Methods and Performance Analyses for Design and Feedback Control of Efficient and Robust Planar Biped Walking

Nihar Talele and Katie Byl

Abstract— Practical bipedal robot locomotion needs to be both energy efficient and robust to variability and uncertainty. In this paper, we build upon recent works in trajectory optimization for robot locomotion with two primary goals. First, we wish to demonstrate the importance of (a) considering and quantifying not only energy efficiency but also robustness of gaits, and (b) optimization not only of nominal motion trajectories but also of robot design parameters and feedback control policies. As a second, complementary focus, we present results from optimization studies on a 5-link planar walking model, to provide preliminary data on particular trade-offs and general trends in improving efficiency versus robustness. In addressing important, open challenges, we focus in particular on discussions of the effects of choices made (a) in formulating what is always, necessarily only an approximate optimization, in choosing metrics for performance, and (b) in structuring and tuning feedback control.

I. INTRODUCTION

Humanoid walking should be both energy efficient and robust to perturbations. This paper explores methods and trade-offs in achieving these two goals through local (gradient-based) optimization of a 5-link planar walking model.

Both energy efficiency and robustness have long been goals for robot walking, and both mechanical design and control strategy play important roles in each objective. Below is a selective summary of relevant prior work.

Toward improved mechanical design, biped robots built on passive dynamic principles drew significant attention over a decade ago [1], but their success at reducing required energy has seemed to be coupled with fragile dynamics, yielding susceptibility to falls. Design of mechanical properties, i.e., lengths and mass distribution, clearly play an important role in enabling efficient legged locomotion, but they also arguably affect stability.

To improve controlled walking strategies, a range of work has focused on both trajectory optimization and control theory. Trajectory optimization through direct collocation [2] is one promising approach. In 1999, for example, Hardt et al. formulated the problem of minimizing energy of a planar 5-link biped, both with and without ankle torque, using DIRCOL software [3] to solve a nonlinear optimization subject to contact constraints. Two years later, Paul et al. looked at simultaneous optimization of both mass distributions (robot design) and nominal motion trajectories, to be tracked via a simple proportional controller (with

saturation limits), using simple neural networks to learn efficient locomotion [4].

In 2002, Westervelt and Grizzle highlighted the importance of optimizing walking motions while simultaneously guaranteeing asymptotic stability [5], as opposed to a still-dominating paradigm of sequential design, first optimizing a nominal trajectory and subsequently adding feedback control in a more ad hoc way. As in [3], they also use DIRCOL, and they solve a sequential quadratic programming (SQP) problem to optimize the sum of u^2 across all four actuators. Note that [5] uses a hybrid zero dynamic (HZD) approach, which parameterizes joint trajectories on a monotonic, geometric variable. In a similar spirit, [6] produce energy-optimal gaits for the 5-link walker using polynomial trajectories in which the gait is defined as $q(s)$, as a function of geometry rather than time, by solving for optimal polynomial coefficients.

Various works have instead focused on optimizing robustness. Dai and Tedrake [7] optimized a measure of robustness that quantifies variation from a nominal trajectory during rough terrain locomotion, for both the spring-loaded inverted pendulum (SLIP) and compass gait (CG) walker planar legged locomotion models. In [8], Saglam and Byl explored Pareto trade-offs between energy use and robustness, using a weighted metric that balances rewards for both low energy use and high mean-time to failure (aka mean first-passage time), using value iteration to optimize across a meshed approximation of the reachable state space for the system.

Recent work by Hamed, Buss and Grizzle also focuses on robustness, tuning control parameters to ensure not only stable eigenvalues of the Jacobian of the period-one return map of limit-cycle walking but also reduced sensitivity of this Jacobian to parameter variation [9]. Here, they decouple the selection of a nominal periodic orbit from that of optimizing a parameterized controller, e.g., torques include both the necessary feedforward terms exactly compatible with the limit cycle of interest, along with some flavor of feedback law (e.g., perhaps but not necessarily HZD) that has no effect along the exact limit cycle trajectory.

Finally, a few other recent works emphasize applicability of legged locomotion optimization to an expanding range of problems. Recent work by Ma, Hereid, Hubicki and Ames on the DURUS robot employs the HZD framework to optimize energy efficiency for stable 3D walking [10]. Xi, Yesilevskiy and Remy employ direct collocation (DC) to optimize energetic cost for gaits without a prescribed sequence of foot contacts with the grounds [11], and within our own group, we have used trajectory optimization to predict the theoretical cost of added mass in exoskeleton design [12]

and to discover nonintuitive locomotion strategies for an underactuated, acrobot-based rolling system [13].

