Do Limit Cycles Matter in the Long Run? Stable Orbits and Sliding-Mass Dynamics Emerge in Task-Optimal Locomotion

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

***Abstract*— We investigate the task-optimality of legged limit cycles and present numerical evidence supporting a simple general locomotion-planning template. Limit cycles have been foundational to the control and analysis of legged systems, but as robots move toward completing real-world tasks, are limit cycles practical in the long run? We address this question both figuratively and literally by solving for optimal strategies for long-horizon tasks spanning as many as 20 running steps. These scenarios were designed to embody practical locomotion tasks, such as evading a pursuer, and were formulated with minimal constraints (complete the task, minimize energy cost, and don't fall). By leveraging large-scale constrained optimization techniques, we numerically solve the trajectory for a reduced-order running model to optimally complete each scenario. We find, in the tested scenarios in flat terrain, that near-limit-cycle behaviors emerge after a transient period of acceleration and deceleration, suggesting limit cycles may be a useful, near-optimal planning target. On rough terrain, enforcing a limit cycle on every step only degrades gait economy by 2-5% compared to optimal 20-step look-ahead planning. When perturbing the scenario with a single "bump" in the road, the model converged in a manner giving the appearance of an asymptotically stable orbit, despite not explicitly enforcing asymptotic stability. Further, we show that the transient periods of acceleration and deceleration may be near-optimally approximated by planning with a simple "sliding mass" template. These results support the notion that limit cycles can be useful approximations of task-optimal behavior, and thus are useful near-term targets for long-term planning.**

I. INTRODUCTION

Applied robotics is a task-driven enterprise, but our best mathematical formulations rarely conform precisely to the task at hand. So by and large, robot control formulations are developed in constant compromise between what is task-relevant and what is mathematically tractable. Legged locomotion is no exception.

For highly dynamic locomotion in particular, limit-cycle locomotion is a prominent task-simplifying framework. Enforcing limit-cycle stability reduces an otherwise large-scale planning operation down to a state-regulation problem; far friendlier territory for modern control theory. Further, legged animals in hosts of experiments appear to exhibit and regulate limit cycles, thus teasing a biological basis for this control simplification.

But what is the price of this mathematical convenience in practical tasks? Particularly, in a robotics field struggling to mitigate the poor energy economy of its machines [1], perhaps non-limit-cycle strategies could yield meaningful efficiency gains. In short, if we restrict our control to the limit-cycle framework, would we miss the solutions being cast aside?

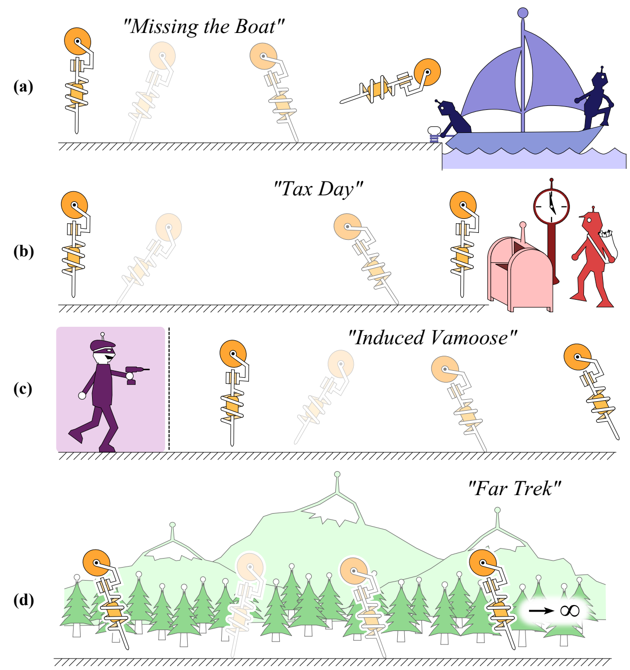

Fig. 1. Scenarios used to investigate properties of task-optimal locomotion strategies. **a)** *"Missing the Boat"* requires crossing a finish line within a specified time limit after starting from a rest position. **b)** *"Tax Day"* specifies the robot cross a finish line within a time limit and must both start and end in a specified rest posture. **c)** *"Induced Vamoose"* mandates a minimum distance from a constant speed pursuer for a specified time frame, finishing in an upright posture. **d)** *"Far Trek"* demands locomotion exceeding a specified average speed for a 20-step maneuver while satisfying equivalence in the boundary conditions, approximating infinite-horizon locomotion planning.

We specifically address the following questions about the role of steady (limit-cycle) locomotion in the context of longer-time-horizon tasks:

Question 1. Does limit-cycle locomotion emerge as an optimal strategy for practical locomotion tasks?

Question 2. What is the additional energy cost of enforcing a limit cycle in rough terrain when compared to best-possible step planning?

Question 3. Does task-optimization yield a stable-appearing limit cycle upon task perturbation?

Question 4. Is there a heuristic for planning transient non-

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: hubickic@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

steady maneuvers?

We investigate these questions by numerically searching for energy-optimal strategies to simple and practical task scenarios using trajectory optimization techniques. These scenarios are designed such that transient behaviors are required, such as starting and stopping (Fig. 1), in order to see if limit-cycle locomotion emerges in the long run, despite not being explicitly demanded. We also further complicate scenarios with rough terrain (both a single mid-course bump and randomly generated rough terrain), in order to compare optimal performance to control that directly targets a limit cycle.

