# The Role of Reflexes Versus Central Pattern Generators in Dynamical Legged Locomotion

Eric Klavins, Haldun Komsuoglu, Robert J. Full and Daniel E. Koditschek

# 1 Introduction

## *1.1 A Research Agenda: Functional Biomimesis*

Animals execute locomotor behaviors and more with ease. They have evolved these breath-taking abilities over millions of years. Cheetahs can run, dolphins can swim and flies can fly like no artificial technology can. It is often argued that if human technology could mimic nature, then biological-like performance would follow. Unfortunately, the blind *copying* or mimicking of a part of nature [Ritzmann et al., 2000] does not often lead to the best design for a variety of reasons [Vogel, 1998]. Evolution works on the "just good enough" principle. Optimal designs are not the necessary end product of evolution. Multiple satisfactory solutions can result in similar performances. Animals do bring to our attention amazing designs, but these designs carry with them the baggage of their history. Moreover, natural design is constrained by factors that may have no relationship to human engineered designs. Animals must be able to grow over time, but still function along the way. Finally, animals are complex and their parts serve multiple functions, not simply the one we happen to examine. In short, in their daunting complexity and integrated function, understanding animal behaviors remains as intractable as their capabilities are tantalizing.

"Functional biomimesis" describes a design methodology embraced within our Computational Neuromechanics Project (CNM) [Koditschek, 1999] inspired by, but not slavishly yoked to, biological example. The trick in deriving practicable engineering design from biological understanding is to develop analogies at the appropriate level of abstraction. Biological materials are vastly different from any available to current engineering practice. Biological constraints on computation and signal processing are quite distinct from those facing contemporary control engineers. But the problems of performing work upon a dynamical environment are common, and even the simplest animals are masters at the kind of multi-functional capabilities yet absent from the horizon of contemporary robotics. These considerations have focused our attention on the **mechanical** end of the animal behavioral spectrum where the biology is most constrained by the physical world. By this approach, the engineering benefit should be maximized with the result of high performance machines. At the same time, we anticipate a substantial benefit to pure science arising from the ability to leverage necessary conditions resulting from these physical constraints.

## *1.2 Functional Biomimesis at Work: the RHex Robot*

RHex, our recently prototyped autonomous hexapod robot [Buehler, Saranli and Koditschek, 2000], depicted in Figure 1, represents a significant success for functional biomimesis. Weighing 7 kg and roughly the size of a shoebox (about 50 cm x 30 cm), this machine aims toward autonomous function in the rubble and debris caused by disasters such as earthquakes, fires and bombings. Despite its distinctly unnatural appearance and apparent kinematic simplicity, RHex is designed to mimic certain functions exhibited by sprawled posture, many-legged runners such as the cockroach species *Blaberus discoidalis*.

[Place Figure 1 Approximately Here]

Decades of accumulating evidence from biomechanics suggest that running animals – whether bipeds, quadrupeds, hexapods, or myriapods – perform work on their center of mass so as to exhibit the dynamics of a pogo stick in the sagittal plane [Full and Farley, 1999]. Their muscles are tuned and postures are selected so as to engage the ground in a compliant manner and bounce from stride to stride. Along with this general strategy, cockroaches and other sprawled posture animals – arguably among most the successful runners – combine an apparent commitment to compliant operation in the horizontal plane, simplifying significantly problems of pitch and yaw stabilization. Inspired by the cockroach example, we have designed RHex as a sprawled posture mechanism with passively compliant legs.

Recent studies [Kubow et al., 1999] underscore dramatically the previously reported [Delcomyn, 1985] evidence for highly stereotyped cockroach limb motion, quantifying the very few degrees of freedom actually used in steady state running (notwithstanding the potential for complex high dimensional motions afforded by these animals' many muscle groups and joints). Since the power densities of commercial electro-mechanical actuators are so limited, RHex borrows from considerations of biological function on this count as well replacing all but one degree of freedom per leg with a fixed (compliant) shape, locating just one actuator, a 20 W brushless DC servo motor at each of the machine's six "hips."

Accumulating empirical evidence confirms that the functional aspects of fast cockroach locomotion we have identified above do indeed appear to confer a high level of performance, morphological and even kinematic details notwithstanding. In a recently reported set of experiments [Buehler, Saranli and Koditschek, 2000], RHex negotiates widely varied and dramatically adversarial terrain at speeds exceeding one body length per second. Moreover, initial runs over a force platform bear out the presumption that RHex's controller may be tuned so that leg coordination patterns supporting center of mass trajectories indistinguishable from a pogo stick in the sagittal plane [Full et al., 2000] result.

## *1.3 From Work to Tasks: Coordination Control as Phase Regulation*

In running insects, the leg patterns that support the pogo stick template take the form of an alternating tripod gait wherein the first and third ipsi-lateral legs move in phase with the contralateral middle leg, all three anti-phase to the opposite tripod. In this fashion, a tripod acts as a *virtual* single leg, and the tripod pair is coordinated in the manner of a virtual biped [Full and Tu, 1990]. The simplest biologically plausible model of how the cockroach generates the alternating tripod running gait takes the form of a "CPG" hypothesis wherein the legs simply track a properly shaped feedforward periodic reference signal. In fact, animals do in general match this model with the very regular stereotyped pattern of their steady state limb motions. They tend to have the shortest swing period possible and as speed increases swing period changes little and stance period decreases. In this view of the animal's coordination control, the pogo stick template [Full and Koditschek, 1999] emerges as an immediate mechanical response, or "preflex", before any classical reflexes contribute [Loeb and Brown, 2000].

Naï ve though it may seem, such a clock driven preflex-stabilized mechanism is exactly the coordination controller presently employed in RHex. As depicted in Figure 2, a two-phase periodic signal describes the conceptual operation of a wheel borne on a single spoke. The two phases are determined by a speed profile (elaborated in Figure 6) intended to match the physical state of a leg. A slow speed phase corresponds to the contact ("retraction") state of a leg during which it acts like the spoke of a wheel to propel the axle forward. The much faster phase corresponds to the recovery ("protraction") state of a leg, during which the spoke must cover a much greater arc in a comparable amount of time as it swings around over the axle to return for its next contact. This idealized signal functions as a reference trajectory tracked by a proportional-derivative (PD) controller driving the hip motor at each leg. Leg coordination is simply imposed by delivering identical copies of the clock reference to the participating motors, respectively in the phase and anti-phase tripods, as depicted in Figure 2 (We defer to the discussion in Section 2.2.1 regarding oscillator models depicted in Figure 5: the explanation for the single and double circle icons used in Figure 2.)

[Place Figure 2 Approximately here]

Our initial performance studies of RHex [Buehler, Saranli and Koditschek, 2000] document that clock-driven preflex-stabilized coordination works quite well when in a steady state mode of operation. There is apparently significant "mechanical intelligence" exhibited by the variously coupled analogue computers that coordinate the intermittent application of ground reaction forces to the robot's body by virtue of its clock excited leg spring compliance. Yet, however capable, this intelligence is surely very local in nature. It seems intuitively clear that some additional feedback from the leg to its controller and then on to its neighbors should confer substantial benefits in the face of significant perturbations, rapid maneuvers, and, in general, operation away from steady state. Of course, the imperatives of autonomous operation in highly unpredictable environments demand exactly this sort of general operational capability. Thus, the question arises as to how to build a "reflex" layer over the existing preflex controller that takes advantage of and extends the mechanical intelligence rather than fighting it.

In the specific context of RHex's laboratory development, this issue emerges most sharply as we take the first steps toward improving and effectively utilizing the robot's sensory suite. The current implementation does not afford any task level perception whatsoever. However, future versions will employ a broad array of sensors proposed to increase RHex's energy efficiency and afford autonomous locomotion without a human operator. Clarifying the role of feedback and feedforward in simulation studies and RHex experiments will lead to the development of sufficient mathematical conditions for the stable coordination of dynamically effective gaits. Such formal results should stimulate novel refutable hypotheses for biology.