In this paper, we focus on several related, open challenges in simultaneously optimizing for energetics and robustness. We highlight important choices made in differentiation and integration that improve speed and accuracy, since local optimization provides only approximate results. With an aim toward improving both energetics and robustness, we explore how variations in mass distribution affect metrics for each of these goals and observe a natural trade-off (between metrics) that results from tuning of feedback control. Our results demonstrate that choice of both mass distribution and feedback control structure have important, and apparently coupled, effects on both energy use and stability, providing evidence for the hypothesis that more comprehensive frameworks are needed for simultaneous optimization across system parameters and desired metrics.

The rest of this paper is organized as follows. Section II describes the 5-link planar walking model we study, while Section III outlines our choices on optimization framework, feedback structure for subsequent control of trajectories and definition of cost metric. Results are given in Section IV, followed by discussion and conclusions in Section V.
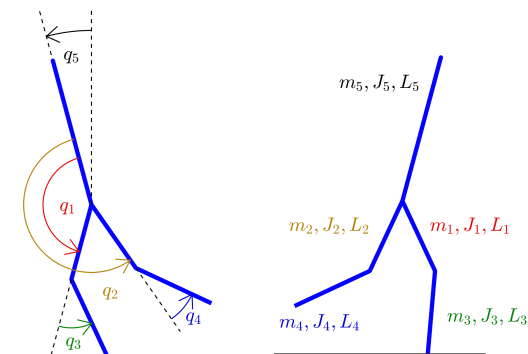
II. SIMULATION MODEL

Fig. 1. 5-link biped model. At left, q_5 is an absolute angle, measured with respect to vertical, while all other angles are relative. The lefthand image is drawn to clearly illustrate angles are positive in the counter-clockwise direction, throughout. At right, a typical pose, while walking to the right.

We use a 5-link model, shown in Fig. 1, for our simulations. The dynamics of this system are constrained to the sagittal plane only. We study several mass distributions, always enforcing that the total mass of the model is 70 (kg). The model has actuators at hips and knees, and the nonlinear dynamics can be written in the matrix form as

$$D(q)\ddot{q} + C(q,\dot{q})\dot{q} + G(q) = Bu. \quad (1)$$

This well-studied model [5] has 5 degrees of freedom, corresponding to 5 joint angles given by $q := [q_1\ q_2\ q_3\ q_4\ q_5]^{\mathsf{T}}$, but due to the passive point-foot contact at the ground, the model still remains underactuated with $u \in \mathbb{R}^4$. We model the impact dynamics between the swing leg and the ground as instantaneous and inelastic [14] to obtain joint velocities just after the impact. Since the impact model we use assumes

inelastic collisions, some amount of energy is lost when the stance foot impacts the ground.

III. TRAJECTORY OPTIMIZATION

A. Framework for Direct Collocation

We discretize the dynamics and use direct collocation (DC) to generate trajectories for the walking motion. Our approach is close to that suggested in [15], except that we use trapezoidal integration instead of backward Euler. Trapezoid rule integration can potentially result in lack of convergence. However, in our work, this had not been an issue, and results with trapezoid rule are significantly more accurate, when comparing the discretized (and thereby approximate) solutions inherent in this framework with subsequent high-resolution (1e-9) simulations of dynamics in Matlab.

The optimization problem is formulated as

$$\min_{q,\dot{q},u} \quad COT \quad (2)$$

$$\text{such that} \quad D(q)\ddot{q} + C(q,\dot{q})\dot{q} + G(q) = Hu \quad (3)$$

$$\Phi(q,\dot{q}) \leq 0 \quad (4)$$

where $q \in \mathbb{R}^n$ is the vector of generalized coordinates, $D(q) \in \mathbb{R}^{n \times n}$ is the mass inertia matrix, $C(q,\dot{q}) \in \mathbb{R}^{n \times n}$ represents the Coriolis forces, $G \in \mathbb{R}^n$ contains the gravitational forces and $H \in \mathbb{R}^{n \times n-1}$ is the input (torque) mapping. $\Phi(q,\dot{q})$ is a vector of constraints. Constraints are imposed to make sure that the normal reaction at the point of contact with the ground is always positive. The optimization problem is set up such that at the end of the trajectory an impact at the ground happens. An additional constraint is added that the state of the model after the impact should match the initial condition in order to obtain a limit cycle behavior.