In multiple tasks, the optimized strategies often approached, or exactly rendered, a limit cycle (often after a transient acceleration period). On rough terrain, we found that directly enforcing a limit cycle had little energetic consequence when compared to an optimally-planned sequence of over 20 steps. Further, when a single bump in the road was placed, the resulting planned trajectory approximated that of an asymptotically stable limit cycle. Lastly, while inspecting for the emergence of limit cycles, we found that the transient behaviors in task-optimal motions matched the dynamics of a sliding mass. We suggest that planning whole robot motions, using a simple sliding mass as a template, could provide computationally simple and near-optimal locomotion task planning.

II. BACKGROUND

Task-optimal locomotion planning has both engineering and scientific applications. However, it represents a significant ongoing computational challenge; thus, it continues to be approached in a number of ways, with varying degrees of task-optimality and dynamical agility. Broadly speaking, long-term planning is achieved by either 1) constraining the scope of the planning to achieving task-feasibility [2], 2) simplifying robot dynamics to a template dynamical model [3], or 3) approximating the task with shorter-horizons [4], *e.g.*, limit cycles.

Finding an optimal, or at least efficient, task solution is notably more computationally demanding than achieving a feasible solution, which may be all that is required for many tasks. Zero-moment point methods [5], perturbation theory [6], low-order action spaces [7], and feedback linearization [8] have all been implemented to develop real-time task-feasible trajectories. Techniques that plan using a reduced-order robot model typically employ either the linear-inverted pendulum (LIP) [9] or the spring-loaded inverted pendulum (SLIP) models [8].

Operating on short-term targets, limit cycles have long been a focus of legged locomotion dynamics [10] as they can be used as a straightforward state target for sustained locomotion. Robots that successfully stabilize a limit cycle can theoretically move indefinitely [11], and many machines can even demonstrate this stability passively [12], albeit for small state regions. Techniques for regulating limit cycles are myriad [13], and have been used to efficiently walk [14] and run [15] and recover successfully in the face of significant terrain perturbations. Further, humans exhibit an apparent limit-cycle stability as well [16], teasing some biological relevance to the approach. Model-predictive control techniques have been able to plan with increasingly distant time horizons [17], but it's still not clear what a useful near-term state target would be on for longer time-horizon tasks.

Task-level control optimization can also shed light on biological debates. Several overarching hypotheses have been posited as frameworks for animal locomotion [18]. The apparent stability of limit-cycle locomotion in animals can be viewed as a self-stable mechanical phenomenon with largely feed-forward control stabilization. Additionally, the cyclical motions of biology could be a product of their periodic neuromechanical architecture [19]. However, an alternative view predicts that features of animal locomotion emerge from the completion of practical tasks [20]. Effective task-level optimizations and insights from their solutions would bear directly on this "task-optimality" hypothesis.

In this investigation, we seek insights from long-horizon strategies for legged locomotion tasks that are as close to optimal as possible, evaluating the relative optimality of limit cycle targets. Therefore, our optimization methods can eschew the need for any real-time trajectory generation [17], and use slower but more precise constrained optimization techniques. In this way, our approach mirrors that of prior optimization investigations seeking insight into locomotion strategies [21], [22].

III. PROBLEM SETUP

Our problem-solving methodology consists of a simple hopping model, a series of defined task scenarios, and a trajectory optimization procedure to find optimal solutions.

A. Model

For our plant model, we employed an actuated and dissipative variant of the Spring-Loaded Inverted Pendulum (SLIP), shown in Fig. 2, which despite being a simple model, still captures key mechanical properties of legged systems. In terms of mathematical properties, the simple model resembles more complex robots in that it is nonlinear, hybrid dynamical, underactuated, and has no exact closed-form solution. However, the reduced number of dynamical states and actuators make longer-horizon trajectory optimization of this model somewhat tractable.

Importantly, unlike the typical SLIP model, this model has inherent dissipative losses, a requirement for meaningful implementation of energy-optimal control. This model variant, when asked to run steadily with minimal actuator work, also generates dynamics and ground-reaction forces that are highly similar to running animals, including the asymmetry in ground-reaction forces [23].

We focused our analysis on running gaits for both computational and scientific reasons. In terms of computation, running can be formulated by a one-legged model, resulting in fewer design variables to optimize. Also, since the aerial running phase can be solved in closed-form, the "flight-to-stance" sequencing allows for more steps to be encoded

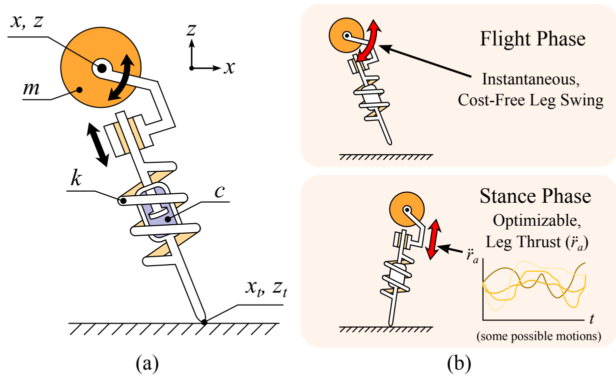

Fig. 2. a) The actuated spring-mass model tasked with optimally completing the specified scenarios and b) A diagram of the actuation schemes in flight and stance.

with fewer variables than walking (and much fewer than multipedal gaits). Further, in the scientific literature, reduced-order modeling of biology is less controversial in running as even unsprung running models will exhibit spring-mass dynamics [21], while models for bipedal walking are further from consensus [24], [25].