## 1.4 Two Models for Comparing and Challenging Control Architectures

At the present time a significant gap in our understanding of animal and machine locomotion surrounds the issue of sensor recruitment for coordination. As the central goal of the simulation study presented in this paper, we wish to understand controller capability along two architectural dimensions that appear to loom hugely for any complex system designed to perform work within and upon its environment. First, we seek to understand the manner and extent to which sensory data about the environment should be used to control actuators along an axis that might be termed "clock" vs. "reflex." Second, along an axis that might be termed "central" vs. "peripheral," we seek to understand the advantages and disadvantages of distributed coordination. The present paper explores the first of these axes in some depth, touching upon the second only in some speculative concluding remarks.

To this end we introduce two simple mathematical models of locomotion that we have endowed, in simulation, with various coordination controllers. The first model, depicted in Figure 3, is a bipedal bead on a rail (BBR) and the second, depicted in Figure 4, is a bipedal spring loaded inverted pendulum (BSLIP) similar to a pogo stick, but with two alternating legs.

[Place Figure 3 Approximately here]

The BBR model consists of a bead on a rail pulled by springy legs whose toes are alternatively fixed ("retraction" phase) or allowed to slide freely ("protraction" phase) on parallel rails. The mechanism propels its "body" directly through the work of its alternating leg springs when they are fixed to the parallel "ground" rails as we detail mathematically in the appendix. The only losses occur due to friction between the bead and the rail. This simple (essentially one degree of freedom) model eliminates the issues of posture, and almost all other extraneous degrees of freedom, focusing, thereby, on the role of coordination for the underlying task of moving around a body. In contrast, the BSLIP model consists of a mass affixed to two spring-loaded legs that may rotate around the mass. The mass is constrained to lie in the sagittal plane and is acted on by gravity. It is an extension of the spring loaded inverted pendulum model (SLIP) that has been proposed as a model for the motion of the center of mass of an animal (or robot) while running.

[Place Figure 4 Approximately here]

We will use the BBR model to illustrate comparisons - points of similarity and difference between various control architectures. We will use more realistic BSLIP model to illustrate challenges – aspects of the locomotion task that will defeat control architectures inadequately informed about the physical world.

For both models, the ground reaction forces imparted to their bodies by their legs determine center of mass acceleration. In turn, the body's twice-integrated position (together with a leg coordination strategy) determines the degree to which the leg springs are compressed while in retraction. The BBR model has only one body degree of freedom: fore-aft progress along the uniform rail. Hence, absolute position has no consequence. Since the exchange of body kinetic energy with leg spring potential energy does not play an important role in body translation, this model operates effectively in the quasi-static regime. In contrast, adding a vertical degree of freedom in the BSLIP model introduces the key interchange of kinetic and gravitational potential energy: in the BSLIP model, the body, operating in a true dynamical regime, can be tripped by poorly executed leg behaviors. A coordination controller for BSLIP must therefore work in harmony with a body stabilization controller, which is not the case for the BBR model.

# 2 Review: General Models of Mechanism and Environment

## *2.1 Mechanism: Design Along the Feedforward/Feedback and Centralized/Decentralized Axes of Control Architecture*

Two distinct schemes for coordination have generated significant discussion in the animal motion control literature. The first is coordination due to a *central pattern generator* (CPG) wherein the state of an actuator, independent of its environment, follows a pattern produced by some centralized, oscillatory mechanism. The second is reflex driven coordination wherein the state of an actuator is a function of its environment, which may include the surface of the ground or the states of other actuators. In the engineering control literature, these notions correspond roughly to *feedforward* or *open loop* and *feedback* or *closed loop* operation, referring to the flow of information between the animal or robot and its environment.

The biological reality of pattern generators – units of identified neural tissue that exhibit fictive locomotion in isolation; excite true motor activity in more intact animals; but cannot sustain the oscillation upon further subdivision – has been established [Orlovsky, Deliagina and Grillner, 1999; Pearson, 1976]. Beginning with the seminal work of [Cohen, Holmes and Rand, 1981/82] and continuing through more elaborated refinements [Cohen, Rossignol and Grillner, 1988; Cohen et al., 1992; Kopell, 1995; Kopell, 2000], a flourishing applied mathematical literature has employed the language and methods of dynamical systems theory to model the nature of such feedforward motor control signals.

In the robotics literature, purely feedforward driven control schemes have been examined by a small, but growing, community as well [Berkemeier and Desai, 1999; Swanson, Burridge and Koditschek, 1995; Ringrose, 1997; Komsuoglu and Koditschek, 2000].

In parallel, a long history of behavioral animal science has resulted in an equally large literature concerned with the role of reflexes in motor coordination. In his seminal analysis of insect leg coordination, Wilson [Wilson, 1966] proposed simple local rules capable of generating a great diversity of steady state insect gaits. Some of the most compelling recent work within this tradition is reported by Cruse [Cruse, 1990] respecting the disturbance recovery properties of stick insect locomotion. Cruse and his colleagues have discovered several mechanisms that, in a *kinematic* model of the stick insect, *Obrimus asperrimus* [Cruse et al., 1998], reproduce the gaits and recovery characteristics exhibited by the animal. These models, emphasizing the role of reflex, represent a feedback-based approach to coordinated locomotion. Research by Beer and colleagues [Chiel, 1992] has successfully demonstrated that such reflex generated coordination mechanisms can elicit robust locomotion in hexapod robotic platforms operating in the quasi-static regime.

Feedback coordination schemes not necessarily inspired by the natural world have, of course, been employed in robot control as well. For example, the "virtual model controllers" of Pratt and colleagues [Pratt, Dilworth and Pratt, 1997] impose joint torques on their walking machines' limbs in such a fashion that the center of mass exhibits the compliance of some desired quasi-static template – typically, an overdamped spring. The juggling machines built by the last author and his former students [Buehler, Koditschek and Kindlmann, 1994; Rizzi, Whitcomb and Koditschek, 1992; Rizzi and Koditschek, 1996] constitute another such class of purely reflex controlled robots. In the work on juggling, there is an explicitly dynamical (one degree of freedom) template – a mass bouncing on an active vertical post – designed to self-excite stable periodic motion at a prescribed recurrent height. A different one degree of freedom dynamical template, the hanging pendulum, is used in [Nakanishi, Fukuda and Koditschek, 2000] to elicit stable locomotion in a brachiating robot, again by self-excitation to a prescribed total (i.e., kinetic plus potential) energy.

The design axis we call centralized versus decentralized is less well understood and we only comment on it briefly, leaving further exploration as future work. On one hand, decentralization of control may be a matter of implementation convenience and on the other, a matter of computational necessity. The reflex control model for the stick insect, proposed in [Cruse, 1990], is completely decentralized, for example. There is no central processing unit that collects the state of the constituent limbs and redirects their actuation relative to some global prescription. Individual actuators are allowed very limited information concerning the state of their neighbors and the body's coordinated gait emerges from their mutual interaction via the body and ground. In contrast, the reflex schemes described in the previous paragraph presume the availability of a global view of limb and body state that can be accessed freely by the constituent actuators in making their contribution toward the global reference template. This same distinction between centralized and decentralized architectures can be made in the context of purely

feedforward control schemes. For example, RHex's feedforward controller presents a completely centralized architecture, as clearly depicted in Figure 2. In contrast, lamprey locomotion models of the kind championed by Cohen and colleagues [Cohen et al., 1992] have the form of coupled peripheral pattern generators. Recent research [Cohen and Boothe, 2000; Lewis et al., 2000] suggests that similar completely local feedforward coordination can yield robust legged robot locomotion in the quasi-static regime. The challenge is to understand whether and how such an approach might be used more generally – for example, in the dynamical regime.

In this paper we concentrate on the feedback/feedforward axis, developing a framework for thinking about coordination control that permits an integration of the two extremes in both the robotics and the animal motion control literature. We will do so by re-interpreting the reflex model in a form that promotes its incorporation within the more traditional mathematical framework of coupled oscillators.