We implement this framework in Matlab making use of CasADi [16], which lets us calculate gradients for optimization using algorithmic differentiation to machine precision. (CasADi uses *C*omputer *a*lgebra *s*ystem syntax to perform *Algebraic Differentiation*; thus the name.) Using algebraic (and not numerical) differentiation greatly increases the stability and convergence properties of our optimization, while also reducing run time considerably.

Using CasADi to improve automated gradient calculation, the nonlinear programming (NLP) optimization itself is solved using IPOPT [17]. This choice (vs use of SNOPT, Matlab's fmincon, etc.) is made based both on improved speed during our own in-lab testing experience and similar external benchmarking results [18].

The DC framework evaluates Eqs. 2-4 only at discrete time intervals, t_k, resulting in an approximation of the desired optimization problem. We use $\Delta t = h = 0.01$ (s) and integrate using the standard trapezoid rule. Also, we assume $u(t)$ is held via a zero-order hold for each time step (as opposed to a first-order hold). Our integration scheme is then

$$q_{k+1} = q_k + \frac{h}{2}[\dot{q}_k + \dot{q}_{k+1}] \quad (5)$$

$$\dot{q}_{k+1} = \dot{q}_k + \frac{h}{2}[f(q_k,\dot{q}_k,u_k) + f(q_{k+1},\dot{q}_{k+1},u_k)] \quad (6)$$

where $f(q,\dot{q},u) = D(q)\backslash(-C(q,\dot{q})-G(q)+Bu)$. We found that using a first order hold for input, i.e., replacing u_k with u_{k+1} at far right in (6) above, leads to trajectories that are undesirably oscillatory and not smooth. This problem is easily rectified by adding some level of regularization. However, the zero order hold method still converges more rapidly. Also important to note is that regularization increases the optimal cost by a small amount.

B. Trajectory Stabilization

After we obtain the trajectories from the optimization framework, we simulate them in matlab using ode45 with an error tolerance of 1e-9. In order to stabilize these trajectories, we use partial feedback linearization (PFL). The total input to the system is $U = U_{ff} + U_{fb}$ where

$$U_{fb} = (SD^{-1}B)^{-1}(v + SD^{-1}(C\dot{q}+G)), \qquad (7)$$

S is given by

$$S = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \qquad (8)$$

and u_{ff} contains the feedforward torques compatible with the nominal dynamics.

Given a passive contact of the stance leg with the ground, PFL allows us to directly set the accelerations of 4 out of 5 angles using v. We set $v = [\ddot{q}_{2des}, \ddot{q}_{1des} + \ddot{q}_{3des}, \ddot{q}_{4des}, \ddot{q}_{5des}]^{\mathsf{T}}$ where $\ddot{q}_{des} = -K_p(q-q_{des}) - K_d(\dot{q}-\dot{q}_{des})$. We set $K_p = \omega_n^2$ and $K_d = 2\zeta\omega_n$. For all our simulations we set $\zeta = 1$ and test across a range of ω_n values. We get $U_{ff}, q_{des}, \dot{q}_{des}$ by interpolating the trajectories from the optimization framework.

C. Cost of Transport (COT)

We optimize the trajectories for cost of transport (COT) which is calculated as

$$COT = \frac{\int_0^t \sum_{n=1}^5 |P_n(t)| dt}{Mgd} \qquad (9)$$

which we implement using trapezoid rule with zero order hold for input as

$$COT = \frac{\sum_{k=0}^{N-1} \sum_{n=1}^5 \widetilde{P}_n(k)\Delta t}{Mgd} \qquad (10)$$

where there are N discrete time steps $t(k)$, d is the stride length, and M is the mass of the model. Rate of work (power) at joint n during time step k is approximated as

$$\widetilde{P}_n(k) = \frac{\widetilde{P}_{n,1}(k) + \widetilde{P}_{n,2}(k)}{2} \qquad n = \{1,2,...,5\}, \qquad (11)$$

where $\widetilde{P}_{n,1}(k)$ is the regularized version of $P_{n,1}(k) = \tau_n(k)\omega_n(k)$ and $\widetilde{P}_{n,2}(k)$ is the regularized version of $P_{n,2}(k) = \tau_n(k)\omega_n(k+1)$. Also, $\widetilde{P}_{n,i}(k)$ is defined to penalize both the positive as well as the negative mechanical work. We smooth the cost function by using a regularization factor, so that