As per standard practice, the dynamics of the model were split into stance and flight phases, respectively defining the dynamics when the leg is on and off the ground. The stance equations of motion are given as:

$$\ddot{x} = \frac{F(x - x_t)}{mr}$$
$$\ddot{z} = \frac{F(z - z_t)}{mr} - g$$

where x and z represent the Cartesian position of the point mass, F is the axial leg force, m is the robot point mass, g is gravitational acceleration, r is the leg length (the Euclidean distance between the toe and the mass position), while x_t and z_t encode the stationary toe position. The force is produced by a massless actuated spring-damper leg, where:

$$F = k(r_a - r) + c(\dot{r}_a - \dot{r})$$

where r_a is the rest length of the leg as defined by the actuator position, and k and c are stiffness and damping coefficients respectively. Actuation is driven by the acceleration of the leg length, $\ddot{r}_a$, and represents the control input during stance. Non-zero compression of the stance leg at the instant of touchdown is prohibited (*i.e.* $r \neq r_a$). All parameter quantities ($m = 1$, $k = 20$, $c = 0.5$, $g = 1$) are normalized to the Froude number ($v/\sqrt{gL_{norm}}$) and are thereby dimensionless ($L_{norm} = 1$). The stiffnesses are chosen to be similar to reported human values [26] with dissipation values similar to other running biped studies [23].

Flight dynamics initiate when F intersects zero and are defined as a purely ballistic trajectory. The swing of the massless leg in flight is modeled such that it can move instantaneously, achieving desired touchdown conditions within set functional limits of the model[1]. We justify this assumption as it accounts for the fact that a bipedal robot has a second leg to swing, positioning itself in anticipation of the next touchdown (and it is a common assumption in other SLIP studies).

B. Task Scenarios

We formulated real-world-inspired tasks which represent practical locomotion functions while minimally constraining possible strategies. In each of these scenarios, we used the loosest "stability" constraint possible for this legged model, only prohibiting falls (*i.e.*, the center of mass (CoM) must always be above the terrain). This freedom permits the optimizer to explore all manner of counter-intuitive strategies, such as moving backward.

All of our specified tasks demand that energy costs are minimized as the sole term of the objective function, and for three key reasons. First, minimizing transport costs is an important goal for robots toward achieving functional autonomy [1]. There is also copious evidence from studies of models [21] and experiments [23] that minimizing energy costs is a high priority for legged locomotion in animals. Most critically, adding multiple objectives to the problem muddies the water in terms of interpreting the resulting solutions since this will require the experimenter to decide (somewhat arbitrarily) how to weight each objective. The energy-cost objective is defined as the specific mechanical cost of transport:

$$COT = \frac{\sum_{i=1}^{N}\left(\int_0^{t_s} |F\dot{r}_a|\, dt\right)_i}{mgd}$$

where N is the total number of stance phases and t_s is the stance duration. Due to the inherent nonsmoothness resulting from the absolute value, a smooth approximation was used to facilitate more reliable solving of the nonlinear program (NLP).

$$|x| \approx \sqrt{x^2 + \epsilon^2} - \epsilon$$

where ϵ is small (0.001).

Again, these problems are always formulated with a single objective, and therefore, have no weighting coefficients that can be tweaked which shape the solution. All task specifications other than minimizing energy costs, such as avoiding falls and passing finish lines, are enforced as hard equality constraints or hard inequality constraints as appropriate, and must be satisfied within defined tolerances (10^{-6}).

1) Missing the Boat: Getting from A to B is a ubiquitous colloquial definition of locomotion. The "Missing the Boat" task demands that the hopper jump on the boat before its colleagues cast off. Operationally speaking, the task requires surpassing a point a predefined distance away in less time than a predefined time limit. The hopper starts at zero

[1]Functional limits for the model are set as follows: $-1 \leq \ddot{r}_a \leq 1$; $0.5 \leq r \leq 1$; and $0 \leq F \leq 3$.

velocity and is dropped from a $1.05L_{norm}$ height, but has no additional constraints on its terminal state. With no posture-related terminal state constraints, the hopper is likely to plan a less-than-graceful arrival (Fig. 1a).

2) Tax Day: More often than not, arriving at a desired locomotion is a necessary, yet insufficient, condition for accomplishing a locomotion task. The "Tax Day" scenario requires the hopper arrive *upright* at a mailbox to mail a letter before the postbot arrives, thereby meeting an unspecified important filing deadline (Fig. 1b). This task is identical to the prior "boat" task except that it places a constraint on the final state. The hopper must finish its maneuver in a final apex condition (velocity and height) equivalent to the initial apex condition (zero velocity and $1.05L_{norm}$ high).

3) Induced Vamoose: Sometimes, locomotion is less about arriving at a location than it is leaving one; and perhaps with haste. In the "Induced Vamoose" task, the hopper must maintain a safe distance from an approaching adversary: an unsavory robotic assailant brandishing an unsavory robotic weapon (Fig. 1c). For a specified time duration, the hopper's CoM must stay out of a danger zone that sweeps forward in time with constant velocity. Further, the robot's CoM must also be above a defined vertical threshold at the end of its sprint ($1.05L_{norm}$), as not to lie down and be easy prey for the merciless attacker.

