A Self-Exciting Controller for High-Speed Vertical Running

Goran A. Lynch†, Jonathan E. Clark‡, Daniel Koditschek†

†Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, USA

‡Department of Mechanical Engineering, FAMU & FSU College of Engineering, Tallahassee, FL, USA

{goran,kod}@seas.upenn.edu, clarkj@eng.fsu.edu

Abstract— **Traditional legged runners and climbers have relied heavily on gait generators in the form of internal clocks or reference trajectories. In contrast, here we present physical experiments with a fast, dynamical, vertical wall climbing robot accompanying a stability proof for the controller that generates it without any need for an additional internal clock or reference signal. Specifically, we show that this "self-exciting" controller does indeed generate an "almost" globally asymptotically stable limit cycle: the attractor basin is as large as topologically possible and includes all the state space excluding a set with empty interior. We offer an empirical comparison of the resulting climbing behavior to that achieved by a more conventional clock-generated gait trajectory tracker. The new, self-exciting gait generator exhibits a marked improvement in vertical climbing speed, in fact setting a new benchmark in dynamic climbing by achieving a vertical speed of 1.5 body lengths per second.**

I. Introduction

Two interconnected problems arise in the design and control of legged climbing robots: attachment and force production [1]. Attaching to the wall requires a mechanism which allows the robot both to cling to a vertical surface (while producing forces tangent to that surface) and then to release its grip rapidly and smoothly in order to re-circulate for the next touchdown. Substantial progress has been made toward developing attachment mechanisms for a range of substrates. Specialized approaches using magnets or a vacuum [2], [3], [4], have recently given way to more general purpose mechanisms such as microspines [5], [6], dry adhesives [7], [8], dactyls [6], and bracing/friction [9] to provide robust, reliable attachment to an ever growing diversity of vertical surfaces when driven appropriately by a well tuned, algorithmically controlled set of actuators [10]. The second challenge, given a functional attachment mechanism, is to achieve a force pattern which results in desirable climbing behavior. Here again, growing empirical evidence indicates that the generation of robust and efficient force patterns requires both passive mechanical and active algorithmic design. Just as legged walking and running machines require a combination of tuned passive mechanisms and algorithmically controlled actuators to locomote quickly and efficiently [11], [12], the first legged machines to achieve mobility on unstructured outdoor walls and trees rely on a combination of passive mechanical and actively powered sources of force patterns [1]. This paper addresses the second aspect of the second problem: the development of algorithms for actively powering force patterns in dynamical climbing.

The past few years have seen the development of a small number of climbing robots that rely on their body dynamics to generate the motions and forces necessary for climbing. These robots have utilized their body dynamics to climb more efficiently [13] and to augment friction-only based contacts to generate upward forces [9]. Our dynamic climber, DynoClimber, depicted in Fig. 1, was inspired by biological studies which document remarkable similarities in the force patterns generated by rapidly climbing animals of different species [6], [14]. These findings suggest that animals employ large lateral in-pulling forces and body rotations to achieve fast, self-stabilizing gaits [14], [15]. DynoClimber was designed to isolate the dynamics of climbing force production from the attachment problem. As such, it was built to ascend a prepared (carpeted) vertical surface using curved aluminum claws which grip the substrate during leg retraction, but can freely slide during leg recirculation. By yawing as it pulls itself upwards, DynoClimber reproduces the force and motion profiles shown in animals and demonstrates self-stabilizing upward motions. The combination of an appropriate body morphology, series passive compliance introduced to mediate attachment impacts, and parallel passive compliance introduced to smooth the power draw on the actuators was shown to generate favorable body dynamics, allowing climbing at documented speeds approaching 1 bodylength per second (30 cm/s) [15].

Though impressive, the performance of this climber had not yet reached the levels predicted by scaled versions of the animal-based dynamical "template" [14]. One cause of this previously documented performance gap is that the robot exceeded its design weight. In the present revised design we have tuned the drive mechanism and passive spring constants, as well as increased actuator power, to address this. More fundamentally, under scrutiny it became clear that there were inefficiencies in the patterns of actuator recruitment and coordination developed by the machine's control algorithm. Here, we replace the work-directed self-exciting hybrid controller presented in [15] with one that is based on similar principles, but implemented in a more rational manner as a smooth vector field. Our new smooth controller permits an analytical proof of stability, and it improves performance by reducing former inefficiencies. Indeed, DynoClimber has proven to be an ideal platform for developing and testing alternative legged control paradigms: whereas the resultant climbing behavior is strongly dependent on the dynamical interactions between the electromotive and mechanical subsystems, its

Fig. 1.   Picture of the robot with annotations.

state space is sufficiently low dimensional to promote analytical tractability.

This paper is organized as follows. A comparison of various alternative algorithms used in the active control of legged locomotion is summarized in Section II. Section III introduces the central focus of this paper, the work-directed controller, and presents a proof of the stability of this new family of active legged climbing force production algorithms. Section IV offers our first, preliminary experimental results comparing a more traditional clock-based reference position generating controller to the new paradigm implemented on the DynoClimber platform. Section V summarizes the results and discusses future work.

