multiple terrain heights and feedback delay was formulated and solved. This problem is markedly similar to perturbed ground height experiments performed with guinea fowl which demonstrated notable intrinsic stability [2].

The optimization problem is outlined in Fig. 3. The problem considers a single step, defined from one apex – the moment of zero vertical velocity during flight – to the next apex. Each ground height condition creates a sequence of events; these sequences are optimized in parallel. These sequences begin with a flight phase, transition into a delay phase at touchdown, then transition into a recovery phase after a 100ms delay. During the flight and delay phases, the preflex inputs are used; these are the same across all ground heights. Recovery inputs are independently optimized for each ground height.

Therefore, the problem solves for two sets of inputs: the preflex and a recovery for each terrain height. The recovery is constrained to produce a limit cycle: the height and velocity at the second apex must equal the height and velocity at the first apex, forcing complete gait recovery in one step. The height and velocity of the apex may be adjusted by the optimization. This reflects a running robot with a feedback delay encountering terrain of unknown height.

Unsigned mechanical work was used as a proxy for energy cost in this experiment. However, the large-scale nonlinear optimization techniques used require a smooth objective, but the unsigned mechanical work calculation requires the use of the nonsmooth absolute value function. Therefore, based on an approximation in [25], the smooth approximation $|x| \approx \sqrt{x^2 + \epsilon^2} - \epsilon$, $\epsilon = 0.1$ was used for this experiment. This gives the following expressions for power for each actuator:

$$P_{ang} = \sqrt{(u_{ang} \cdot \tau)^2 + \epsilon^2} - \epsilon \tag{7}$$

$$P_{len} = \sqrt{(u_{len} \cdot f)^2 + \epsilon^2} - \epsilon \tag{8}$$

We may formulate the unsigned mechanical work as a function of height using the above expressions for power:

$$W(h) = \int_0^{t_{end}} P_{ang} + P_{len}\, dt \tag{9}$$

This integral was evaluated using the trapezoidal method.

If $P[h]$ is the probability distribution function of the terrain heights over the interval $[h_{min}, h_{max}]$, then we may calculate the expected unsigned mechanical work as follows:

$$E[W] = \int_{h_{min}}^{h_{max}} P[h] \cdot W(h)\, dh \tag{10}$$

Since the model's weight is fixed and the model is constrained to move at least a minimum distance during the scenario, the cost of transport is approximately proportional to work done. Therefore, the expected work – a proxy for expected cost of transport – is used as the objective for the optimization.

B. Trajectory Optimization

The technique used to perform this trajectory optimization is the direct collocation technique. In direct collocation, all states and control inputs are exposed as optimization variables at each integration timestep. Constraints based on ODE solution techniques are used to enforce dynamic consistency. This technique is described in detail in [26]. The backwards Euler ODE solution method was utilized for this optimization.

Since the flight phase dynamics have closed-form solutions, these solutions were used for the flight phase instead of a numerical ODE solution technique. Further, since the objective and constraints are defined using direct collocation, rather than adaptive integration, we were able to formulate closed-form symbolic gradients for the objective and all constraints to significantly accelerate the constrained optimization.

The resulting large-scale optimization problem was solved using SNOPT, a Sequential Quadratic Programming implementation chosen for its performance and local minima certification capability [27]. However, the methods described here do not depend on SNOPT and may be implemented with other optimization packages.

C. Input Linking

All of the delay phases must share the same inputs (the preflex). However, because their triggering events occur at different times, the delay phases do not have the same timescales; therefore, additional effort is required to match the inputs between the delay phases.

Matching the preflex inputs between the scenarios was done with the use of an absolute timescale. This absolute timescale is a continuous representation of the inputs as a function of time, defined relative to the start of the scenario.

After comparing methods, Bèzier curves were selected as the continuous input representation during the delay phase as they resulted in fast numerical optimization. As shown in Fig. 3, the preflex inputs are constrained to equal each other by using the Bèzier curves, while the recovery phase inputs are left to be optimized independently. The Bèzier curves directly represent actuator position trajectories, which are used to compute the touchdown conditions. The derivatives of the Bèzier curves are used as the actuator rate inputs themselves. Additionally, a time offset variable (t_{offset}) is used to offset the absolute timescale relative to the initial apex time; this improves optimization convergence speed without affecting the space of possible inputs. In total, if the Bèzier control points are $(b_1, \ldots, b_n)$, the touchdown leg angle and length are constrained to equal the Bèzier curves:

$$(r, \varphi)_{TD} \stackrel{!}{=} \text{Bèzier}(b_1, \ldots, b_n, t + t_{\text{offset}}) \tag{11}$$

and the delay phase control inputs are constrained to equal the Bèzier derivatives:

$$u_{\text{delay}}(t) \stackrel{!}{=} \frac{d}{dt}\left[\text{Bèzier}(b_1, \ldots, b_n, t + t_{\text{offset}})\right] \tag{12}$$

The touchdown time is an optimization parameter; touchdown condition constraints ensure the apex height, control inputs, and touchdown times correspond.