4) Far Trek: Finally, when a task destination is sufficiently far away, it is generally assumed to be at an infinite time horizon. In the "Far Trek" scenario, the hopper decides to take an infinite-time horizon voyage, efficiently going where no underactuated spring-mass-damper hopping model has previously gone (Fig. 1d). We specify the hopper to move at a required average velocity over the course of a predefined number of hopping steps, simulating a desired arrival time at a distant location. To approximate the infinite horizon, we enforce a cyclic constraint at the boundary conditions of the multi-step maneuver, requiring equivalence in apex height and velocity.

C. Terrain Perturbations

This investigation also seeks insight into how terrain affects task-optimal locomotion. For the "Far Trek" scenario, we perturbed the robot model with two terrain conditions to test the effect on two quantities of interest: apparent trajectory stability and energy economy. To probe for apparent stability of these task-optimal trajectories, we confronted the robot model with a single obstacle in the middle of the task[2] (step 11 of 21). An obstacle height of 0.1 L_{norm} was selected as it was big enough to potentially change the optimal strategy, but small enough to be completely rejected by crouching if deemed desirable by the optimizer (results depicted in Fig. 5a).

We also investigated the energy costs of enforcing limit cycles by subjecting the hopper to "rough terrain," in order to see if the additional computation of planning many steps ahead is worthwhile energetically. To model rough terrain for completing the "Far Trek" task, we modified the scenario such that the terrain height of each step varies in a uniform distribution between 0-0.05 L_{norm}. Further, we iterated the optimization to command a range of average speeds (v_{avg} from 0.5 to $3\sqrt{L_{norm}g}$) in order to compare transport costs across speeds (results depicted in Fig. 5b). In all optimizations, the optimizer has full awareness of upcoming terrain, and thus has the ability to plan accordingly.

D. Optimization

While a 20-step journey may not intuitively sound like a "long run," finding an optimal strategy for a long-time-horizon problem can be a challenging computational task, ultimately requiring the optimization of many thousands of variables and constraints. To accomplish multi-step planning, we formulated a large-scale directly collocated nonlinear program and solved it using large-scale constrained optimization techniques. Constrained optimization techniques like Sequential Quadratic Programming (SQP) are well suited to task optimizations since task constraints can be formulated neatly, independently, and can be strictly enforced. To first formulate the problem to be solved, we used a optimization-parsing package for MATLAB we title: **C**ontrol **O**ptimization **A**ccelerated by the **L**arge-scale **E**xport of **S**ymbolic **C**ollocated **E**lements (COALESCE). COALESCE builds a fast and smooth optimization problem using a direct collocation approach [27], which treats system dynamics as constraints to be enforced directly by the optimizer, not by a time-marching integrator [28]. After symbolically generating an objective, constraints and a sparse Jacobian, we solved the problem with SNOPT, an SQP implementation designed for large-scale constrained optimizations [29].

In formulating the direct collocation constraints, we used an implicit Euler integration scheme to accelerate solving. Switching to Hermite-Simpson integration, a more accurate method, required significantly more computation time (approximately 10 minutes instead of 20 seconds), but did not yield any significant differences in optimal strategy. Presented solutions discretized the problem to 50 nodes per stance phase, with flight dynamics enforced as jump maps since they are ballistic and thus solvable in closed form. The phase ordering was scheduled *a priori*, chaining stance and flight phases in succession. For the "Missing the Boat" and "Tax Day" tasks, the optimal number of steps was unknown, so individual solutions were computed with between 4-21 steps and the lowest cost solution was chosen *post hoc*.

This optimization procedure was very reliable in finding and certifying locally optimal results. In terms of the "locality" of the optima, solutions were routinely found with arbitrary initial guesses (all zeros or ones). Different initial guesses were spot checked, not revealing notably different solutions from previously certified solutions (small differences would vanish upon tightening optimality tolerances, which are ordinarily 10^{-6}). Further, the optimization was benchmarked on problems with known optimal solutions, such as a minimum work problem with no dissipation nor energy-changing task requirements.

[2] The optimizer was able to plan for this obstacle, so it could be considered a perturbation to the task, not the gait.