II. Controllers for Active Production of Climbing Forces

In this paper, we examine two extremes of the feedforward-feedback spectrum in gait generation. The new class of controllers constructed here explores the feedback end of the spectrum [16], [17]; in our instantiation, the robot computes a pair of actuator terminal voltages solely as a function of the instantaneous limb positions (as measured by the instantaneous motor shaft angles) with neither the intermediating construction of a reference trajectory nor with any recourse to an internal clock. In this sense, our construction bears closest resemblance to the work-directed algorithms of Raibert's hoppers [18] or Buehler's Scout [19] that scheduled actuator energy expenditures as a function of limb state, using no additional clock state. Feedback-based excitation ("self-excitation") has also been explored outside the realm of legged robots in the context of dynamically

dexterous manipulation [20], [21], [22]; this work is characterized, as are [18] and [19], by a scheduled and intermittent application of force. The central difference between these earlier work-directed locomotion schemes and ours is that our use of discrete event-triggered hybrid control is limited to the introduction of a smooth, piecewise analytic function applied to a very general actuator model rather than a more extensive (and generally non-smooth) "case-based" logic applied to a specific body model.

Many of the most successful dynamic legged robots have occupied the other end of the spectrum, in which the gait is primarily excited by a feedforward Central-Pattern Generator (CPG). At the heart of these gait generation schemes lies a periodic signal generator, whether a CPU-originated reference phase trajectory [23] or a combination of neural oscillators [24]. This signal, the clock, is "shaped" to produce a position- (and, often, velocity- and acceleration-) based a reference signal for the limbs which, in consequence of their actuators's efforts to track the reference signal, transmit the appropriate ground reaction forces to the mass center as they interact with the substrate during the stance phase of a stride. Mechanical limb compliance modulates the feedforward based energy input to produce self-stabilizing behavior. Running robots such as Tekken[25], RHex [12], and Sprawlita [26] have all utilized variants of this approach to run effectively, and feedforward approaches scale readily to systems of higher dimensions. CPGs also facilitate the incorporation of both gait transitions and feedback into a robot's behavioral repertoire. However, it has proven challenging and time-consuming to tune up highly optimized peformance in these systems because: (i) the use of position tracking errors to call up actuator torques represents an indirect approach to recruiting the active power available within a robot's energy reservoir; (ii) the appropriate timing and magnitude of the errors that achieve this indirect recruitment require significant empirical tuning [27] because the analytical basis for their stability remains imperfectly understood [28], [29], [30]; and (iii) small variations in the clock-generated reference signal have large and somewhat unintuitive effects on the resulting behavior [28], [31], [32].

Our controller is further differentiated from those discussed above by its flexibility in specifying a desired work outcome. Specifically, we attempt to encode the control task "apply the maximum possible mechanical power while maintaining synchronization". This "work-directed target" can not be achieved explicitly with the above schemes. Buehler's Scout, for instance, relies on a feedforward torque controller when a leg is in ground contact, partially accomplishing the explicit work goal, but then relies on a PD controller to track a reference position signal designed to encode the coordination task during the robot's flight phase [19]. The recourse to reference signal tracking mitigates against a machine's achieving maximal power output. In our experience, with extensive tuning, errors generating near-optimal power output can be engineered throughout the gait. However, the dependence of actuator power output upon tracking error precludes the direct specification of a work-target, as tracking

error can not be neatly forced to the desired level.

A. DynoClimber's Control Space

Dynoclimber's electromechanical state space is quite simple as it possesses only two actuators - one motor per leg. The extension of the leg is slaved to the rotation of the motor via the kinematics of a crank-slider four-bar mechanism [15]. The kinematic configuration space of its actuators, then, is $\mathbb{S}^1 \times \mathbb{S}^1$, the 2-torus, $\mathbb{T}^2$. Of course, its motors have inertia and its limbs have mass, so the full configuration space is the tangent bundle $T\mathbb{T}^2 \approx (\mathbb{S}^1 \times \mathbb{R}^1) \times (\mathbb{S}^1 \times \mathbb{R}^1)$. The robot itself has many more degrees of freedom (body rotations and passive deformation of elastic elements among them), but the electromechanical state space just described is most immediately affected by the control policy while reflecting back enough of the "load" (body-substrate) mechanics to present a useful abstraction of the machine's climbing state. In contrast, a careful analysis of the complete, hybrid switched, high degree of freedom, compound pendulum that couples this electromechanical abstraction to the physical environment goes well beyond the scope of our present understanding. The output of the controller is specified as a pair of voltages applied to the two motor terminals of the robot at each time step. The position of each limb is sensed by a high-resolution optical encoder attached to each joint, and the velocity of each limb is computed by smoothing the computed changes in limb positions.

B. Baseline Controller

To provide a basis of comparison for our exclusively feedback-driven controller, we construct a benchmark feedforward controller. This controller is adapted from [14], in which the legs of the robot in simulation are kept exactly 180 degrees out of phase. Specifically, we command a constant frequency trajectory in motor shaft space ($\mathbb{T}^2$):

$$r(t) = \left[\begin{array}{c} \mathrm{mod}(2\pi f t + \pi, 2\pi) - \pi \\ \mathrm{mod}(2\pi f t, 2\pi) - \pi \end{array} \right] \quad (1)$$

This trajectory moves both reference legs at a constant velocity through motor-shaft space such that the shafts rotate with frequency f, while keeping both leg targets 180 degrees out of phase. The reference trajectory specified here is geometrically identical, for some f, to the emergent limit cycle generated by our self-exciting controller, making this reference trajectory the exact feedforward analog of our feedback controller.