$$\widetilde{P}_{n,i} = \sqrt{P_{n,i}^2 + \varepsilon^2} \qquad i = \{1,2\} \qquad (12)$$

as suggested in [19]. We use $\varepsilon^2 = 0.01$, which works well for our optimization. Lower values of ε led to stability issues with the solver.
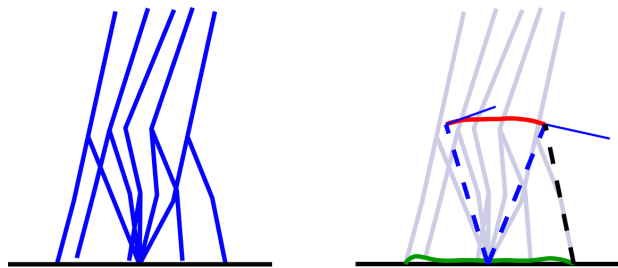
IV. RESULTS

Fig. 2. Typical motion for the 5-link walker, using trajectories generated from the optimization framework. The figure on the left shows snapshots of motion. On the right, trajectories of the COM and the end of the swing foot are overlaid.

Fig. 2 shows a typical motion generated for the mass distribution $m_5 = 50$ (kg), $m_1 = m_2 = 7$ (kg), $m_3 = m_4 = 3$ (kg), which corresponds roughly to a human mass distribution. Fig. 3 and Fig. 4 show the corresponding angle and angular velocity trajectories for that motion, and Fig. 5 shows the joint torques. All the results we present here are generated for a walking motion of stride length = 0.6 (m) over a time interval of 0.6 (s), resulting in a velocity of 1 (m/s).

A few details in these four figures are worth pointing out. First, note that all trajectories are divided into two subplots for better resolution and clarity, since q_1 and q_2 remain close to π, while the other joints are near 0. Upper plots correspond to upper leg segments; solid (blue) lines correspond to stance leg segments (femur and tibia).

Several characteristics seen in this example are common among optimizations we performed across a range of mass distributions, as itemized below:

1) The swing leg follows a very low trajectory, as depicted by the solid green line in Fig. 2. We enforce a minimum ground clearance of 2 (cm), except at the first and last 5 points of the trajectory, and the swing leg tip overshoots near the ends and grazes this value mid-gait. Without adequate ground clearance (e.g., 2 (cm)), resulting limit cycle gaits could not be stabilized.

2) The center of mass (COM) trajectory "flattens out" mid-stride, as opposed to following an arc, which is a feature also seen in human walking. This trajectory requires work in bending of the stance leg but results in less acceleration and deceleration vertically (against gravity), for lower accelerations overall of COM.

3) Also, rapid changes in velocities just before impact, as seen in Fig. 4, deflect the COM velocity at the end of gait slightly "upward", reducing kinetic energy losses at impact. The velocity vector, depicted as a solid blue line in Fig. 2 is close to orthogonal with the dashed line drawn from COM to stance leg tip at the start of the step (at 92.8°) and more obtuse at the end (100.0°

with respect to current stance contact, and $66.0°$ wrt upcoming stance leg, as a black dashed line).

4) Also, the velocity vector is longer (i.e., faster speed) at the end, showing kinetic energy has built up, to compensate for dissipation at impact.
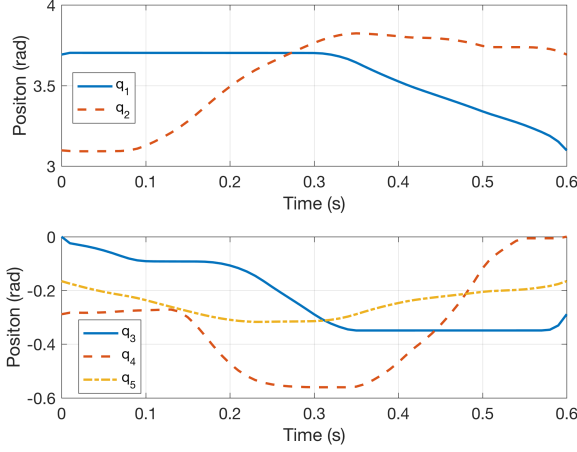