## *2.2 Environment: Phase Regulation Mechanisms for Dynamical vs. Quasi-static Tasks*

### 2.2.1 Oscillatory Systems that Perform Work Can Only Regulate Phase by Adjusting Energy

We depict in Figure 5 two oscillator models commonly encountered in the literature on biological motor control but whose crucial differences represent a significant obstacle to the kind of unified perspective we hope to form. A mass whose position (length $l$ in the figure) describes a periodic function of time may be considered an oscillator possessed of a "phase" variable that measures the percentage a given cycle has completed relative to some agreed upon fixed reference such as its point of maximal extension along the line. Continuous signals cannot be both periodic and non-reversing (i.e., exhibiting nonnegative speed) unless they lie on a circle. This motivates the terminology "clock," as suggested by the icons appearing in the right hand column of the figure. Quasi-static models are characterized by masses whose velocities can be altered directly by applied forces as depicted in Figure 5(a). First order systems (i.e., their state is one dimensional), cannot be made to oscillate unless some regulatory switch is used to mediate the reversal of applied force as a function of extension relative to some reference lengths (that we simply denote by $l_+$ and $l_-$ for now). The resulting hybrid (i.e. mixed continuous and discrete state) system exhibits immediately an exactly periodic oscillation. Such models yield intrinsically stable clocks that may be instantaneously reset and which are tuned by position variables, the reference lengths, $l_+$ or $l_-$.

[Place Figure 5 Approximately here]

In contrast, models suitable for studying mechanical work have state defined by their velocities as well as the positions of their masses. Applied forces alter the acceleration of the body and, hence, can only affect position via a cumulative integrated effect. In the presence of positive stiffness potential forces such as the Hooke's law spring depicted in Figure 5(b), such systems oscillate with an intrinsically (neutrally) stable phase velocity

determined by their energy level. When they perform work upon a lossy environment (as in this study: recall that the leg springs of the BBR and BSLIP models are subject to viscous damping) then their energy gradually leaks out and the phase velocity decays to zero – effectively stopping the clock. There is no other way to maintain a lossy mechanical oscillator other than by pumping in energy, for example, as depicted here, by tuning the potential energy stiffness, $k$, or by adjusting the reference lengths $l_+$ and/or $l_-$. However, because phase results from the integrated effects of these energy changes, the resulting lag can destabilize the clock if the regulatory policy is not carefully applied.

### 2.2.2 Example: The Reflex as a Quasi-static Mechanism for Feedback Based Phase Regulation

A major objective of this paper is to recast a class of reflex models introduced by Cruse and colleagues [Cruse, 1990] in a form amenable to regulating the exchange of total energy rather than merely prescribing kinematic relationships between limb positions and velocities. In Cruse's model, each leg alternates between *protraction* (swinging the leg forward) to an anterior extreme position (*AEP*) and *retraction* to a posterior extreme position (*PEP*). Each leg receives a combination of signals, which we describe in more detail in the next section, from certain neighboring legs, which alters the *PEP* of the leg. Changing *PEP* effectively changes the period of the leg, synchronizing it in or out of phase with its neighboring legs. For example, in one mechanism, a leg in protraction signals the next rostral, ipsilateral leg to decrease its *PEP* thus causing the receiving leg to retract further. The combination of mechanisms operating in parallel according to an experimentally observed connection scheme, results in a decentralized controller model.

Beer and his group [Chiel et al., 1992] have successfully used three of Cruse's mechanisms to control the gait of a hexapod robot operating in the quasi-static regime. Furthermore, Beer and Calvitti [Calvitti and Beer, 1999] have begun to show analytically how these mechanisms might operate in a model that is, roughly speaking, a logical abstraction of the kinematics presented by our BBR model. In particular, they examine a two-leg arrangement wherein one leg (the master) influences the other (the slave) via one of the three mechanisms. There is no reciprocal coupling. The legs are not attached to a body nor do they ever touch the ground, providing no means of mechanical, or environmental, coupling.

The authors are able to prove that each mechanism results in a neutrally stable, attracting region of phase relationships centered near 180° out of phase. That is, any phase relationship in a certain range is possible, phase relationships outside the range are *pushed* into the stable region, but no particular phase relationship within the stable range is preferred. Such neutral stability is consistent with the hypothesis in [Wilson 1966] that one gait may smoothly transfer to another, although other mechanisms may produce this behavior as well. It remains to analyze situations where legs are mutually coupled, coupled with more than one of the mechanisms Cruse describes, or where there are more than two legs. Furthermore, a model that includes the body of the insect coupled with its environment warrants analysis to relate the stability of body attitude and velocity to the stability of gait.

### 2.2.3 Example: The CPG as a Dynamical Mechanism for Feedforward Based Phase Regulation

For the purposes of this paper, a "CPG" is an internally generated periodic signal used to produce a motor reference trajectory that is tracked by the actuators of the robot. Coordination results in a manner similar to that of a conductor leading an orchestra. In the control of RHex, we have chosen to realize this reference generator by passing a "clock" – a constant speed periodic signal source - through a monotone output function that "shapes" the commanded speed profile as depicted in the upper portion of Figure 5. The speed profile presently used in RHex is graphed in Figure 6

[Place Figure 6 Approximately here]

We model the clock as a harmonic oscillator, denoting its phase by $\phi$. The CPG thus effectively introduces an additional state whose dynamics is governed by the differential equation

$$\dot{\phi} = \frac{2\pi}{T_c},$$

where the parameter $T_c$ represents the period of the CPG, and the "shape" of the speed profile is set by adjusting such constants as those depicted in Figure 6 used in RHex.

Notice that the period of the harmonic oscillator essentially corresponds with its energy. Thus, we must explicitly manipulate the parameter $T_c$ in order to change the phase velocity of the clock. In other words, clocks effectively show up as first order (a single circle in the icon system introduced in Figure 5) oscillators in the study we report here. Although second order oscillators (e.g., Fitzhugh-Nagumo or other two dimensional nonlinear approximations to the Hodgkin-Huxley equations [Holmes, 1997]) have been proposed as models of neural dynamics, we prefer to adopt the reasoning of Cohen, Holmes and Rand [Cohen, Holmes and Rand, 1981/81] and operate in the more neutral "phase coordinates" known to exist in the neighborhood of any limit cycle oscillator[1].

The CPG approach to control, avoiding any active measurements, is by its nature immune to sensor noise issue, which is a big advantage in practical implementations. On the other hand, steady state behavior depends very strongly on the environmental parameters, and therefore, external disturbances to the system cause the system to vary its steady state body speed wildly, thereby eliminating coordination between legs. We have observed that the leg synchronization is maintained fairly well if the external changes are slow and do not drive the system into regions where no stable fixed operation exist.

In summary, coordination via a CPG does not require sensory information to generate control signals but instead drives the underlying cyclic mechanism in a purely feed-forward fashion.

# 3 Simulation Study: Application of Reflex and CPG Control to Two Dynamical Models

RHex is presently controlled via the centralized pattern generator depicted in Figure 2 that provides no mechanism for incorporating body, leg, or ground state feedback. On the other hand, the feedback mechanisms discussed in Section 2.2.2 target the quasi-static regime of operation. In this section we use the "cartoon" BBR and BSLIP models to provide simulation evidence that the reflex coordination mechanism must indeed be reworked for successful operation in the dynamical regime. We then offer a modified representation of the reflex mechanism and simulation evidence suggesting that it works quite effectively in the dynamical regime. We also use the models to explore the feedback/feedforward axis of dynamical coordination control, presenting a means of interpolating between the two.