In order to track the reference trajectory, the following control inputs are used, where $\Theta(t)$ is a 2×1 vector containing the motor shaft angles at time t:

$$\left[\begin{array}{c} V_1(t) \\ V_2(t) \end{array} \right] = k_p(r(t) - \Theta(t)) + k_d(\dot{r}(t) - \dot{\Theta}(t)) \quad (2)$$

The gains k_p and k_d are empirically chosen to yield desirable behavior and tracking of the reference trajectory, while $V_i(t)$ is the voltage applied across the i^{th} motor terminal. The tracking controller presented here does not claim to provide asymptotically exact tracking. In order to do so, it would

require an accurate model of the system parameters. Given the difficulty of establishing an accurate system model in the context of a legged robot, we implement and empirically tune this PD controller. Indeed, the application of a carefully tuned PD controller is a routine practice, applied to the tracking of a periodic reference trajectory on both the RHex[11] and RiSE[6] platforms.

C. Self-Exciting Controllers

The first controller used on DynoClimber was a "mirror law" [15]. Inspired by an earlier generation of juggling robots [33], [22], the controller generated periodic motion of the legs without using an internal clock. Moreover, this algorithm guaranteed that the robot's stance leg would be commanded to apply the maximum possible voltage. However, the hybrid controller was specified by a non-smooth function that switched based upon the stance and swing states of the limbs, which sometimes resulted in jerky, halting motions during certain hybrid state transitions when climbing.

We now "replace" the hybrid mirror-style excitation with a smooth family of controllers which is similarly self-excited, yet incorporates a number of relative advantages. Our new controller admits an analytical proof of stability with a very general system model and, unlike its mirror-style predecessor, does not require any gain tuning to insure tracking. Because the mirror-law employed a PD tracking controller to generate motor voltages during leg recirculation, an inadequately tuned controller would result in a recirculation which is either too slow or too abrupt, affecting the phase difference between the legs as well as the overall velocity of the system. In contrast, the work-directed scheme only requires tuning to alter the transient behavior of the system, without affecting the eventual limit cycle, as demonstrated by our proof of stability. Moreover, we note that simply pulling "as hard as possible" in stance – as the former controller did – does not guarantee maximum phase velocity, as a recirculating leg can (and often did) "lag" the opposing stance leg. To increase overall phase velocity and therefore climbing speed, our new framework does not prioritize stance over recirculation and instead attempts to enforce an antiphase relationship between the legs, limiting power to the leading leg, as necessary. A user must specify only a base voltage; from this base voltage, the controller will generate a limit cycle with some constant velocity (the actual velocity depends on physical parameters of the motors and mechanisms). We hypothesized, in consequence, that a user could achieve higher performance dynamical climbing (higher velocities) with reduced time spent tuning.

III. THE WORK-DIRECTED, SELF-EXCITING CONTROLLER

A. Control Objective

There are three principles at work in our new controller for DynoClimber. First, we aim to maintain antiphase orbits of the robot's legs. This is not strictly necessary, as we really only need the assurance that at least one leg will maintain ground contact at all times (the leg in "flight" can vary

speeds with impunity). However, an antiphase relationship should, given our passive mechanical energy storage (each leg stretches a spring to store energy during recirculation [15]), be a solution which generates nearly constant motor-shaft velocity at a constant motor power output level. Second, we want to generate the most rapid climbing possible by injecting as much energy at the highest rate possible. Third, given the machine's target speeds and dynamic environment, we strive to build a controller which is largely model-independent; speaking practically, the less a controller design relies upon inevitably imperfect and uncertain dynamical plant models, the easier it will be to implement in any dynamic environment.

We demonstrate below that our work-directed ("pull as hard as possible") controller generates the desired limit cycle, and requires little knowledge of specific plant parameters.

B. Motor Model

We first introduce a simplified model of the physical motor system in order to provide some analytical basis for the success of our controller. This model is not intended to be highly accurate to our specific robot - rather, we construct a very general actuator model and prove that our controller functions as desired if applied to any actuator chosen from the general class.

The robot has two legs, each with its own identical motor, and we will denote the pair of motor shaft angles as $\Theta = (\theta_1, \theta_2) \in \mathbb{T}^2$, subject to the traditional second-order linear motor model:

$$\frac{\ddot{\theta}_i J R}{k_\tau} + \frac{\dot{\theta}_i}{k_v} = V \tag{3}$$

where θ is output shaft angle, J is the moment of inertia of the commutator, output shaft and mechanism, R is the winding resistance of the motor, k_τ is the torque constant of the motor, k_v is the speed constant of the motor (the back-emf term), and V is the terminal voltage. This is an equivalent model to that employed in [34], with the caveat that we base our analysis on a system which supplies voltage, not current, to the motors. It is important to note here that variable loading and frictional effects from the dynamics of climbing (manifested as substantial time variations in J and k_v) dominate the behavior of the system, and any forces applied to the foot are reflected through a highly backdriveable mechanism as torques applied to the motor. As we design a controller, then, we do not want to be heavily dependent on an accurate system model, as the parameters of this model could vary widely based on the operating regime of the robot. We therefore generalize the motor model from Eq. 3 to include all constant inertia, Rayleigh-damped, Hooke's Law spring potential mechanical systems of the form

$$k_2 \ddot{\theta} + k_1 \dot{\theta} = V \tag{4}$$

where $k_1, k_2 > 0$. We construct a controller which will achieve its goals regardless of the choice of k_1 and k_2, and believe that the generality of the model given here

reinforces the connection between theoretical guarantees and the behavior of the robot in a dynamic environment.