Fig. 3. Angle trajectories for the walking shown in Fig. 2. The top plot shows the positions of the swing and the stance thighs, while the bottom plot shows the positions of the stance knee, swing knee and torso.
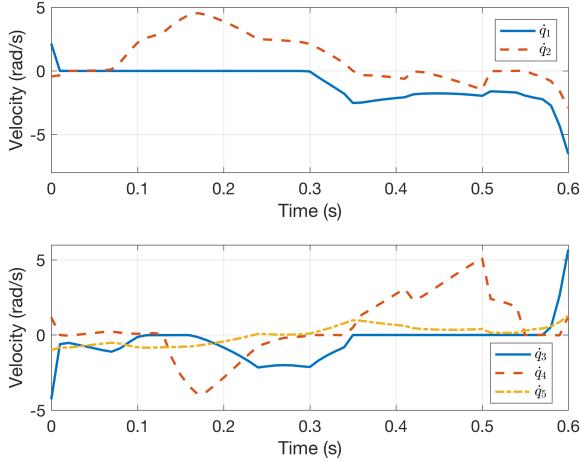

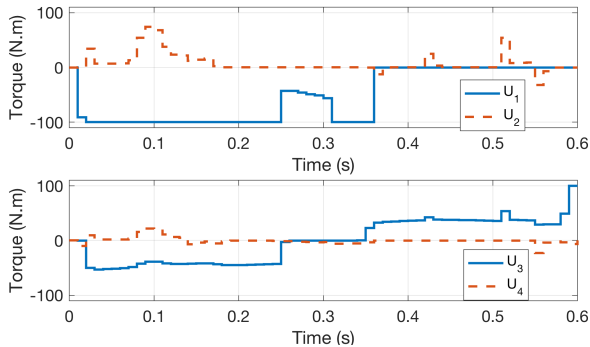Fig. 4. Velocity trajectories for the walking motion shown in Fig. 2.

Fig. 5. Torques for the walking motion in Fig. 2.

The velocity (Fig. 4) and torque (Fig. 5) trajectories also show repeatable features that were not anticipated but are (retrospectively) intuitively in agreement with our cost function, as noted in the rest of the list of features, below.

5) *Magnitude* of velocities of the stance leg segments (solid blue) increase *quite rapidly* at the very end of the step, in achieving the COM deflection upward (item 3).

The "toe-off" described above has a well-known benefit for energetics [20], but it also causes a problem:

6) We observe that rapid increases in velocity near the end of a step, required for toe-off, also have an important (and unfortunate) negative effect on stability.

For example, when we attempted to use linear quadratic regulation (LQR) to provide feedback control, the resulting gait was not stable. Specifically, we linearized about each "knot point" of the optimal solution, and then controlled motions using the nominal feedforward torque (U_{ff}) added to feedback of the form $U_{fb} = -K(X - X_{nom})$, where $K = K(t)$ and $X_{nom} = X_{nom}(t)$ were interpolated between their values at the discrete points of the optimal solution. During continuous motion, the trajectories definitely converged toward the nominal trajectories as expected, but the effects through impacts were too destabilizing, resulting in falls after 3 to 6 steps.

This is a particularly interesting result, as it demonstrates evidence for a strong hypothesis that trajectories and feedback policies should be optimized as a concurrent problem, rather than planning ad hoc feedback subsequent to solving for a nominal trajectory. Finally,

7) values of u_n and the corresponding relative angular velocity $\dot{q}_n$ show a complementary behavior: when one is significant in magnitude, the other is near zero.

This makes sense, given $P_n = u_n \dot{q}_n$, from Eq. 11. In Figures 4 and 5, note in particular the solid blue lines, for the stance thigh (upper) and knee joints. The relative thigh angle (between torso and stance leg) is nearly zero while torque is at its maximum magnitude (-100 (N·m)), with the associated bobbing motion of the torso in Fig. 2 during the first 0.25 (s) of the step. For the stance knee, a period of negative velocity for $\dot{q}_3$ (i.e., knee bending) at mid-stance corresponds to a flat region in which $u_3 \approx 0$, i.e., bending almost passively during the gait. Near the end of the step, push-off with the stance knee is concentrated in particular at the last time step, perhaps exploiting the approximate nature of the optimization problem somewhat.

A different choice of cost function would result in somewhat different solution characteristics. We also tested a simple quadratic cost on control effort (i.e., cost = u^2); here, the "toe-off" behavior, with a spike in torque and velocity of the stance knee at the end of the step, also occurred. Overall, the trajectories are qualitatively more smooth for this cost function, however, and the overall COT is noticeably higher than when optimizing for COT specifically.