D. No-Preflex Baseline

A no-preflex baseline was created in order to evaluate the effectiveness of the optimized preflex. This baseline was designed to represent a trajectory generated without considering disturbances. The actuator inputs were optimized for flat terrain, then optimal recoveries for each ground height were generated by running optimizations with the same delay-phase inputs. These optimizations allowed for evaluation of the energy cost of a replanning-based strategy which lacks an optimized preflex.

E. Numerical Experiment

For the preflex optimization, 6 Bèzier control points were used to define each input. 10 equally-spaced backwards Euler timesteps were used for the delay phase, and 24 were used for the recovery phase, producing timesteps less than or equal to 10 milliseconds. The step heights used were 0.02 m, 0 m, -0.02 m, and -0.04 m. The optimization objective, as defined in (10), was evaluated by summing the costs for each step height, corresponding with a uniform probability distribution.

Seven Bèzier control points were utilized for the no-preflex baseline; the baseline had timesteps of at most 5 milliseconds and the same step heights as the preflex optimization.

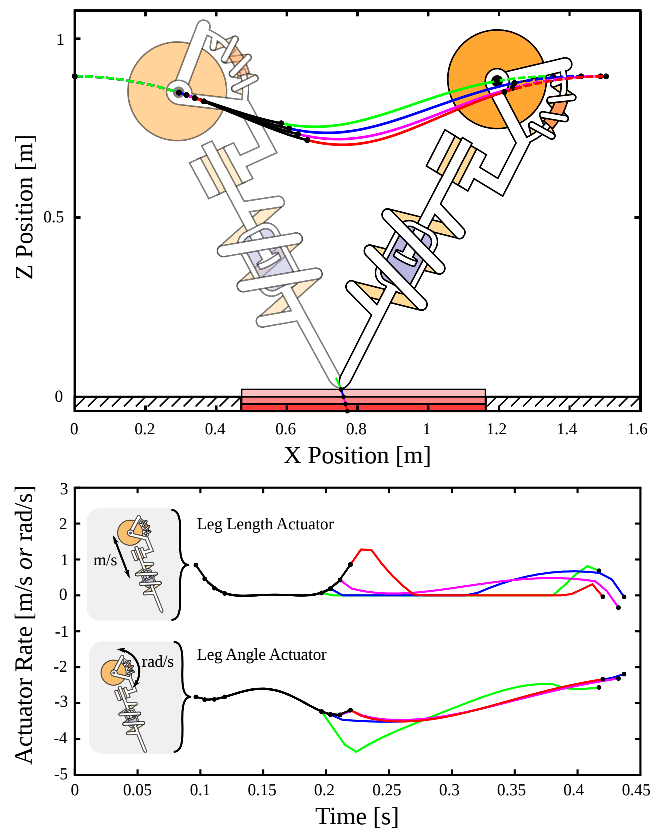

Fig. 4. CoM trajectories and actuator inputs for each step height with an optimized preflex. The toe trajectory near touchdown is nearly vertical, and the actuator motions during the recovery phase are relatively small

V. RESULTS

The center of mass trajectories and inputs for the preflex-and-replan strategy for each ground height are shown in Fig. 4. For comparison, the baseline versus preflex costs are shown in Fig. 6. The costs for each ground height are listed in Table I.

The total cost of all heights in the baseline test was 760.3 Joules; the total cost of all heights in the preflex test was 441.0 Joules, 42% smaller than the baseline cost.

VI. DISCUSSION

Overall, the preflex-and-replan strategy uses significantly less energy than a policy optimized for a single terrain height. While the baseline performs best on level ground – the height it was optimized for – the preflex-and-replan strategy performs much better under disturbances.