## *3.1 The BBR and BSLIP Dynamical Models*

The BBR and BSLIP models are detailed mathematically in the appendix and it will suffice here merely to sketch the architecture of their dynamics and control blocks as in Figure 7. Each leg periodically forces the body, defining second order mechanical oscillators[2] - the double circles, according to the icon conventions established in Figure 5. The dynamics are *hybrid* continuous/discrete, discrete modes being defined by whether one, two, or no legs are touching the ground. Since that condition may be read off the pair of leg phase angles, the block labeled "body dynamics" may be correctly imagined as represented by an appropriate dynamical system on a torus (the cross product of two circular phases), as we have detailed in [Klavins, Ghrist and Koditschek, 2000]. Conceived of in this light, the job of the block labeled coordination controller is to regulate a pair of angles.

[Place Figure 7 Approximately here]

In the case of the BBR model, the coordination controller acts to regulate that pair of phase angles by prescribing the PEP for each leg, determining in this manner the next angle of leg liftoff. In the case of the BSLIP model, each of the double circles should properly be repeated as in Figure 8 since there is a second degree of freedom (vertical as well as fore-aft or, equivalently, radial and angular) that must be coordinated as well. In the present study we have, following Raibert's lead in the control of hopping robots [Raibert, 1984], employed a feedback controller whose radial and angular components are decoupled[3]. The radial degree of freedom is controlled according to a PD law that seeks to regulate height by adjusting the spring stiffness independently of the leg angle. The angular degree of freedom is controlled, by adjusting the touchdown angle (coincident with the liftoff angle magnitude at steady state). In our study, we connect the output of the coordination block to the BSLIP angular controller by interpreting *PEP* as a phase coordinate prescription from the former and as a liftoff angle prescription to the latter block.

[Place Figure 8 Approximately here]

Throughout the remainder of this section, we will report on the results of simulating the dynamical system depicted in Figures 7 and 8 with different coordination controllers substituted in the blocks marked "coordination control".

## *3.2 Dynamical Reflex Coordination Control*

### 3.2.1 Quasi-Static Reflexes Support "Walking" in the BBR Model

We first apply a reflex coordination scheme similar to that described by Cruse and studied by Beer and Calvitti (see Section 2.2.2). Of the coupling mechanisms described by Cruse, three are considered to be at work in regulating gait. In each, the *PEP* of one leg (the receiver) is offset based on the state of the other leg (the sender) as described in Beer and Calvitti's analysis of Cruse's work [Calvitti and Beer, 2000]. In our model, each leg serves both as sender and receiver. In the first mechanism, the *PEP* of the receiving leg is decreased below its nominal value by some small amount $\alpha_1$ when the sending leg is protracting, thereby delaying the start of protraction in the receiving leg. In the second mechanism, the *PEP* of the receiving leg is increased beyond its nominal value by a small amount $\alpha_2$ when the sending leg is in retraction and its length is within some small distance of its *AEP*, thereby shortening the period of the receiving leg. In the third mechanism, the *PEP* of the receiving leg is increased from its nominal value gradually, throughout the retraction mode of the sending leg, by the value $\alpha_3(AEP - l)$ where $\alpha_3$ is a tunable parameter and $l$ is the length of the sending leg.

As a kind of conceptual calibration to the previous literature – particularly the papers of Cruse and Beer mentioned above - we simulated the BBR model with each of the mechanisms in turn. With two-way coupling, we found, as did Calvitti and Beer in their model with one way coupling, neutrally stable gaits for each mechanism. In addition, we found that the forward velocity of the body was regulated to a steady state. We found the third mechanism to be the most aggressive and we present simulation data for it in Figure 9(a). Note that the average velocity seems to stabilize after an initial acceleration and that the legs approach out of phase alternation but do not quite achieve it.

[Place Figure 9 Approximately here]

### 3.2.2 Quasi-Static Reflexes Do Not Support "Running" in the BSLIP Model

We now replace the BBR model with the BSLIP and retain the reflex in the "coordination controller" block. It is not obvious that Cruse's model of insect reflexes, concerned as it is with the behavior of decidedly quasi-static insect walking, should extend to the control of insect running. Specifically, because it does not address the exchange of potential and kinetic energy, that coordination mechanism must be supplemented with some additional

scheme that can account for the second order dynamics of running captured by the BSLIP model.

Initially, we attempted to supplement a Cruse style reflex mechanism with a Raibert style gait controller. This simulation coordinates the BSLIP legs by adjusting their liftoff length - the analog of PEP in this setting. Control of the remaining gait parameters (hopping height, forward velocity and duty factor) is accomplished using mechanisms similar to Raibert's as discussed above. However, the domain of stability of this control scheme was quite small. Essentially, we found that changing the liftoff conditions interfered with the control of forward velocity and hopping height unless the synchronization gain, $k_{sync}$, was quite small, while a too-small $k_{sync}$ resulted in a lack of synchronization — meaning that at some point the mass hit the ground because neither leg was "ready" to touchdown. See the appendix for an explanation of $k_{sync}$. Examples of these behaviors are shown in Figures 10(a) and 10(b).

[Place Figure 10 Approximately here]

### 3.2.3 Dynamical Reflexes Use the Language of Coupled Oscillators

The quasi-static view of the world presumes that forces effect change in position. This is a useful fiction when operating in a highly damped regime, or in the context of behaviors that do not entail the exchange of kinetic energy. To understand intuitively the importance of managing explicitly the exchange of total energy intrinsic to mechanical systems that perform work in demanding regimes, consider the following analogy. When driving a car, there is a delay between the control input and the position output of the system. The driver of the car experiences this delay, for example, when he or she steps on the brakes: the car does not stop immediately but rather decelerates, gradually coming to a stop.

In this section we describe another, similar, coordination mechanism resulting from a formalization [Klavins, Koditschek and Ghrist, 2000] of our phase regulation algorithm for vertical juggling [Buehler, Koditschek and Kindlmann, 1994; Rizzi, Whitcomb and Koditschek, 1992] that presents (when viewed with the appropriate abstraction) an almost identical coordination problem to our present concern arising from locomotion. Implementation within the simple BBR and BSLIP models suggests that this representation of local reflexes is capable of imitating key properties of the Quasi-static style reflex coordination mechanisms introduced by Cruse, described above. Expressed within the traditional framework of coupled oscillators, it enjoys the added benefit of straightforward analysis via standard tools from dynamical systems theory. Moreover, as we will soon show, for this reason, it can be extended in a straightforward manner to handle the exigencies of the dynamical regime represented by our more complex BSLIP model.

The method, summarized in Figure 11, is based on defining an *artificial energy function, V,* on the difference in the phases of the two legs. Specifically, if we let $\boldsymbol{j}_1$ and $\boldsymbol{j}_2$ be the

phases of two legs, one suitable definition for $V$ is $\cos(\varphi_1 - \varphi_2)$. Minimizing the energy between the two phases via changing $PEP$ in a suitable fashion requires that $\varphi_1$ and $\varphi_2$ differ by $\pi$ – that is, that the two legs be exactly out of phase with each other. The resulting dynamics bears a strong resemblance to the provably stable system described in [Klavins, Koditschek and Ghrist, 2000]. We note that in our simulations we do not have direct access to the actual phase of a leg. Rather, we define an approximation to phase based on the length of the leg rather than on time. Unless the body moves backwards, this quantity increases monotonically with time and is an effective stand in for phase in the present application.

[Place Figure 11 Approximately here]

In simulation, the system regulates the forward velocity of the body to a steady state and regulates the legs to an alternating gait *exactly* 180° out of phase. The alternating gait appears to be stable.

The energy function may be defined by $\max\{\alpha, \cos(\varphi_1 - \varphi_2)\}$ for some number $\alpha$ between *–1* and *0*, to simulate the third mechanism described by Cruse, and results in neutrally stable gaits centered around 180° out of phase. In Figure 9 we compare the return maps, obtained numerically, of the model with different energy functions and with Cruse's third mechanism in place. We conclude that the phase-regulation method can mimic Cruse style coordination and provides and obvious way to specify gaits exactly by providing the appropriate energy function.