In order to represent our robot, we use two identical but independent actuator models, each standing in for one of the robot's motors and linkages:

$$k_2 \ddot{\Theta} + k_1 \dot{\Theta} = \begin{bmatrix} V_1(\Theta) \\ V_2(\Theta) \end{bmatrix} \tag{5}$$

The controller is designed to dynamically "couple" these putatively independent motors through a memoryless nonlinear output feedback law that respects their terminal voltage magnitude constraints and guarantees that in the absence of external perturbations they will converge as a coupled system to the desired limit cycle on the torus of paired shaft angles and its tangent space of paired velocities from almost every initial condition. In employing this abstraction we admittedly neglect the motors' crucial mechanical coupling through the body, and relegate the actual task-related properties of body state to the role of "noise" felt as unmodeled "load" perturbations on independent motor shafts. We turn to the mechanical design of DynoClimber [15] to demonstrate effective climbing as long as its legs are maintained in a roughly antiphase relationship. In further defense of our coarse abstraction we observe that these models are sufficiently complex that so far the only analytical results for work-directed controllers encompassing physical actuator models explicitly coupled to the physical body state model have been obtained for one degree of freedom bodies (e.g. such as [35]) and that we see the present analysis as a first step along the way to that more informative but far less tractable problem. We also observe that no smooth work-directed scheme has heretofore been shown to converge even on $\mathbb{T}^2$.

C. Controller Definition

Formally, letting $V := (V_1, V_2)$ be the voltage command signal and $\delta := \theta_1 - \theta_2$ we take

$$V(\Theta) = V_{Max} \begin{bmatrix} 1 \\ 1 \end{bmatrix} - h(\delta) \cdot \begin{bmatrix} \mathrm{u} \circ \sin(-\delta) \\ \mathrm{u} \circ \sin(\delta) \end{bmatrix} \tag{6}$$

where the unit step function, u, outputs the scalar value 1 if its argument is positive and outputs 0 elsewhere, while $h : \mathbb{S}^1 \to \mathbb{R}^1$ is any smooth, even, positive function that vanishes if and only if its argument is 0 or π. Conceptually, the step functions serve to retard exactly one leg of the robot at a time, choosing to weaken the leading leg as necessary to guarantee convergence of the two legs into a limit cycle. For purposes of the present implementation we have chosen the specific "weakening" function

$$h(\delta) := k_{\mathrm{diff}} \sin^2(\delta)$$

as it is simple, tunable (giving a choice of $0 < k_{\mathrm{diff}} \le 2$), and meets the criteria imposed upon $h(\cdot)$.

Combining controller and plant, our system is

$$k_2 \ddot{\Theta} + k_1 \dot{\Theta} = V_{Max} \begin{bmatrix} 1 \\ 1 \end{bmatrix} - h(\delta) \cdot \begin{bmatrix} \mathrm{u} \circ \sin(-\delta) \\ \mathrm{u} \circ \sin(\delta) \end{bmatrix} \tag{7}$$

To verify that our control input is smooth, we show that our term containing step functions,

$$v(\delta) = h(\delta) \cdot \begin{bmatrix} \mathrm{u} \circ sin(-\delta) \\ \mathrm{u} \circ sin(\delta) \end{bmatrix}$$

is differentiable. First, for $\delta \in (0, \pi)$), noting that $\mathrm{u} \circ sin(\delta) = 1$ and $\mathrm{u} \circ sin(-\delta) = 0$,

$$\frac{dv}{d\delta} \mid_{\delta \in (0,\pi)} = dh/d\delta \cdot \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

and for $\delta \in (-\pi, 0)$, similarly,

$$fracdvd\delta \mid_{\delta \in (-\pi, 0)} = dh/d\delta \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Because h is nonnegative and smooth with isolated zeroes when its argument is 0 or π, $v(0) = v(\pi) = 0$, and $dh/d\delta \to 0$ as $\delta \to 0 \, or \, \pi$ from either side. Since the derivative of a step function is undefined at 0, we define $(dv/d\delta)(0) = (dv/d\delta)(\pi) = 0$. This makes $dv/d\delta$ continuous everywhere and demonstrates that our control input is smooth despite the presence of step functions.

A final informal observation about our controller: since the controller specifies motor voltages directly, it keeps at least one motor operating along its speed-torque curve at all times. Our present implementation provides no guarantee that it will not exceed the motor's sustainable current rating; using this control framework to specify voltages, as we have done here, can indeed require the motors to overheat. In DynoClimber, this has not been a problem throughout our testing, a trait we attribute to appropriate gearing and sufficiently powerful motors. In a more general case, identical analysis to that presented here can be used to specify motor current, with the substitution of a maximum desired current I_{Max}, for V_{Max}. Given this substitution, the present analysis carries through in a substantially similar manner, presuming the system contains mechanical damping (ie. it is crucial that the motor model is damped), as it surely does. In such a scenario, motor safety is guaranteed at the expense of maximal power output; we do not apply this approach on DynoClimber, as controlling for maximum power has not damaged our motors.

D. Proof of Correctness

For $\alpha \in \mathbb{S}^1$ denote the α-translate of the diagonal in $\mathbb{T}^2$ as

$$\Delta_\alpha := \{(\theta, \theta + \alpha) \mid \theta \in \mathbb{S}^1\}.$$

Proposition 1: The anti-diagonal tangent space,

$$T\Delta_\pi := \{(\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2) \mid \theta_1 = \theta_2 + \pi, \dot{\theta}_1 = \dot{\theta}_2\} \quad (8)$$

is an attracting invariant set whose domain includes $T(\mathbb{T}^2) - T\Delta_0$.