A. Effect of system parameters on energy and stability

We repeated the same optimization on five different mass distributions. Link lengths and masses are shown in Tables I

and Table II, respectively. Each link is modeled as a simple rod, with COM at mid-length.

Segment	$Length(m)$	$COM(m)$
$L_1 = L_2$	0.4	0.2
$L_3 = L_4$	0.43	0.215
L_5	0.77	0.385

TABLE I

LENGTH AND COM PARAMETERS USED IN SIMULATION EXPERIMENTS

Set	$m_1 = m_2(kg)$	$m_3 = m_4(kg)$	$m_5(kg)$	COT_{opt}
1	7	7	42	0.0992
2	5	7	46	0.0996
3	7	5	46	0.0861
4	5	5	50	0.0853
5	7	3	50	0.0705

TABLE II

MASS PARAMETERS USED IN SIMULATION EXPERIMENTS

Fig. 6 shows how changing ω_n affects COT. First, recall that discretized trajectory optimization solutions of true dynamics are always approximate ([2], [3], [9], [11], [13], [15], [19], etc.), and also that ω_n is analogous to a "stiffness" (or proportional controller, "K_p"), within the PFL framework. Intuitively, increasing control gains would increase control effort, in following an arbitrary reference trajectory.

However, our reference trajectories are far from arbitrary: they are precalculated to achieve a locally optimal (i.e., minimal) cost of transport – within some approximation errors. Therefore, increasing w_n monotonically decreases resulting COT in our higher-accuracy Matlab simulations.
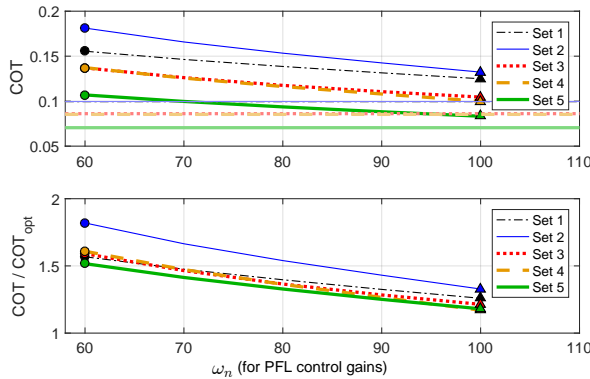

Fig. 6.   Cost of Transport (COT) as a function of ω_n.

Figure 6 both illustrates how COT converges exponentially downward as ω_n increases and also how different parameter sets result in a range of different errors, in comparing optimization results to more accurate simulations. For example, the lowest COT is for Set 5 from Table II, which is also closest to a typical human mass distribution. The lower subplot of Fig. 6 shows the ratio of "actual" COT (from simulation) to the value output from optimization, also as a function of ω_n, where Set 5 also shows the lowest approximation error.

Figure 7 shows the evolution of energy over time for different values of ω_n. The steep rise at the start is due to high peak torques as the controller tries to pull the actual trajectory back to the optimal trajectory.
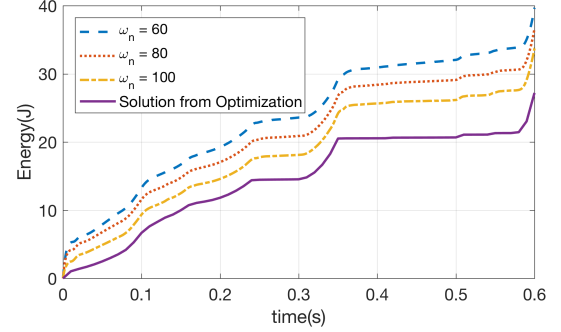

Fig. 7.   Variation of energy as ω_n changes.
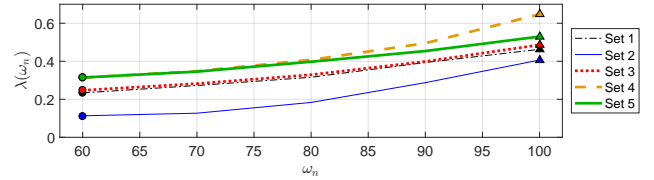
B. Energy vs Stability Tradeoff

Fig. 8.   Rate of convergence, λ, as a function of ω_n, for push recovery.