### 3.2.4 Dynamical Reflexes Support "Running" in the BSLIP Model

An alternative means of regulating the legs in this model is to use the fact that phase regulation is based on a *continuous* reference field which is obtained from the energy function *V* as in Figure 11. Thus, the legs may be synchronized at times other than liftoff events. Specifically, instead of changing *PEP*, which as we have shown substantially affects the motion of the body, we change the swing phase dynamics by modifying the stiffness of the hip springs, $k_i$, to be $k_{nom} - k_{sync}\sin(\theta_i - \theta_j)$ – that is, a nominal stiffness modified by a gain times the negative gradient of the "pseudo-energy" function (See the appendix for an explanation of the gain $k_i$). Thus, we suppose that the leg angles correspond roughly to phase and use the fact that $k_i$ controls the velocity of leg *i* during swing. Since we only change the speed of swing, we do not affect the underlying control, which affects the touchdown angle and the leg spring stiffness during decompression. Data from a typical run is shown in Figure 12. It is not obvious how such a coordination controller could have been constructed using the Cruse style reflex framework directly and in a natural way, although careful consideration of the effects of changing the liftoff length on the horizontal velocity may prove fruitful. We believe the proper route is to use the mechanisms described by Cruse as inspiration, but to work in the more informative language of coupled oscillators when designing controllers.

[Place Figure 12 Approximately here]

## 3.3 Exploring the Feedback/Feedforward Axis of Dynamical Coordination Control

### 3.3.1 CPG Based Coordination Control

We aim to develop a means of interpolation between pure feedforward and pure feedback coordination mechanisms. We have, just above, established a pure feedback mechanism that implements (Cruse inspired) leg angle corrections by means of phase velocity corrections (as described in Figure 11). To establish the desired correspondence with pure feedforward coordination, it is helpful to take the obverse step of reworking an absolute reference velocity based controller (inspired by the RHex coordination mechanism of Figure 2) into an absolute reference position based controller as follows. One can effectively realize the two-speed RHex velocity profile of Figure 6 by running a constant speed clock and issuing a "leg lift-off" command as its internal reference phase crosses a marker at $f = f_{pep}$. To coordinate two legs of an alternating gait, one places two markers, $f_{pep1}$ and $f_{pep2}$ on the circle which are $180^o$ out of phase.

As a feedforward mechanism, such a position-based clock has no better "idea" of the leg's true state than it's velocity reference generating equivalent of Figure 6. But the physical fact of whether or not the leg is in contact with the ground determines the physical consequences of the command. For example, the leg lift-off command has a physical effect if it arrives while the leg is in retraction. That command can have no physical effect if the leg is already in protraction when it is received. This observation hopefully now makes it clear that notwithstanding the absence of an explicitly implemented sensory pathway from the leg's mechanical oscillator back to the CPG, the following two relationships between actuation effect and phase difference provide a "hidden" mechanical feedback channel – the "preflexive" mechanism. First, the leg liftoff command may cause the leg to begin protracting at any point in retraction and thus alters the effective *PEP* length, and changes the leg cycle period. Second, the change in phase difference is a monotonic function of the difference between the periods of the leg and CPG. This mutual dependency causes the phase difference, sampled at *AEP* for example, to converge to a steady state constant value. We observe that the losses in the mechanical system and the above phase-actuation relationships are sufficient for asymptotic convergence to a steady state operating point.

Applied to a single legged version of the BBR model, this feedforward controller would yield a system very similar to one analyzed formally in [Komsuoglu and Koditschek, 2000] wherein sufficient conditions are given for the stability of a feedforward "clock" controlled hopping gait. In that work, stabilization can be formally demonstrated to obtain via the implicit interaction of the mechanical oscillator with the clock – the "preflexive" mechanism intuitively described above. [4] Applied to the complete BBR model of interest here, the feedforward controlled two legged systems are best thought of as a pair of one-way coupled oscillators - one internal (the CPG) and one mechanical (the leg) - as

depicted in Figure 2, corresponding to the setup in Figure 7 where the connection between the body dynamics block and the coordination controller (that block now occupied by the coordinated pair of reference marked clocks from Figure 6) has been severed.

Our simulation study of the CPG mechanism applied to the BBR model suggests that feedforward coordination of a pair of preflex stabilized legs functions as effectively as a preflex stabilized single leg. [4] Each leg converges to its respective "synchronized" operating point, the transient effects of the other leg are rejected as essentially external disturbances, and the body converges to a steady velocity. At the steady state operation, the legs are $180^o$ out of phase in accordance with the markers on the CPG, Figure 13(a).

[Place Figure 13 Approximately here]

### 3.3.2 Comparing CPG and Reflex Coordination of BBR

Having related reflex and clock style coordination through the language of coupled oscillators, we are now in a position to compare their performance in various versions of our BBR and BSLIP model. Although a detailed comparison across many parameters of interest lies beyond the scope of the present paper, it is worth presenting one specific example of the questions that might now be addressed, both analytically as well as computationally.

As might be expected, the CPG suffers serious limitations on the achievable gaits in an uncertain world. For a single legged version of BBR model, one can compute quite straightforwardly a CPG period that has a fixed operation point yielding a desired body velocity: choose a *PEP* leg length that does not yield net body speed change and integrate the differential equation over a complete leg cycle. This computation assumes there is at most one leg applying force to the body at any given time, hence, absent a "double stance" juncture, it applies as well to the full (two-legged) model provided that there is no double stance at the desired steady state. As this tuning procedure suggests, the CPG period cannot be chosen arbitrarily but rather needs to be in proper accord with such task parameters as the natural stiffness of the legs and rail width. In support of this, our numerical studies, sampled in Figures 13(a) and 13(b), indicate that the resulting fixed operation may be stable or unstable depending on the period of the CPG relative to the natural period of the BBR model. In general, the feedforward CPG controller cannot dictate an oscillation that is faster than the legs are physically capable of realizing. The equivalent problem in RHex would occur if its CPG oscillated faster than its motors were capable of spinning: the legs would end up playing a never ending, possibly unstable, game of catch-up with the CPG.

The previous intuitive arguments may be confirmed by direct simulation. Figure 14 depicts the relation of the steady state response (forward velocity $x^*$) of the BBR system with CPG control to the distance $w$ from the rail to the outer guide-rails. The black regions indicate that there is no fixed point for the $(w,x^*)$ pair. Gray regions indicate that there is a fixed point that is unstable. And, finally, the white region is the stable fixed

points. This figure exhibits the fixed-point dependence to the environmental settings in a feedforward control system. In contrast, Figure 15, exhibiting the same relationship for the reflex coordinated BBR, suggests the value of feedback as a hedge against environmental uncertainty. In contrast to the previous case, we see that the model can operate successfully at every rail width, with no a priori information about what that parameter might be.

[Place Figure 14 Approximately Here]

[Place Figure 15 Approximately Here]

.

# 4 Conclusion: Stabilizing Feedback Pathways for Mechanical Locomotion

RHex (Figure 1) functions quite ably in unusually demanding terrain on the strength of its "preflexes" alone. At present, RHex's CPG (Figure 2) receives no information regarding the state of the actuators, the legs they move, the body that translates in consequence, nor the terrain that body visits. It seems clear that significant advantage can derive from feedback, but it is not yet clear how best to insert sensory information in the purely feedforward control scheme. In this paper we have reported on work in progress addressing this question - the beginnings of a simulation study that promises to shed light on how to do so.

Specifically, we have introduced two simple computer models of bipedal locomotion: one, the BBR model, where the exchange of potential and kinetic energy of the body does not significantly impact the locomotion task; and the other, the BSLIP model, where it matters a great deal. We have explored within these two contrasting examples the applicability and efficacy of two general paradigms of biological motor control: CPGs and reflexes. To carry out this exploration we have found it useful to modify the quasi-static representation of reflex coordination originally proposed by Cruse, bringing it into correspondence with the traditional literature on coupled oscillators. Whereas the past representation succeeds in coordinating the BBR model, it leads to tripping and crashing in the BSLIP. The revised representation we introduce here affords a version of the reflex mechanism capable of coordinating the total energy as well as kinematic variables relevant to running.