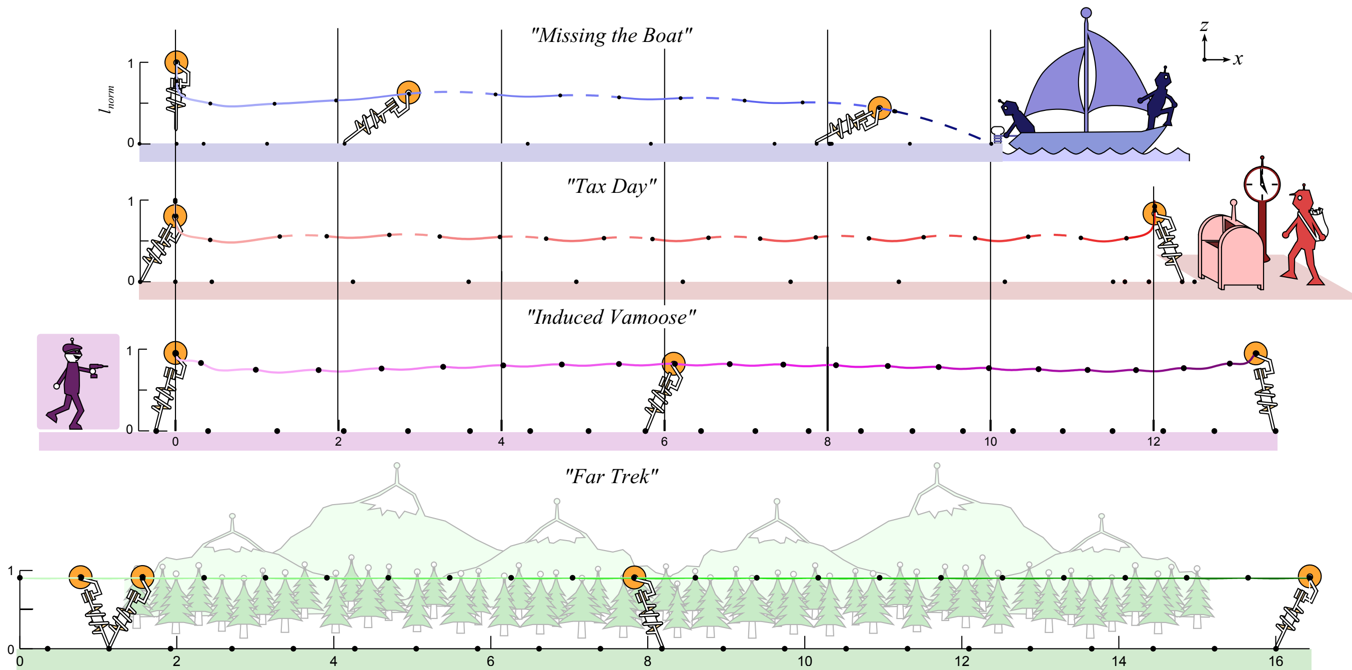

Fig. 3. Optimal trajectories plotted in Cartesian space for each task scenario, designed to give an intuitive impression of the optimized strategies. Solid and dashed lines signify stance and flight phases respectively and black dots mark the position of the toe and CoM at the beginning and end of each stance phase. The opacity of the plotted lines increases with time.

IV. NUMERICAL RESULTS

We generated optimal solutions for each task scenario[3] (Fig. 3), inspecting the state trajectories for limit cycles and other general patterns. In general, each of the solutions appears to take all obvious measures toward minimizing energy costs (e.g., landing on the boat, head-first, right at the finish line as not to waste any gravitational potential energy), suggesting the optimizer is successfully finding optimal solutions to the problem.

Fig. 4 visualizes the phase plots of the optimized maneuvers so that limit cycles can be seen directly (as overlapping orbits). Each of the first three tasks appear to accelerate and asymptotically approach a steady gait before performing stopping maneuvers (Fig. 4a-c). In the case of "Tax Day," the robot appears to be nearly maximally accelerating and decelerating (i.e., bang-bang) throughout the task, so it likely furthest from achieving a repeated cycle. After the first 3-4 steps, the "Missing the Boat" and "Induced Vamoose" tasks appear to hover in the vicinity of a limit cycle before rapidly braking (or crashing) at the end of the task.

Most notably, the "Far Trek" task was very steady to the point of near-numerical accuracy (Fig. 4d), in spite of having the freedom to choose any number of non-steady solutions. This suggests that once a satisfactory limit cycle gait is found, there is likely little utility adjusting the gait unless speed changes are desired. Put another way, it suggests that a unique optimal gait corresponds to each desired speed.

In the single-bump terrain perturbation, the hopper chose to vault slightly higher onto the obstacle and quickly (and asymptotically) return to a limit cycle (Fig. 5a). In the "rough terrain" experiment, the cost of transport for an enforced limit cycle was only slightly more expensive (2-5%) than the optimal planned strategy, across all tested speeds.

[3]Specific task parameters: *"Missing the Boat"* task finish line, $d_{final} = 10L_{norm}$, and time limit, $t_{lim} = 10\sqrt{g/L_{norm}}$; *"Tax Day,"* $d_{final} = 10L_{norm}$, $t_{lim} = 10\sqrt{g/L_{norm}}$; *"Induced Vamoose"* adversary velocity, $v_{adv} = 1\sqrt{L_{norm}g}$, number of steps, $N_{steps} = 21$; *"Far Trek"* mean velocity, $v_{avg} = 2\sqrt{L_{norm}g}$, $N_{steps} = 21$.

V. DISCUSSION

Question 1: Does limit-cycle locomotion emerge as an optimal strategy for practical locomotion tasks? In short, yes, to varying degrees. As can be seen in numerous phase plots, many orbits emerge that trace over each other very closely. This most pronounced in Figure 4d, the "Far Trek," where the limit cycle is nearly exact. In the other scenarios, while a clear drift a the limit cycle can be seen in this modeling experiment, such a subtle change might be dismissed as noise in robot measurements (or in animal experiments).