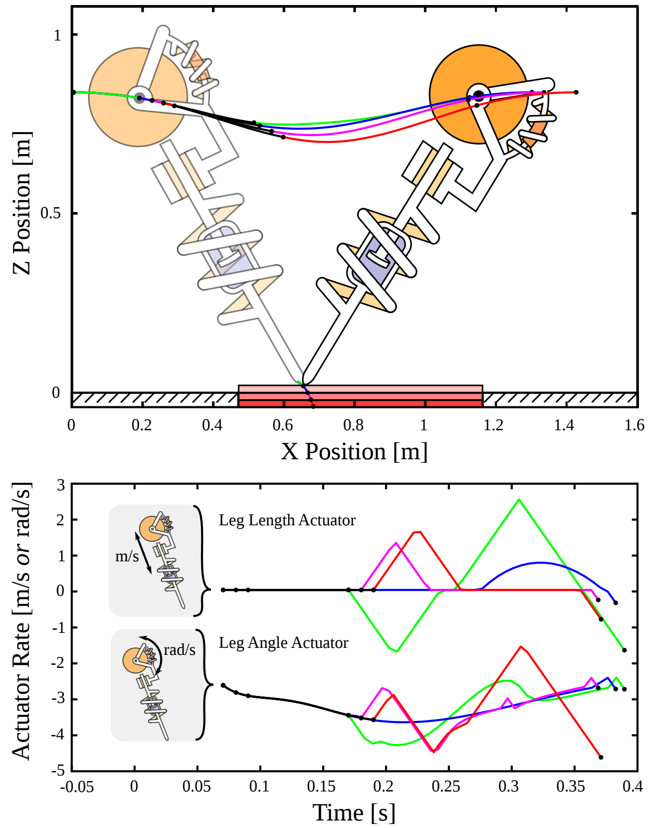

Fig. 5. CoM trajectories and actuator inputs for each step height from the no-preflex baseline. Recovery requires large actuator motions; these motions are limited by the actuator's acceleration limits

TABLE I

BASELINE VERSUS PREFLEX ENERGY COSTS. BOLDED HEIGHTS WERE USED IN THE PREFLEX OPTIMIZATION OBJECTIVE

Height	Baseline Cost	Preflex Cost	% Change
0.02 m	261 J	59.3 J	-77.3 %
0.01 m	123 J	55.9 J	-54.7 %
0.00 m	35.9 J	56.6 J	+57.9 %
-0.01 m	41.0 J	58.9 J	+43.5 %
-0.02 m	65.4 J	62.8 J	-4.06 %
-0.03 m	98.8 J	69.1 J	-30.1 %
-0.04 m	135 J	78.3 J	-42.0 %

One notable property, which may be seen in Fig. 4 and Fig. 5, is that the preflex-and-replan controller spent much less time with the actuators at their maximum acceleration. Therefore, this method may also be utilized to more effectively handle actuator bandwidth limitations, typically incurred as a result of large rotor inertias.

One of the strengths of this technique is it may be adjusted for different scenarios. The objective for this experiment corresponded with a uniform probability density function over the interval $[-0.04, 0.02]$ m, resulting in a preflex that sacrificed flat-ground performance for performance on varied terrain. However, one could alternatively assign level-ground performance a higher probability to obtain a gait that is more

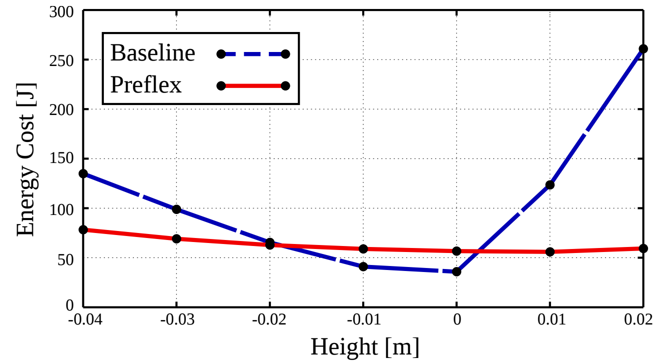

Fig. 6. Preflex and baseline costs versus step height. This demonstrates that the preflex allows for much more efficient recovery from disturbances

efficient on flat ground yet remains reasonably efficient on unexpectedly varying ground.

For robots that operate in a variety of environments, the preflex design could be implemented online. A robot may store the variation in recently-encountered terrain to create an estimate of the terrain's probability density function. This estimate may be used adaptively to generate new preflexes optimized for the current terrain, permitting a single controller to operate robustly on highly varied terrain without sacrificing efficiency on flat terrain.

Tuning preflexes in simulation may also yield insight into the control strategies and evolution of terrestrial animals. As an example hypothesis, quail and other small bipeds live in relatively rough terrain compared to their size. This environment may be the source of their characteristically crouched nominal gait and posture. Using this method on a bird locomotion model could answer whether such crouched gaits confer an advantage on rough and uncertain terrain.