It would be most desirable to select from a spectrum of possible coordination models that introduce more or less sensory information at will. However, engineers will confidently deploy an appropriate combination of these general coordination mechanisms only when similarly general principles addressing the efficacy of such combinations have been developed. The daunting modeling and analysis task requisite to such general understanding is arguably best approached by initial appeal to "cartoon" instances developed that capture the essence of the problems, yet which are more amenable to formal inquiry. Such a formal inquiry is the target of our longer-term research effort in this area.

# Acknowledgements

The Computational Neuromechanics Project is supported by DARPA/ONR under grant N00014-98-1-0747. The CNM Hexapod project is supported by DARPA/SPAWAR under contract N66001-00-C-8026. We thank Martin Buehler, Uluc Saranli, Phil Holmes, Devin Jindrich , and the other members of these projects for their shared insights and help with the specific questions addressed in this paper. We thank Art Kuo as well for a number of useful discussions bearing on the topics addressed.


# Appendix: Details of the Simulation Models

## The Bipedal Bead: A Simple Model

This model consists of a point mass, constrained to lie on a rail, with two massless legs connecting the body to one of two parallel rails as in Figure 3. The legs act as springs and pull the body along the rail. A leg can be either protracting (swinging forward, not grasping an outer rail) or retracting (swinging backward, grasping the rail and, therefore, pulling the body). In protraction, the length of a leg is less than *AEP* and increasing. In retraction, the length of a leg is greater than *PEP* and decreasing. (Note that in this model *AEP* and *PEP* are not angles but instead are considered, equivalently, to be leg lengths.).

A leg protracts until its length is greater than *AEP*, at which point it attaches to the rail and begins to retract. It retracts until its length is less than *PEP*, at which point it releases the rail and begins to protract again. The length of time between the beginning of the *k*th retraction and the *(k+1)*st retraction is called the *period* of the *k*th cycle. The length of the *k*th period is a function of the *PEP* of the leg and the velocity of the body at the beginning of the *k*th retraction. Thus, *PEP* is a natural candidate for a control input as it affects the period of leg, which is essential for gait regulation.

We denote by $s_i \in \{p, r\}$ the discrete state of leg $i$ and by $x$ the position of the body. The symbols $l_1$ and $l_2$ denote the lengths of the legs, and $x_{toe,1}$ and $x_{toe,2}$ denote the positions of the toes of the legs. Finally, $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$ denote the angles from the horizontal to leg one and leg two respectively. We assume that the body has a mass of one and that it is constrained in the plane to move along the line $y = 0$ and the toes along the lines $y = -w$ and $y = w$ respectively, where $w$ is the distance from the center rail to an outside rail.

Let $l_0$, $PEP$, $AEP$, and $l_{ex}$ be constant parameters such that

$$l_0 < w < PEP < AEP < l_{ex}.$$

The constant $l_0$ is the rest length of a leg in retraction. The constant $l_{ex}$ is the rest length of a leg in protraction. The constant $PEP$ is the nominal posterior extreme position of a leg. The actual posterior extreme position of leg $i$ is $PEP_i$ — a value that can be used as a control input and which is constrained to lie between $w$ and $AEP$. The constant $AEP$ is the anterior extreme position. If leg $i$ is protracting ($s_i = p$), then $l_i$ increases until it exceeds $AEP$ at which time it begins to retract, $s_i = r$, until its length is less than $PEP$, at which point it begins to protract again.

When a leg is in protraction, its movement does not affect the movement of the body. The leg simply extends and the toe glides along the rail away from the body. Since the legs are massless, we model a protracting leg as a first order spring:

$$\text{if } s_i = p \text{ then } \dot{l}_i = -k_{pro}(l_i - l_{ex})$$

where $k_{pro}$ is a constant spring stiffness. When a leg is touching the ground, its movement and that of the body are coupled. We have

$$\ddot{x} = \sum_{s_i = r} \left[ k_{ret}(l_i - l_0)\cos\theta_i - b\dot{l} \right],$$

and the length of leg $i$ is just a function of the position of the body:

$$l_i = \sqrt{w^2 + (x_{toe,i} - x)^2}\ .$$

Here, $k_{ret}$ is the spring stiffness and $b$ is a damping constant.

## The Bipedal SLIP: A More Complex Model

The bipedal SLIP model of Figure 4, may be familiar in its resemblance to Raibert's two-legged hopping robot [Raibert, 1986]. However, with a view toward relevance to RHex, this model differs from Raibert's hoppers in that the legs must circulate backwards and return above the body toward their next touchdown.

Let *(x,y)* denote the position of the center of mass of the BSLIP robot, which we assume to have a mass of one. Let $\theta_1$ and $\theta_2$ denote the angles of the two legs, which we assume to be essentially massless. The legs may not exceed a length of $L$. We denote by $x_{toe}$ the horizontal position of the toe of the robot when it is in stance and by r, the distance $\| (x, y) - (x_{toe}, 0) \|$. A leg can either be in stance or in swing mode. In stance, it is not controlled directly but instead a slave to the body dynamics. Thus, if leg $i$ is in stance, then $\theta_i = \sin^{-1}\left[(x_{toe} - x)/r\right]$.

In flight, the dynamics of the center of mass are

$$\begin{pmatrix} \ddot{x} \\ \ddot{y} \end{pmatrix} = \begin{pmatrix} 0 \\ -g \end{pmatrix} \tag{A1}$$

and while in stance they are

$$\begin{pmatrix} \ddot{x} \\ \ddot{y} \end{pmatrix} = -\frac{DU(r) + b\dot{r}}{r} \begin{pmatrix} x_{toe} - x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ -g \end{pmatrix}. \tag{A2}$$

Here $U(r)$ is the spring potential law. For our simulations we chose $U(r)=k/r^3$ where $k$ is the stiffness of the spring, tunable online. In particular, we may change $k$ during stance to adjust the amount of thrust supplied by the leg. The constant $b$ determines the amount of damping in the spring and $g$ is the gravitational constant. We assume that the masses of the legs are much less than the mass of the robot so that their velocities in swing are directly controllable: $\dot{\boldsymbol{q}}_i = u_i$.

The discrete dynamics are as follows. The body's dynamics are ballistic (A1) and both legs swing according to their control inputs until one of the legs touches down, which happens when $y < L\cos(\boldsymbol{q}_i)$ for some $i$. Then the dynamics of the body are given by (A2) and the leg that touched down becomes a slave to the body dynamics. The leg that is still swinging is still controllable. When the leg reaches maximal compression, when $\dot{r} = 0$, then the spring stiffness $k$ may be adjusted. When the value of $r$ exceeds $L$, the leg lifts off again and the cycle repeats. In the version of the synchronization controller where $PEP$ can be altered, liftoff occurs when r exceeds $PEP(t)$, a control input.

Discussion of the details of the body control algorithm can be found in [Saranli, Schwind and Koditschek, 1999] Essentially, we choose the next touchdown angle to be some nominal value plus $k_v(v^*-v_{prev})$ where $k_v$ is a gain, v* is the desired horizontal velocity during flight corresponding to the nominal touchdown andgle and $v_{prev}$ is the horizontal velocity during the last flight mode. Similarly, we choose the spring stiffness $k$ to be some nominal value plus $k_h\ (h^*-h_{prev})$ where $h^*$ is the vertical hopping height corresponding to the nominal value and $h_{prev}$ is the hopping height of the last flight mode.

The leg control during swing is via an artificial spring. That is,

$$\dot{\boldsymbol{q}}_i = k_i(\boldsymbol{q}_{td} - \boldsymbol{q}_i)/(\| \boldsymbol{q}_{td} - \boldsymbol{q}_i \| + \boldsymbol{e})$$

where $\boldsymbol{q}_{td}$ is the desired touchdown angle of the leg and $k_i$ is a tunable gain. In the continuous control version of the synchronization problem, we have $k_i(t)$ as a control input, essentially affecting the velocity of the swing.