What is the additional energy cost of enforcing a limit cycle in rough terrain when compared to best-possible step planning? With randomly bumpy terrain, the additional cost of enforcing a limit cycle was quite small (Figure 5b), inflating the cost of transport just 2-5%. This suggests that the computational challenge of carefully optimizing a strategy for every step of upcoming rough terrain may not yield much energetic benefit. Previous metastability analyses suggested a similar conclusion/recommendation, but primarily for the purpose avoiding falls [7]. As such, assuming a

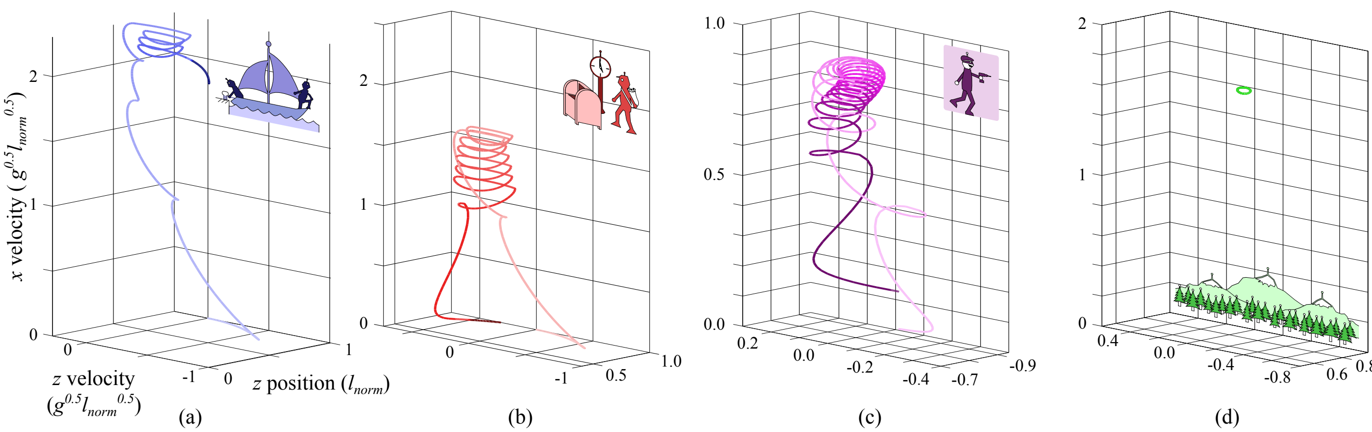

Fig. 4. Optimal trajectories plotted in phase space for each scenario (horizontal velocity, vertical velocity, and vertical position). Such phase plotting lends to the easy visualization of potential limit cycles as overlapping orbits, with time increasing with the opacity of the plotted trajectory. a) The "Missing the Boat" strategy shows near-limit-cycle behavior at high speed, just before landing on the boat (head first). b) "Tax Day" behavior accelerates steadily, approaches a potential limit cycle, and decelerates quickly to a stop. c) "Induced Vamoose" behavior spends most of its steps near a limit cycle before braking to a stop. d) The "Far Trek" solution was a near-numerically-perfect limit cycle for all 21 steps.

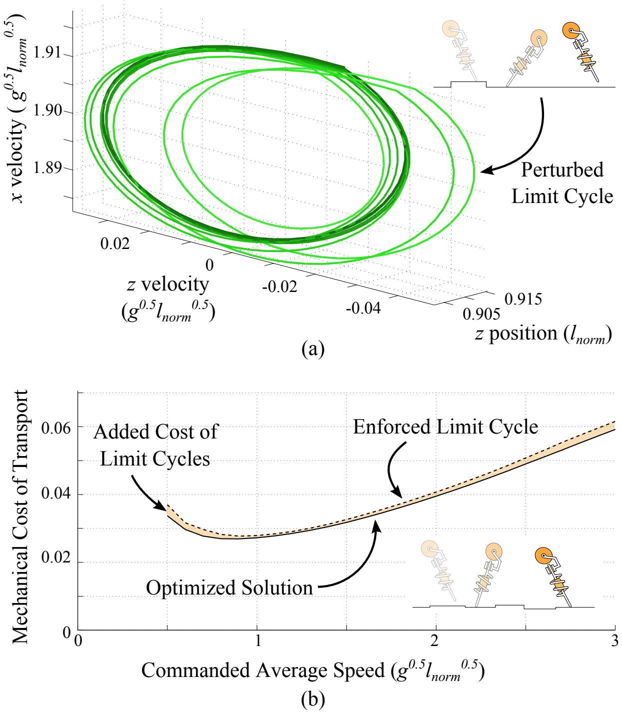

Fig. 5. The effect of planning with terrain perturbations on optimal strategies. a) A phase plot (horizontal velocity, vertical velocity, and vertical position) of the result of a single obstacle "bump" placed into the middle of the Far Trek scenario. Beginning with the on-obstacle step (lightest line shading), the planned trajectory quickly and asymptotically returns to a nominal limit cycle as an emergent behavior. b) A specific mechanical cost-of-transport (COT) plot (Energy cost/distance/weight) for the Far Trek scenario with randomly generated bumpy terrain. This plot compares two strategies, one optimized with full knowledge of every forthcoming obstacle over 21 steps, and the other constrained to a limit cycle for those 21 steps (defined by apex conditions). The cost of transport across many different commanded average speeds shows only a small energetic benefit (2-5%) in meticulously planning a strategy to optimally negotiate a sequence of small obstacles.

limit cycle in the presence of "randomly" rough terrain may be adequate for practical energetic purposes.

Question 3: Does task-optimization yield a stable-appearing limit cycle upon task perturbation? In Figure 5a, the evolution of the terrain-perturbed trajectory looks like the convergence of an asymptotically stable limit cycle. This is notable since this "apparent asymptotic stability" is emergent, as the optimization was not explicitly tasked with finding asymptotic convergence to limit cycles. This behavior was produced by long-time-horizon task constraints and energy-cost minimization. This suggests that encoding asymptotic stability is not the only approach for achieving control that exhibits stable-looking behaviors.