Fig. 8 shows the change in rate convergence (i.e., largest/slowest discrete-time eigenvalue) as a function of ω_n when recovering from push perturbations. Three impulsive push tests were simulated for each set of parameters by applying an impulse of 10 (Ns) (e.g., effectively a pulse of 1000 (N) for 0.01 (s)) at the hip, stance knee, or torso, respectively. The velocities post impact were calculated as

$$D(q)\dot{q}^+ - D(q)\dot{q}^- = E \cdot F_{ext}\Delta t \qquad (13)$$

where $F_{ext}\Delta t = 10$ (Ns) is the applied pulse and $E = \frac{\partial p(q)}{\partial q}$ where $p(q)$ is the point at which force is applied.
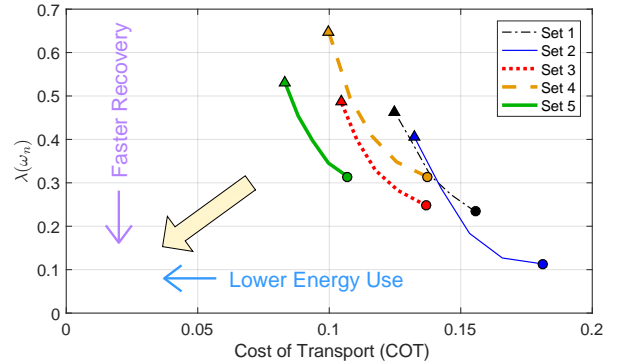

Fig. 9.   Cost of Transport vs rate of convergence $\lambda(\omega_n)$, for push recovery. Each line sweeps across results from $\omega_n = 60$ to $\omega_n = 100$.

As with Figure 6, the trend seen in Figure 8 as a function of ω_n is not immediately intuitive. Usually, it would be expected that higher values of ω_n should lead to faster rates of convergence as higher gains should pull back the trajectory to the optimal reference trajectory faster.

Combining data from Figures 6 and 8 to investigate any potential relationship between stability and energy further, Fig. 9 shows a plot of COT vs rate of convergence for the parameter sets given in Table II. A clear trade-off between energy efficiency and rate of convergence is evident here.

V. DISCUSSION AND CONCLUSIONS

Our simulation results demonstrate that although optimizing for energy alone and then stabilizing the trajectories can work, there is a need for a more cohesive framework that takes into account both energy as well as stability of the system for optimization simultaneously. Furthermore, because different physical robot design parameters have important effects on stability and energy properties, it would be prudent to include these parameters as open variables, via appropriate optimization frameworks.

Toward better understanding of the effects of system parameters on the energy and stability of biped systems, we presented simulation data for 5 different sets of parameters. For the mass distribution sets we chose, the one which is most similar to human parameters (Set 5) does in fact have the lowest energy consumption. However, the set with the fastest rate of convergence is set 2. These phenomena may in part be a result of other factors, such as choice of feedback control structure and certainly warrant further study.

Also, these and other simulations we have performed show that while increasing the mass of the lower leg (with constant upper-leg mass and total mass) results in a fairly linear increase in COT, as shown in the lower subplot of Figure 10, energy use remains close to flat as mass of the upper leg increases (while holding lower-leg and total mass constant). Corresponding trends relating mass distribution to stability are less apparent and deserve further investigation.
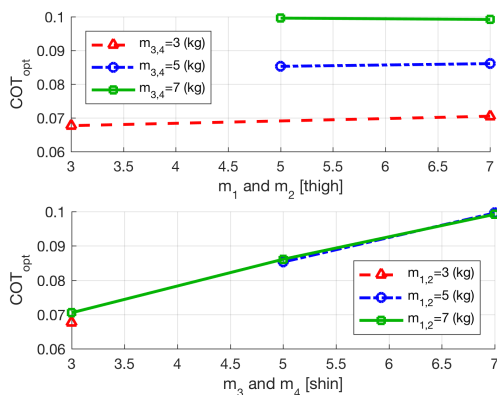

Fig. 10.   COT variations as upper or lower leg mass varies.

To explore trade-offs, we presented COT vs rate of convergence in Figure 9, which illustrates a Pareto frontier, formed essentially by sets 2, 3, and 5, while sets 1 and 4 provide poor trade-off characteristics by comparison. We hypothesize that a more comprehensive framework that includes energy, robustness, and physical parameters in a single optimization can improve the limits of such performance trade-offs, and developing such a framework is a goal for our future work.