# Footnotes

[1] This models exactly the present day RHex architecture. From the view of biology, the approach can be justified as reflecting the (plausible) presumption that the settling time of the electrical-chemical energy components inherent in the more complex (second and higher order) models of the biological pattern generators should be much faster than those of the body mechanics. This first cut simplification significantly eases the translation of the Cruse reflex model into the language of coupled oscillators that we introduce in Section 3.2.3. That translation in place, a more physiologically sophisticated view of the biological pattern generator can be substituted in future studies as needed.

[2] To be precise, we simulate the legs using first order velocity dynamics when they are not in contact, modeling the common situation that they have negligible mass and encounter significant viscous damping. Adding second leg flight dynamics would have a negligible effect on the qualitative outcome but add significantly to the complexity of the simulation.

[3] Even though the Raibert-style architecture decouples the two controller degrees of freedom, the body degrees of freedom are intrinsically coupled through the dynamics of a spring loaded inverted pendulum. If space and the reader's patience allowed, we would substitute for each double circle in Figure 7 a copy of the block diagram itself to suggest the manner in which the Raibert controller cleverly coordinates the phases of the vertical and horizontal single leg degrees of freedom.

[4] We have conducted extensive simulations of the single legged bead on a rail, but have not thought the results worth reporting in the present paper since they are so reminiscent of the single vertical hopping leg studied in [Komsuoglu and Koditschek, 2000]. We believe but have not yet demonstrated that the stability proof for that case can be readily applied to the single legged bead on a rail as well. In the sense we have tried to depict in Figure 7, the single legged SLIP represents a parallel composition of this vertical hopping leg together with the single legged bead on a rail. The formal question of whether a CPG controller can stabilize this two degree of freedom composed mechanism (i.e., the single legged SLIP) is a matter of ongoing mathematical study.

# References


Berkemeier, M. D. and K. V. Desai. "Control of Hopping Height in Legged Robots Using a Neural-Mechanical Approach." Paper presented at the International Conference on Robotics and Automation, Detroit, Michigan, May 1999.

Buehler, M., D. E. Koditschek, and P. J. Kindlmann. "Planning and control of robotic juggling and catching tasks." *International Journal of Robotics Research* 13, no. 2 (1994): 101-18.

Buehler, M., U. Saranli, and D. E. Koditschek. "Design, Modeling and Preliminary Control of a Compliant Hexapod Robot." Paper presented at the International Conference on Robotics and Automation, San Francisco, CA, April 2000.

Calvitti, A, and Beer R. D. "Analysis of a distributed model of leg coordination. I. Individual coordination mechanisms." *Biological Cybernetics* 82, no. 3 (2000): 197-206.

Chiel, H. J., R. D. Beer, R. D. Quinn, and K. S. Espenschied. "Robustness of a distributed neural network controller for locomotion in a hexapod robot." *IEEE Transactions on Robotics and Automation* 8, no. 3 (1992): 293-303.

Cohen, A. H., and D. L. Boothe. "Sensorimotor interactions: Principles derived from central pattern generators." In *Handbook of Brain Theory and Neural Networks,* edited by Arbib, M., Boston: MIT Press, 1995.

Cohen, A. H., G. B. Ermentrout, T. Kiemel, N. Kopell, K. Sigvardt, and T. Williams. "Modelling of intersegmental coordination in the lamprey central pattern or for locomotion." *Trends in Neurosciences* 15 (1992): 434-438.

Cohen, A. H., P. J. Holmes, and R. H. Rand. "The nature of the coupling between segmental oscillators of the lamprey spinal generator for locomotion: A mathematical model." *J. Math. Biol.* 13, no. 3 (1981/82): 345-369.

Cohen, A. H., S. Rossignol, and S. Grillner. *Neural control of rhythmic movements in vertebrates.* New York: Wiley, 1988.

Cruse, H. "What mechanisms coordinate leg movement in walking arthropods?" *Trends in Neurosciences* 13, no. 1 (1990): 15-21.

Cruse, H., T. Kindermann, M. Schumm, J. Dean, and J. Schmitz. "Walknet--a biologically ispired network to control six-legged walking." *Neural Networks* 11 (1998): 1435-1447.

Delcomyn, F. "Walking and running." In *Comprehensive Insect Physiology, Biochemistry, and Pharmacology. Nervous System: Structure and Motor Function*, 439-466. New York: Pergamon Press, 1985.

Full, R., U. Saranli, M. Buehler, Brown. B., N. Moore, D. Koditschek, and H. Komsuoglu. "Evidence for Spring Loaded Inverted Pendulum Running in a Hexapod Robot." Paper presented at the International Symposium on Experimental Robotics, Honolulu, Hawaii, 2000.

Full, R.J., and C.T. Farley. "Musculoskeletal Dynamics in Rhythmic Systems - A Comparative Approach to Legged Locomotion" in *Biomechanics and Neural Control of Movement*. Edited by J.M. Winters and P.E. Crago. New York: Springer-Verlag, 1999.

Full, R. J., and D. E. Koditschek. “Templates and anchors: Neuromechanical hypotheses of legged locomotion on land.” *Journal of Experimental Biology* 202, no. 23 (1999): 3325-3332.

Full, R. J., and M. S. Tu. “Mechanics of six-legged runners.” *Journal of Experimental Biology* 148 (1990): 129-146.

Grillner, S. “Neurobiological bases of rhythmic motor acts in vertebrates.” *Science*, Apr 12 1985, 143-149.

Holmes, P., and D. A. Taylor. “Simple models for excitable and bursting neural networks.” *J. Math Biol*. (1997).

Klavins, E., R. Ghrist, and D. E. Koditschek. “Toward the Regulation and Composition of Cyclic Behaviors.” Paper presented at the Proceedings of 4th Int. Workshop on Algorithmic Foundations of Robotics, Dartmouth, NH 2000.

Koditschek, D. Computational Neuromechanics: *Programming Work in Biological and Machine Systems University of Michigan*, Aug 1999. Available at http://ai.eecs.umich.edu/CNM/.

Komsuoglu, H., and D. E. Koditschek. “Preliminary Analysis of a Biologically Inspired 1-DOF 'Clock' stabilized Hopper.” Paper presented at the World Multiconference on Systemics, Cybernetics and Informatics, Orlando 2000.

N. Kopell, “Chains of Coupled Oscillators”, in *Handbook of Brain Theory and Neural Networks*, M. Arbib, ed., MIT Press, Cambridge (1995) 178-183.

Kopell, N. “We got rhythm: Dynamical Systems of the Nervous System.” *Notices of the AMS*, January 2000, 6-16.

Kuo, A. D. “The relative roles of feedforward and feedback in control of rhytmic movements.” In preparation.

Leob, and Brown. “A reductionist approach to creating and using neuromusculoskeletal models.” In *Biomechanics and Neural Control of Posture and Movement*, edited by J. M. Winters and P. E. Crago. New York: Springer-Verlang, 2000.

Lewis, M. A., R. Etienne-Cummings, A. H. Cohen, and M. Hartmann. “Toward Biomorphic Control using Custom a VLSI Chips.” Paper presented at the International Conference on Robotics and Automation, San Francisco, April 2000.

Nakanishi, J., T. Fukuda, and D. E. Koditschek. “A brachiating robot controller.” *IEEE Transactions on Robotics and Automation* 16, no. 2 (2000): 109-123.

Orlovsky, G. N., T. G. Deliagina, and S. Grillner. Neural Control of Locomotion. New York: Oxford University Press, 1999.

Pearson, K. “The Control of Walking.” *Scientific American* 1976, 72-86.

Pratt, J., P. Dilworth, and G. Pratt. “Virtual model control of a bipedal walking robot.” Presented at the International Conference on Robotics and Automation (1997): 193-8.

Raibert, M. H. *Legged Robots That Balance*, MIT Press series in artificial intelligence. Cambridge, Mass.: MIT Press, 1986.

Raibert, M. H., H. B. Brown, and M. Chepponis. “Experiments in balance with a 3D one-legged hopping machine.” *International Journal of Robotics Research* 3 (1984): 75-92.