Proof: Rewrite (7) in the new coordinates,

$$\begin{bmatrix} \rho_1 \\ \rho_2 \end{bmatrix} = \begin{bmatrix} \theta_1 - \theta_2 \\ \theta_1 + \theta_2 \end{bmatrix} \quad (9)$$

yielding

$$\begin{bmatrix} V_1(\Theta) \\ V_2(\Theta) \end{bmatrix} = \begin{bmatrix} \frac{\ddot{\rho}_1 + \ddot{\rho}_2}{2} \\ \frac{\ddot{\rho}_2 - \ddot{\rho}_1}{2} \end{bmatrix} k_2 + \begin{bmatrix} \frac{\dot{\rho}_1 + \dot{\rho}_2}{2} \\ \frac{\dot{\rho}_2 - \dot{\rho}_1}{2} \end{bmatrix} k_1 \; .$$

Solving for $\ddot{\rho}_2$ in the second equation, substituting it into the first, and simplifying yields

$$\ddot{\rho}_1 k_2 + \dot{\rho}_1 k_1 = - h(\rho_1) \cdot \mathrm{u} \circ sin(-\rho_1)$$
$$+ h(\rho_1) \cdot \mathrm{u} \circ sin(\rho_1) \quad (10)$$
$$\ddot{\rho}_2 k_2 + \dot{\rho}_2 k_1 = V_1(\Theta) + V_2(\Theta) \quad (11)$$

Noting that ρ_1 is decoupled from ρ_2, we introduce a LaSalle function over $T\mathbb{S}^1$,

$$E(\rho_1, \dot{\rho}_1) = k_2 \cdot \frac{\dot{\rho}_1^2}{2} - H(\rho_1);$$
$$H(\rho_1) := \int_0^{|\rho_1|} h(x) \, dx \quad (12)$$

$h(\rho_1)$ goes to 0 smoothly as $\rho_1 \to 0$, so H is smooth. By construction, $h(x) > 0 \, \forall x \in \mathbb{S}^1 - \{0, \pi\}$, and $h(0) = h(\pi) = 0$. $H(\rho_1)$ is strictly decreasing in $|\rho_1|$, and therefore takes its minimum at π and its maximum at 0, with no other critical point. It follows that $(\pi, 0)$ is the unique minimum of H.

Taking the time derivative of E along the motions of the system, and recalling that $h(\cdot)$ is an even function, we find

$$\dot{E}(\rho_1, \dot{\rho}_1) =$$
$$k_2 \dot{\rho}_1 \ddot{\rho}_1 + h(\rho_1) \dot{\rho}_1 (\mathrm{u} \circ sin(-\rho_1) - \mathrm{u} \circ sin(\rho_1)) \quad (13)$$

After substituting $\ddot{\rho}_1$ from Eq. 10 and cancelling terms, we obtain

$$\dot{E}(\rho_1, \dot{\rho}_1) = -k_1 \dot{\rho}_1^2 \quad (14)$$

Thus, $\dot{E}$ is negative semidefinite and E is a suitable LaSalle function.

Examining the inverse image,

$$\dot{E}^{-1}(0) = \{(\rho_1, 0) \mid \rho_1 \in \mathbb{S}^1\} \quad (15)$$

we find the only invariant subsets of $\dot{E}^{-1}(0)$ occur at the zero section corresponding to the critical points of H, i.e., when $\dot{\rho}_1 = 0$ and $h(\rho_1) = 0$, which implies that $\rho_1 = 0$ or $\rho_1 = \pi$. Since $(\pi, 0)$ is a minimum of E, while $(0, 0)$ maximizes E in ρ_1, the former is an attractor and the latter a repellor, and the result follows.

QED

Corollary 1: The restriction dynamics on the attracting invariant submanifold $T\Delta_\pi \approx T\mathbb{S}^1$ gives rise to an almost globally asymptotically stable limit cycle.

Proof:

On $T\Delta_\pi$ we have $(\rho_1, \dot{\rho}_1) = (\pi, 0)$, hence, the restriction dynamics are given by $\ddot{\rho}_2 k_2 + \dot{\rho}_2 k_1 = 2V_{Max}$, and the system yields a single attracting limit cycle of the form

$$(\rho_2, \dot{\rho}_2)(t) = (\rho_2(0) + \omega t, \omega) \quad (16)$$

where $\omega := 2V_{Max}/k_1$.

QED

IV. EMPIRICAL VERIFICATION

A. Robot Overview

In order to establish the potential practical value of this new controller, we compare its performance to that of the baseline CPG-clock controller on our experimental platform, DynoClimber. As shown in Fig. 1, DynoClimber is an improved version of the climber described in [15].

The most significant changes to DynoClimber v1.0 as used here, are an increase in mass to 2.6kg due to the addition of new electronics designed to improve data logging capabilities, improve robustness, and allow operation of the motors at voltages up to 35 volts. The basic physical parameters are summarized in Tab. I.

It bears note that DynoClimber employs a roll-stabilization bar. Excessive rolling can cause the robot to lose attachment from the wall, and the robots roll dynamics are not actively controlled. A horizontal bar is affixed to the bottom of DynoClimber, which dramatically reduces the magnitude of roll generated and improves attachment.