REFERENCES

[1] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, "Efficient bipedal robots based on passive-dynamic walkers," *Science*, vol. 307, no. 5712, pp. 1082–1085, 2005.

[2] O. Von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," *Annals of Operations Research*, vol. 37, no. 1, pp. 357–373, 1992.

[3] M. Hardt, K. Kreutz-Delgado, and J. W. Helton, "Optimal biped walking with a complete dynamical model," in *Proc. IEEE Conference on Decision and Control (CDC)*, vol. 3, 1999, pp. 2999–3004.

[4] C. Paul and J. C. Bongard, "The road less travelled: Morphology in the optimization of biped robot locomotion," in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Systems (IROS)*, vol. 1, 2001, pp. 226–232.

[5] E. Westervelt and J. Grizzle, "Design of asymptotically stable walking for a 5-link planar biped walker via optimization," in *Proc. IEEE Int. Conf. on Robotics and Auto. (ICRA)*, vol. 3, 2002, pp. 3117–3122.

[6] D. Djoudi, C. Chevallereau, and Y. Aoustin, "Optimal reference motions for walking of a biped robot," in *Proc. IEEE Int. Conf. on Robotics and Auto. (ICRA)*, 2005, pp. 2002–2007.

[7] H. Dai and R. Tedrake, "Optimizing robust limit cycles for legged locomotion on unknown terrain," in *Proc. IEEE Conference on Decision and Control (CDC)*, 2012, pp. 1207–1213.

[8] C. O. Saglam and K. Byl, "Quantifying the trade-offs between stability versus energy use for underactuated biped walking," in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Systems (IROS)*, 2014, pp. 2550–2557.

[9] K. A. Hamed, B. G. Buss, and J. W. Grizzle, "Exponentially stabilizing continuous-time controllers for periodic orbits of hybrid systems: Application to bipedal locomotion with ground height variations," *The Int. J. of Robotics Research (IJRR)*, vol. 35, no. 8, pp. 977–999, 2016.

[10] A. Hereid, E. A. Cousineau, C. M. Hubicki, and A. D. Ames, "3D dynamic walking with underactuated humanoid robots: A direct collocation framework for optimizing hybrid zero dynamics," in *Proc. IEEE Int. Conf. on Robotics and Auto. (ICRA)*, 2016, pp. 1447–1454.

[11] W. Xi, Y. Yesilevskiy, and C. D. Remy, "Selecting gaits for economical locomotion of legged robots," *The International Journal of Robotics Research (IJRR)*, vol. 35, no. 9, pp. 1140–1154, 2016.

[12] S. Sovero, N. Talele, C. Smith, N. Cox, T. Swift, and K. Byl, "Initial data and theory for a high specific-power ankle exoskeleton device," in *Proc. Int. Symp. on Experimental Robotics (ISER)*, 2016, pp. 355–364.

[13] G. Bellegarda, N. Talele, and K. Byl, "Exploring nonintuitive optima for dynamic locomotion," 2017 (submitted).

[14] Y. Hurmuzlu and D. B. Marghitu, "Rigid body collisions of planar kinematic chains with multiple contact points," *The International Journal of Robotics Research (IJRR)*, vol. 13, no. 1, pp. 82–92, 1994.

[15] M. Posa and R. Tedrake, "Direct trajectory optimization of rigid body dynamical systems through contact," in *Algorithmic Foundations of Robotics X*.   Springer, 2013, pp. 527–542.

[16] J. Andersson, "A General-Purpose Software Framework for Dynamic Optimization," PhD thesis, KU Leuven, 2013.

[17] L. T. Biegler and V. M. Zavala, "Large-scale nonlinear programming using ipopt: An integrating framework for enterprise-wide dynamic optimization," *Comp. & Chem. E.*, vol. 33, no. 3, pp. 575–582, 2009.

[18] H. Mittleman. (2017) Decision tree for optimization software. (Benchmarking for SQP, and other optimization problems). [Online]. Available: http://plato.asu.edu/bench.html

[19] M. Srinivasan, *Why walk and run: energetic costs and energetic optimality in simple mechanics-based models of a bipedal animal*. Cornell University Ithaca, NY, 2006.

[20] A. D. Kuo, "Energetics of actively powered locomotion using the simplest walking model," *Journal of biomechanical engineering*, vol. 124, no. 1, pp. 113–120, 2002.