Question 4: Is there a heuristic for planning transient non-steady maneuvers? When inspecting many of these phase plots, one can see indications of "bang-bang"-like accelerations in tasks requiring fast movement (such as in "Tax Day"). Bang-bang is most iconically attributed to minimum-time control of a sliding mass. As such, we attempted to reproduce the velocity profiles exhibited by our hopper by giving the same task constraints to a simple "sliding mass" subject to an external force. The dynamics for this simple system are defined as:

$$\ddot{x} = \frac{F(t) - c\dot{x}}{m}$$

where x is the horizontal position, $F(t)$ is a bounded time-varying control input, m is the block's mass, and c is a viscous friction coefficient. After some manual iteration of the sliding mass model parameters (upper and lower force limits, and linear viscous friction), we produced the velocity profile for the sliding mass completing the "Tax Day" scenario (Fig. 6).

The resulting velocity profile bears a striking resemblance to our optimized hopper strategy. This makes a good deal of intuitive sense. If our optimizer is successful at mini-

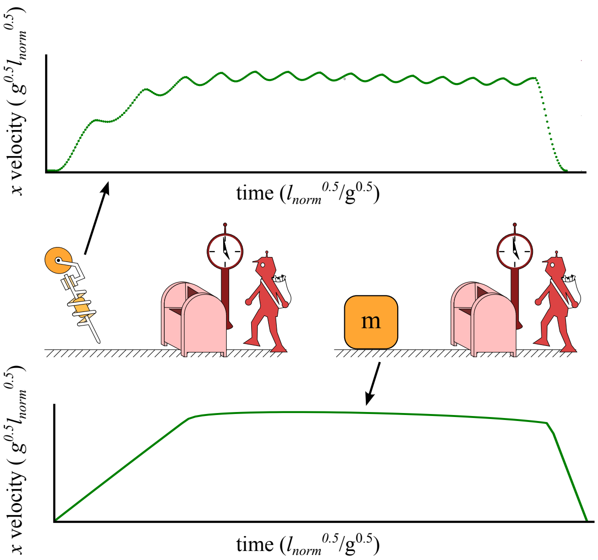

Fig. 6. The optimal velocity profile of the hopper performing the "Tax Day" scenario (top) compared to a sliding mass (with viscous friction subject to a time varying horizontal force) completing the same task (bottom). This close similarity implies that assuming a simple sliding mass may be a "good-enough" approximation for scheduling velocities for optimal locomotion planning.

mizing energy costs, it should be stripping away extraneous costs and reveal the basic physics of moving a mass from A to B. What is not necessarily obvious is that this clear pattern would emerge in spite the nonlinearity and underactuated/hybrid-dynamical mode of propulsion present in our hopping model. We posit that this property may be true of even more complex robots, which may take advantage of fundamental sliding-mass physics for the purpose of long-horizon planning.

A. Application to More Complex Systems

The employed spring-mass-damped hopping model is quite simple. It does not include many potential sources of energy loss which would be seen on practical machines, including rigid-body impacts, swing-leg costs, thermal dissipation, transmission losses and torso-stabilizing costs. However, we argue that the energetics of locomotion transport economy are likely similar across legged machines and cost models. Numerical gait generation studies have suggested that optimal trajectories are somewhat invariant to the details of the energy-cost objective [30]. In more complicated legged systems where transport costs have been rigorously analyzed across speed, the same smooth concave-up speed-cost relationship emerges [31] matching what is shown in Figure 5b. This general relationship also appears to be true of animals, from walking humans [32], to running ostriches [33], to trotting horses [34]. We suspect that the energetics of transport could be mapped to a sliding-mass template by tuning the sliding-mass dissipation as we did to produce Fig. 6.

VI. CONCLUSIONS

We investigated the role of limit cycles and their apparent stability in optimally accomplishing legged locomotion tasks. While the formulated tasks were presented in a somewhat tongue-in-cheek style, they were designed to address serious questions about the role of limit-cycle targets in locomotion planning. We note that in several tasks, near-limit cycle behavior resulted for a portion of the scenario, if not its entirety. Further, adding obstacles as perturbations to the scenario revealed an apparent asymptotic convergence to a limit cycle, despite not directly encoding asymptotic stability into the optimization. These observations suggest an inherent utility to limit cycles in legged locomotion. Further, the apparent asymptotic stability of the limit cycle suggests that such stable-looking orbits can emerge from energy minimization and task constraints, not necessarily any explicit stabilization of limit cycles.

We also suggest that "sliding mass" dynamics appear to approximate observed optimal strategies, despite being among the simplest possible systems to optimize. This suggests that long-term planning motions for analytically complex legged systems may be meaningfully reduced to short-term planning operations. We posit that this template can be used to rapidly compute near-energy-optimal state planning for arbitrarily long locomotion tasks.

Lastly, we make a general note regarding the dynamics of task-optimal legged locomotion. While limit cycles and asymptotic limit-cycle stability are handy and efficient targets of control design, our investigation observes these properties as purely emergent phenomena. This observation contributes to a utilitarian explanation for the ubiquity of periodic locomotion in nature, suggesting that it likely best serves the task at hand. Pressures for task-optimality likely mold both engineering and biology; be it for robot control or animal survivability, in the long run, it's the task that matters.