Further, while the post-delay recovery motions were generated for specific terrain cases in optimization, repeating this process may not be necessary for intermediate terrain conditions. One option is to interpolate between the optimized recoveries: if a disturbance of 0.01 m is detected, then the average of the 0 m and 0.02 m inputs will closely approximate the optimal recovery. Another option is to use an online optimization to generate the recovery inputs. In the tested scenario, the preflex gives the optimizer a time window of up to 100ms to generate an optimal recovery. Online computation may be expedited further by bootstrapping the optimizer with an interpolated recovery.

This technique is not limited to terrain-based disturbances nor single single-step recovery tasks. Preflex optimization can be used to generate gaits robust to randomly timed impulses, an ability which cockroaches have exhibited [4]. The locomotion task can also be simply amended to allow for different stability constraints, such as two-step recoveries. Fig. 7 shows a preflex maneuver optimized such that it can take an extra step to recover in order to lower energy costs.

The preflex-and-replan strategy is particularly applicable to robots which are unable to sense upcoming terrain. This limitation may come as a result of many factors including

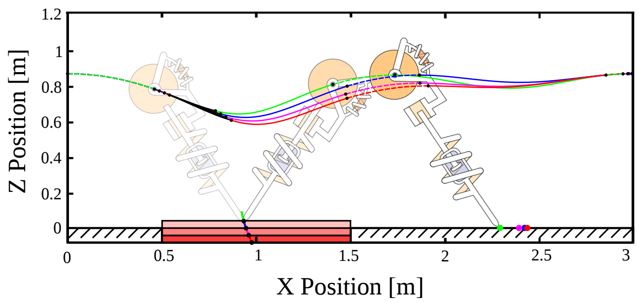

Fig. 7. A demonstration that this technique works for multi-step recoveries

component cost, space constraints, weight, and the complexity of processing 3D sensor data. Robots incorporating this strategy may achieve robust and efficient walking and running without sophisticated environmental sensing.

VII. CONCLUSIONS

A method for generating actuator inputs for preflexes based on the simultaneous coupled optimization of multiple trajectories was introduced and implemented in simulation. This method was used to generate preflexes robust against unexpected changes in ground height. These preflexes were used to form a "preflex-and-replan" framework for disturbance rejection, which was then tested in a numerical experiment alongside a baseline replanning strategy optimized for level ground.

Comparing the experimental energy costs for each strategy revealed that while the baseline strategy was more efficient at handling the scenario it was designed for, the preflex strategy performed much better under disturbances, offering a net improvement of 42%.

This technique has a variety of applications, both as a disturbance-rejection tool and as an optimal gait design tool. It may be used to handle many types of disturbances as well as limitations in state estimation. actuation, and control strategy. Control strategies designed with preflexes in mind will be an important component of making robots run as robustly as animals.

ACKNOWLEDGMENTS

The authors would like to thank Siavash Rezazadeh for his guidance in mathematical modeling.

REFERENCES

[1] H. L. More, J. R. Hutchinson, D. F. Collins, D. J. Weber, S. K. H. Aung, and J. M. Donelan, "Scaling of sensorimotor control in terrestrial mammals.," *Proceedings. Biological sciences / The Royal Society*, vol. 277, pp. 3563–8, Dec. 2010.

[2] M. A. Daley and A. A. Biewener, "Running over rough terrain reveals limb control for intrinsic stability," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, pp. 15681–15686, Oct. 2006.

[3] R. J. Full and D. E. Koditschek, "Templates and Anchors: Neuromechanical Hypotheses of Legged Locomotion on Land," *Journal of Experimental Biology*, vol. 202, pp. 3325–3332, Dec. 1999.

[4] D. L. Jindrich and R. J. Full, "Dynamic stabilization of rapid hexapedal locomotion.," *The Journal of experimental biology*, vol. 205, pp. 2803–23, Sept. 2002.

[5] A. Birn-Jeffery and M. Daley, "Birds achieve high robustness in uneven terrain through active control of landing conditions," *The Journal of Experimental Biology*, vol. 215, no. 12, pp. 2117–2127, 2012.

[6] P. Holmes, R. J. R. J. Full, D. E. D. Koditschek, and J. Guckenheimer, "The Dynamics of Legged Locomotion: Models, Analyses, and Challenges," *SIAM Review*, vol. 48, pp. 207–304, May 2006.