TABLE I
PHYSICAL PARAMETERS FOR THE ROBOT

Body size	$400 \times 116 \times 70$ mm (excluding cables)
Body mass	2.6 kg
Wrist Spring Stiffness	640 N/m
Arm Spring Stiffness	140 N/m
Motor	Maxon RE 25 118752
Gear head	Maxon Planetary Gearhead GP 32A 114473 33:1 Gear ratio
Encoder	Maxon digital encoder HEDS55 110515 500 count/turn
Bevel Gear	2:1 reduction
Leg stroke	120 mm

B. Experimental Procedure

Experiments were conducted on a 12 foot vertical carpeted climbing surface. The robot utilized a tethered power supply at 32v with a peak current capacity of 11A. The external power supply was utilized rather than onboard batteries in order to generate repeatable and controlled runs. Furthermore, in order to facilitate robust and rapid data collection, an ethernet cable was tethered to the robot from a nearby laptop. Finally, the robot was also belayed using a 3mm diameter compliant rope in order to reduce the impulsive force on the robot in the event of a fall. This safety rope also supported both tethers; as the rope was hoisted as the robot climbed, the cables were pulled upward alongside the robot, minimizing the effect of the tethers on the dynamics of the climber. All control computation was conducted on-board, and a controlling laptop was used only to send high level commands and collect data. The robot was labeled with a center of mass marker; video was post-processed to compute the robots trajectory.

C. Specific Controller Chosen

The baseline controller was chosen, as described in section II-B, to be a PD controller driving the system toward a reference trajectory. The reference was a constant-velocity path in motor-shaft coordinates with the two legs exactly 180 degrees out of phase (see Eq. 1), resulting in near-sinusoidal motion of the robot's wrists. Both the frequency of the trajectory and the gains used by the PD controller were chosen empirically to maximize climbing speed. A driving frequency of 3.25hz was determined to be the peak achievable frequency of oscillation of the legs of the robot with this trajectory (while climbing), and the PD gains used to achieve this maximum frequency were $k_p = 0.9$, $k_d = 0.25$.

The self exciting controller described in Section III-C was implemented, as described, with the retarding function $h(x) = k_{\text{diff}} \cdot \sin^2(x)$. The retarding gain, k_{diff}, determines the transient behavior of the system - a larger k_{diff} forces the system to converge more quickly, at the expense of speed of oscillation during the transient period and any time the system is perturbed from its limit cycle. With a k_{diff} near 0, on the other hand, the system returns more slowly to its limit cycle during any transient period, but both motors are, on average, commanded higher voltages while the system is away from its limit behavior. As shown in section III-D, regardless of the choice of k_{diff}, the system provably converges to a limit cycle with a velocity which does not depend on the retarding gain. Moreover, the controller, as long as k_{diff} is kept between 0 and 2, will not exceed the specified maximum voltage, V_{Max}. For the following experiments $k_{\text{diff}} = 0.5$.

D. Experimental Results

The first result worthy of note is a confirmation that the controller induces the desired behavior from the system. Figure 2 plots the configuration space of the robot, $\mathbb{T}^2$, and shows actual encoder data from the robot while climbing (single datapoints) against the control vector field (arrows on graph). The torus is subdivided into quadrants, each quadrant corresponding to a different combination of left and right leg-states and is labeled as such. The beginning of the run is


Fig. 2.   The control vector field and a real trajectory

given by the stray 'strand' just above the middle of the plot. Rapidly, the encoder angles hone in on the solid red line in the plot - this is the line at which $\theta_1 - \theta_2 = \pi$. Although the controller was only demonstrated analytically to function for a simplified model, the generality of that model lent enough strength to the result that it holds quite well on the actual robot. These data were taken while the robot was accelerating from a standstill to full speed dynamical climbing, and despite this the legs converge quickly to the neighborhood of the limit cycle specified by the controller. It bears note that several factors are at work in preventing exact convergence to the limit cycle. First, the robot's legs are subjected to large periodic disturbances as the robot climbs (any unmodeled loading of the motors); hence, we see systematic and periodic differences between the desired limit cycle and the limit cycle achieved. Secondly, the k_{diff} chosen is fairly low - this causes the controller to retard the leading leg less forcefully and effectively places a premium on absolute speed over rapid convergence of the legs to the limit cycle. This trade-off does well for the robot, as demonstrated by its steady-state climbing speed.

Having established that the controller functions as intended, we turn our attention to a comparison between this self-exciting controller and a traditional feedforward alternative. In establishing a basic benchmark with the feedforward controller, we found that at leg frequencies above 3.25hz, the controller faltered. However, using the self-exciting controller, the robot climbed smoothly with leg frequencies up to 3.9hz. This increase in frequency is a demonstration of the judicious use of power in the self-exciting controller. By permitting the legs to spin "as hard as possible," a very substantial increase in stride frequency is achieved.

Finally, the CPG based controller achieved a maximum speed of 54cm/s, while the self-exciting controller reached a top speed of 66cm/s. Thus, the self-exciting controller managed to exceed the traditional controller's performance with 20% faster leg movement and a corresponding 22% increase in top speed. Figures 3 and 4, give a center of mass trace for the climber, and its vertical displacement and velocity, respectively.

V. CONCLUSION AND FUTURE WORK

We have presented a controller which has enabled DynoClimber to break into new climbing regimes. By stripping the controller of any internal states and and attempting to encode in it the task of performing maximal possible work against gravity, our robot achieves the fastest dynamical legged climbing yet seen: vertical locomotion at speeds up to 1.5 body lengths/second (66cm/s).