REFERENCES

[1] S. H. 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] M. Zucker and J. Bagnell, "An optimization approach to rough terrain locomotion," *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 3589–3595, 2010.
[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] T. Erez, Y. Tassa, and E. Todorov, "Infinite-horizon model predictive control for periodic tasks with contacts," in *Robotics: Science and Systems VII*, 2012.
[5] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, H. Hirukawa, and M. Morisawa, "Biped walking pattern generation by using preview control of zero-moment point," in *ICRA'03*, vol. 2, (Taipei, Taiwan), pp. 1620–1626, 2003.
[6] L. Sentis and B. Fernandez, "Perturbation theory to plan dynamic locomotion in very rough terrains," *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2267–2273, Sept. 2011.
[7] K. Byl and R. Tedrake, "Metastable Walking Machines," *International Journal of Robotics Research*, vol. 28, no. 8, pp. 1040–1064, 2009.
[8] G. Piovan and K. Byl, "Two-element control for the active SLIP model," *2013 IEEE International Conference on Robotics and Automation*, pp. 5656–5662, May 2013.

[9] E. C. Whitman and C. G. Atkeson, "Control of Instantaneously Coupled Systems applied to humanoid walking," *2010 10th IEEE-RAS International Conference on Humanoid Robots*, pp. 210–217, Dec. 2010.

[10] Y. Hurmuzlu and G. D. Moskowitz, "The role of impact in the stability of bipedal locomotion," *Dynamics and Stability of Systems*, vol. 1, no. 3, pp. 217–234, 1986.

[11] P. A. Bhounsule, J. Cortell, and A. Ruina, "Design and control of Ranger: an energy-efficient, dynamic walking robot," in *Proceedings of the Fifteenth International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines (CLAWAR)*, pp. 441–448, 2012.

[12] T. McGeer, "Passive dynamic walking," *International Journal of Robotics Research*, vol. 9, no. 2, pp. 62–82, 1990.

[13] 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.

[14] 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.

[15] 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.

[16] J. B. Dingwell, J. P. Cusumano, D. Sternad, and P. R. Cavanagh, "Slower speeds in patients with diabetic neuropathy lead to improved local dynamic stability of continuous overground walking.," *Journal of biomechanics*, vol. 33, pp. 1269–77, Oct. 2000.

[17] 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.

[18] 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.

[19] A. J. Ijspeert, "Central pattern generators for locomotion control in animals and robots: a review.," *Neural Networks*, vol. 21, no. 4, pp. 642–653, 2008.

[20] J. Lee, A. Lamperski, J. Schmitt, and N. Cowan, "Task-level control of the lateral leg spring model of cockroach locomotion," *Fast Motions in Biomechanics and Robotics*, pp. 167–188, 2006.

[21] M. Srinivasan and A. Ruina, "Computer optimization of a minimal biped model discovers walking and running," *Nature*, vol. 439, pp. 72–75, Jan. 2006.

[22] J. D. Karssen, M. Haberland, M. Wisse, and S. Kim, "The optimal swing-leg retraction rate for running," *2011 IEEE International Conference on Robotics and Automation*, pp. 4000–4006, May 2011.

[23] A. Birn-Jeffery, C. Hubicki, Y. Blum, D. Renjewski, J. Hurst, and M. Daley, "Don't break a leg: Running birds prioritise safety and economy, not strict stability," *Journal of Experimental Biology, under review*, 2014.

[24] H. Geyer, A. Seyfarth, and R. Blickhan, "Compliant leg behaviour explains basic dynamics of walking and running," *Proc. R. Soc. Lond. B.*, vol. 273, pp. 2861–2867, Nov. 2006.

[25] M. Srinivasan and P. Holmes, "How well can spring-mass-like telescoping leg models fit multi-pedal sagittal-plane locomotion data?," *Journal of theoretical biology*, vol. 255, pp. 1–7, Nov. 2008.

[26] S. W. Lipfert, G. Michael, D. Renjewski, S. Grimmer, and A. Seyfarth, "A model-experiment comparison of system dynamics for human walking and running," *Journal of Theoretical Biology*, vol. 292, pp. 11–17, 2012.

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

[28] A. V. Rao, "A survey of numerical methods for optimal control," *Advances in the Astronautical Sciences*, vol. 135, no. 1, pp. 497–528, 2009.

[29] 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.

[30] 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.

[31] C. Remy, K. Buffinton, and R. Siegwart, "Comparison of cost functions for electrically driven running robots," in *IEEE/RSJ International Conference on Robotics and Automation (ICRA)*, pp. 2343–2350, Ieee, May 2012.

[32] A. D. Kuo, J. M. Donelan, and A. Ruina, "Energetic consequences of walking like an inverted pendulum: step-to-step transitions.," *Exercise and sport sciences reviews*, vol. 33, pp. 88–97, Apr. 2005.

[33] J. Rubenson, D. B. Heliams, D. G. Lloyd, and P. A. Fournier, "Gait selection in the ostrich: mechanical and metabolic characteristics of walking and running with and without an aerial phase," *Proc. R. Soc. Lond. B*, vol. 271, no. 1543, pp. 1091–1099, 2004.

[34] C. T. Farley and C. R. Taylor, "A mechanical trigger for the trot-gallop transition in horses," *Science*, vol. 253, pp. 306–308, 1991.