[7] H. Park, K. Sreenath, A. Ramezani, and J. W. Grizzle, "Switching control design for accommodating large step-down disturbances in bipedal robot walking," in *IEEE/RSJ International Conference on Robotics and Automation (ICRA)*, pp. 45–50, Ieee, May 2012.

[8] I. R. Manchester, U. Mettin, F. Iida, and R. Tedrake, "Stable dynamic walking over uneven terrain," *The International Journal of Robotics Research*, vol. 30, pp. 265–279, Jan. 2011.

[9] H. Dai and R. Tedrake, "L2-Gain Optimization for Robust Bipedal Walking on Unknown Terrain," in *2013 IEEE International Conference on Robotics and Automation (ICRA)*, 2013.

[10] T. Erez, K. Lowrey, Y. Tassa, and V. Kumar, "An integrated system for real-time Model Predictive Control of humanoid robots," in *IEEE/RAS International Conference on Humanoid Robots*, 2013.

[11] K. Byl and R. Tedrake, "Metastable Walking Machines," *International Journal of Robotics Research*, vol. 28, no. 8, pp. 1040–1064, 2009.

[12] Y. Tassa, T. Erez, and E. Todorov, "Optimal limit-cycle control recast as bayesian inference," in *Proceedings of the IFAC ...*, 2011.

[13] K. D. Mombaur, R. W. Longman, H. G. Bock, and J. P. Schlöder, "Open-loop stable running," *Robotica*, vol. 23, pp. 21–33, Jan. 2005.

[14] A. Spröwitz, A. Tuleu, M. Vespignani, M. Ajallooeian, E. Badri, and A. J. Ijspeert, "Towards Dynamic Trot Gait Locomotion: Design, Control, and Experiments with Cheetah-cub, a Compliant Quadruped Robot," *The International Journal of Robotics Research*, 2013.

[15] M. Ernst, H. Geyer, and R. Blickhan, "Extension and customization of self-stability control in compliant legged systems.," *Bioinspiration & biomimetics*, vol. 7, p. 046002, Dec. 2012.

[16] J. W. Hurst, B. Morris, J. E. Chestnutt, and A. A. Rizzi, "A Policy for Open-Loop Attenuation of Disturbance Effects Caused by Uncertain Ground Properties in Running," in *Robotics and Automation, 2007 IEEE International Conference on*, pp. 1455–1460, 2007.

[17] H. R. Vejdani, Y. Blum, M. a. Daley, and J. W. Hurst, "Bio-inspired swing leg control for spring-mass robots running on ground with unexpected height disturbance.," *Bioinspiration & Biomimetics*, vol. 8, p. 046006, Dec. 2013.

[18] B. Andrews, B. Miller, J. Schmitt, and J. E. Clark, "Running over unknown rough terrain with a one-legged planar robot.," *Bioinspiration & biomimetics*, vol. 6, p. 026009, June 2011.

[19] J. Seipel and P. Holmes, "A simple model for clock-actuated legged locomotion," *Regular and Chaotic Dynamics*, vol. 12, pp. 502–520, Oct. 2007.

[20] R. Altendorfer, D. E. Koditschek, and P. Holmes, "Stability analysis of a clock-driven rigid-body SLIP model for RHex," *The International Journal of Robotics Research*, vol. 23, no. 10-11, pp. 1001–1012, 2004.

[21] S. Gay, J. Santos-Victor, and A. Ijspeert, "Learning robot gait stability using neural networks as sensory feedback function for Central Pattern Generators," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 194–201, Ieee, Nov. 2013.

[22] T. Geijtenbeek, M. van de Panne, and a. F. van der Stappen, "Flexible muscle-based locomotion for bipedal creatures," *ACM Transactions on Graphics*, vol. 32, pp. 1–11, Nov. 2013.

[23] R. Desai and H. Geyer, "Muscle-reflex control of robust swing leg placement," *2013 IEEE International Conference on Robotics and Automation*, pp. 2169–2174, May 2013.

[24] C. Hubicki, J. Grimes, M. Jones, D. Renjewski, A. Sprowitz, A. Abate, and J. Hurst, "ATRIAS: Enabling Agile Bipedal Locomotion with a Template-Driven Approach to Robot Design," *In Submission*, 2014.

[25] M. Srinivasan, "Fifteen observations on the structure of energy-minimizing gaits in many simple biped models," *J. of the Royal Society*, vol. 54, pp. 74–98, June 2010.

[26] O. von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, vol. 37, pp. 357–373, 1992.

[27] P. Gill, W. Murray, and M. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM review*, vol. 47, no. 1, pp. 99–131, 2005.