In order to examine the behavior of the system more thoroughly as well as continue to expand DynoClimber's performance, additional comparisons between this new family of controllers and more traditional schemes are necessary. An experiment examining the robot's behavior under varied payload weights will be used to examine our hypothesis that the self-exciting controller requires less tuning for higher performance. Similarly, the comparative energy efficiency
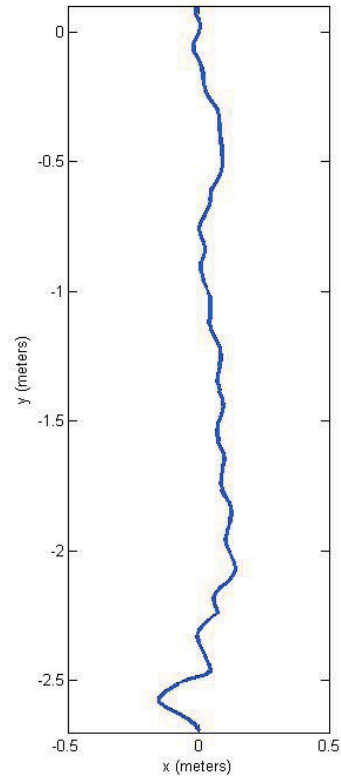

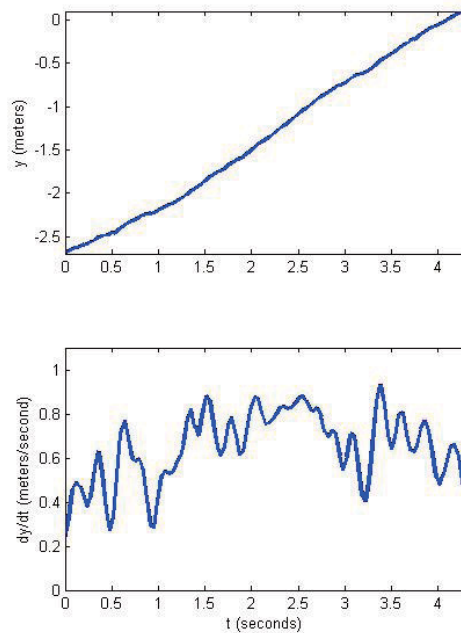Fig. 3. Self exciting controller center of mass. Elapsed time = 4.4s

Fig. 4. Self exciting controller vertical displacement and velocity

of the competing control paradigms may be studied by monitoring power usage during rapid climbing. Finally, the behavior of the robot after an unanticipated perturbation is also of interest and will be tested experimentally.

We believe the self-exciting control scheme presented here holds substantial promise as a controller for other legged platforms. The controller is able to both generate a desired limb coordination pattern and encode the desire for optimal power output, while requiring no parametric knowledge of the motors or system it is controlling. The promise of reduced tuning time coupled with improved performance as well as analytical guarantees make the self-exciting control paradigm an appealing option for other robots as well.

REFERENCES

[1] M. Spenko, J. Saunders, G. Haynes, M. Cutkosky, A. Rizzi, R. Full, and D. Koditschek, "Biologically inspired climbing with a hexapedal robot," *Journal of Field Robotics*, vol. 25, pp. 223–242, 2008.

[2] C. Balaguer, A. Gimenez, J. Pastor, V. Padron, and C. Abderrahim, "A climbing autonomous robot for inspection applications in 3d complex environments." *Robotica*, vol. 18, pp. 287–297, 2000.

[3] "http://www.vortexhc.com/vmrp.html."

[4] J. Xiao, A. Sadegh, M. Elliot, A. Calle, A. Persad, and H. M. Chiu, "Design of mobile robots with wall climbing capability," in *Proceedings of IEEE AIM, Monterey, CA, Jul. 24-28*, 2005, pp. 438–443.

[5] S. Kim, A. Asbeck, M. R. Cutkosky, and W. R. Provancher, "Spinybot ii: Climbing hard walls with compliant microspines," in *Proceedings, IEEE - ICAR, Seattle, WA, July 17-20*, 2005.

[6] K. Autumn, M. Buehler, M. R. Cutkosky, R. Fearing, R. Full, D. Goldman, R. Groff, W. Provancher, A. Rizzi, U. Saranli, A. Saunders, and D. Koditschek, "Robotics in scansorial environments," in *Unmanned Systems Technology VII*, D. W. G. G. R. Gerhart, C. M. Shoemaker, Ed., vol. 5804. SPIE, 2005, pp. 291–302.

[7] S. Kim, M. Spenko, S. Trujillo, B. Heyneman, V. Mattoli, and M. R. Cutkosky, "Whole body adhesion: hierarchical, directional and distributed control of adhesive forces for a climbing robot," in *Proceedings of IEEE ICRA, 10-14 April*, 2007, pp. 1268–1273.

[8] C. Menon, M. Murphy, and M. Sitti, "Gecko inspired surface climbing robots," in *Proceedings of IEEE ROBIO, Aug. 22-26*, 2004, pp. 431–436.

[9] A. Degani, A. Shapiro, H. Choset, and M. Mason, "A dynamic single actuator vertical climbng robot," in *Proceedings of IEEE IROS*, 2007, pp. 2901–2906.

[10] A. A. Rizzi, G. C. Haynes, R. J. Full, and D. E. Koditschek, "Gait generation and control in a climbing hexapod robot," G. R. Gerhart, C. M. Shoemaker, and D. W. Gage, Eds., vol. 6230, no. 1. SPIE, 2006, p. 623018. [Online]. Available: http://link.aip.org/link/?PSI/6230/623018/1

[11] U. Saranli, M. Buehler, and D. E. Koditschek, "Rhex: A simple and highly mobile hexapod robot," *International Journal of Robotics Research*, vol. 20, no. 7, pp. 616–631, 2001.

[12] R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, J. Brown H. B, D. McMordie, U. Saranli, R. Full, and D. E. Koditschek, "Rhex: A biologically inspired hexapod runner," *Autonomous Robots*, vol. 11, no. 3, pp. 207–213, 2001.

[13] S. Jensen-Segal, S. Virost, and W. Provancher, "Rocr: Dynamic vertical wall climbing with a pendulum two-link mass-shifting robot," in *Proceedings of IEEE ICRA, Pasadena CA, May 19-23*, 2008.

[14] D. I. Goldman, T. S. Chen, D. M. Dudek, and R. J. Full, "Dynamics of rapid vertical climbing in a cockroach reveals a template," *Journal of Experimental Biology*, vol. 209, pp. 2990–3000, 2006.

[15] J. E. Clark, D. I. Goldman, P. Lin, G. Lynch, T. Chen, H. Komsuoglu, R. J. Full, and D. Koditschek, "Design of a bio-inspired dynamical vertical climbing robot," in *Robotics, Systems and Science*, 2007.

[16] E. Klavins, H. Komsuoglu, J. Robert, and D. E. Koditschek, *The Role of Reflexes versus Central Pattern Generators in Dynamical Legged Locomotion*. MIT Press, Cambridge, MA, 2002, pp. 351–382.

[17] E. Klavins and D. E. Koditschek, "Phase regulation of decentralized cyclic robotic systems," *The International Journal of Robotics Research*, vol. 21, no. 3, p. 257, 2002.

[18] M. H. Raibert, *Legged robots that balance*, ser. MIT Press series in artificial intelligence. Cambridge, Mass.: MIT Press, 1986.

[19] M. Buehler, R. Battaglia, A. Cocosco, G. Hawker, J. Sarkis, and E. Yamazaki, "Scout: A simple quadruped that walks, runs, and climbs." in *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 5, 1998, pp. 1707–1712.

[20] M. Biihler, D. E. Koditschek, and P. J. Kindlmann, "A family of robot control strategies for intermittent dynamical environments," *IEEE Control Systems Magazine*, vol. 10, pp. 16–22, 1990.

[21] A. A. Rizzi, L. L. Whitcomb, and D. E. Koditschek, "Distributed real-time control of a spatial robot juggler," *Computer*, vol. 25, no. 5, pp. 12–24, 1992.

[22] A. A. Rizzi and D. E. Koditschek, "An active visual estimator for dexterous manipulation," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 697–713, 1996.

[23] G. C. Haynes and A. Rizzi, "Gait regulation and feedback on a robotic climbing hexapod," in *Proceedings of Robotics: Science and Systems*, August 2006.

[24] A. J. Ijspeert, "A connectionist central pattern generator for the aquatic and terrestrial gaits of a simulated salamander," *Biological Cybernetics*, vol. 84, pp. 331–348, 2001.

[25] H. Kimura, Y. Fukuoka, Y. Hada, and K. Takase, "Adaptive dynamic walking of a quadruped robot on irregular terrain using a neural system model," *Advanced Robotics*, vol. 15, p. 859, 2001.

[26] J. G. Cham, J. Karpick, J. E. Clark, and M. R. Cutkosky, "Stride period adaptation for a biomimetic running hexapod," in *International Symposium of Robotics Research*, Lorne Victoria, Australia, 2001.

[27] D. E. Koditschek, R. J. Full, and M. Buehler, "Mechanical aspects of legged locomotion control," *Arthropod Structure and Development*, vol. 33, no. 3, pp. 251–272, 2004.

[28] ——, "Mechanical aspects of legged locomotion control," *Arthropod Structure and Development*, vol. 33, no. 3, pp. 251–272, 2004.

[29] R. Altendorfer, D. E. Koditschek, and P. Holmes, "Stability analysis of legged locomotion models by symmetry factored return maps," *International Journal of Robotics Research*, vol. 23, no. 10, pp. 979–999, 2004.

[30] R. M. Ghigliazza, R. Altendorfer, P. Holmes, and D. Koditschek, "A simply stabilized running model," *SIAM Journal on Applied Dynamical Systems*, vol. 2, no. 2, pp. 187–218.

[31] J. Weingarten, R. Groff, and D. Koditschek, "A framework for the coordination of legged robot gaits," in *IEEE International Conference of Robotics, Automation and Mechatronics, Singapore*, 2004.

[32] *Sensitive dependence of the motion of a legged robot on granular media*, vol. 106, 2009.

[33] M. Buehler, D. E. Koditschek, and P. J. Kindlmann, "Planning and control of a juggling robot," *International Journal of Robotics Research*, vol. 13, no. 2, pp. 101–118, 1994.

[34] I. Poulakakis, J. A. Smith, and M. Buehler, "Experimentally validated bounding models for the scout ii quadrupedal robot," in *ICRA*. IEEE, 2004, pp. 2595–2600.

[35] D. E. Koditschek and M. Buhler, "Analysis of a simplified hopping robot," *International Journal of Robotics Research*, vol. 10, no. 6, pp. 587–605, 1991. [Online]. Available: http://www.scopus.com/scopus/inward/record.url?eid=2-s2.0-0026372031&partnerID=40