Ringrose, R. “Self-Stabilizing Running.” Paper presented at the International Conference on Robotics and Automation, Albuquerque, New Mexico 1997.

Ritzmann, R. E., R. D. Quinn, J. T. Watson and S. N. Zil. “Insect walking and biorobotics: A relationship with mutual benefits.” *BioScience* 50, no. 1 (2000): 23-33.

Rizzi, A. A., and D. E. Koditschek. “An active visual estimator for dexterous manipulation.” *IEEE Transactions on Robotics and Automation* 12, no. 5 (1996): 697-713.

Rizzi, A. A., L. L. Whitcomb, and D. E. Koditschek. “Distributed real-time control of a spatial robot juggler.” *Computer* 1992, 12-24.

Saranli, U., W. J. Schwind and D. E. Koditschek. “Toward the Control of a Multi-Jointed, Monoped Runner.” Paper presented at the International Conference on Robotics and Automation, Leuven, Belgium 1998.

Kubow, T., W. Schwind, R. Full, M. Garcia and D.E. Koditschek. “Posture Selection in a Running Cockroach.” In preparation.

Swanson, P.J, R.R. Burridge and D.E. Koditschek, "Global Asymptotic Stability of a Passive Juggling Strategy: A Possible Parts Feeding Method", Mathematical Problems in Engineering, 3(1), 1995.

Vogel, S. *Cats' Paws and Catapults: Mechanical Worlds of Nature and People*. New York: Norton, 1998.

Wilson, D. M. “Insect Walking.” In *Annual Review Entomology*, 103-122, 1966.

# Illustrations

**Figure 1:** The RHex robot running at a body length (~50cm) per second over rough terrain.

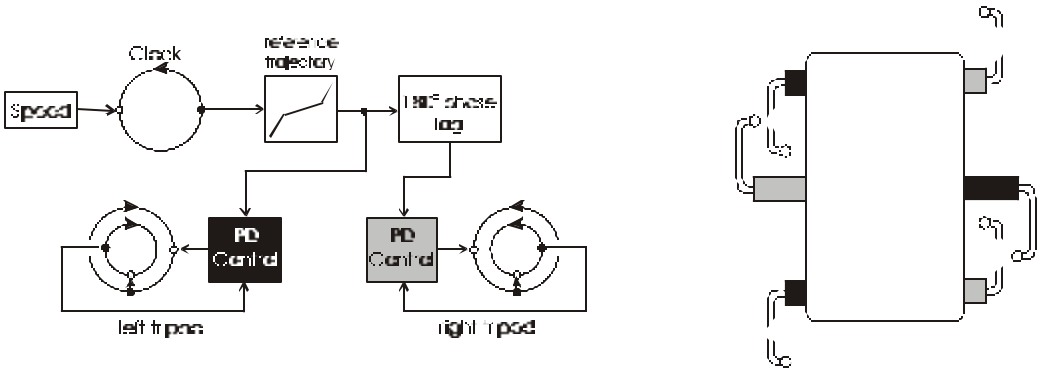


(a) (b)

**Figure 2:** (a) Identical copies of the phase and anti-phase version of the two stroke clock signal are used as reference trajectories by decoupled PD controllers at each of the hip motors comprising the respective tripods. (b) A schematic top view of RHex showing in black the left tripod and in gray, the right tripod.

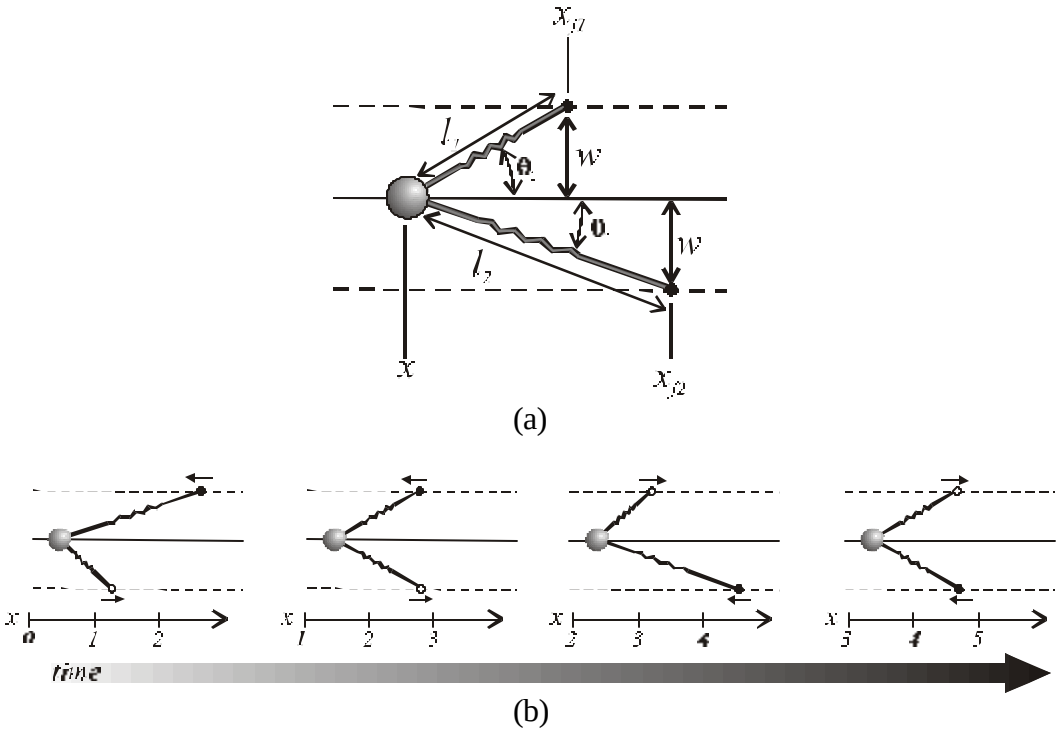


**Figure 3:** The Bipedal Bead on a Rail (BBR) model. (a) A bead is constrained to move on a rail (central horizontal line in the figure). The legs may grab the parallel outer rails to pull the body forward, which they do via their springs. Specifically, during protraction, a leg is not grasping the outer rail and is extending its spring well beyond its rest length. When it grasps the rail in retraction, the spring pulls the body toward the toe position, attempting to restore itself to the rest length. See the appendix for details of the BBR model. (b) A sequence showing the legs alternating between protraction and retraction. Open indicate the leg is not grasping the rail and closed indicate that it is.

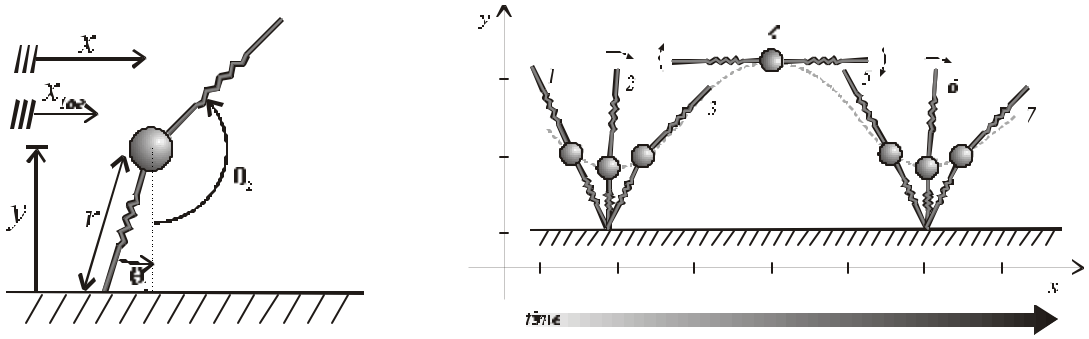


(a) (b)

**Figure 4:** The Bipedal Spring Loaded Inverted Pendulum (BSLIP) model. (a) Two spring-loaded legs alternatively support the body and swing over the top of it. (b) A sequence showing the legs alternating in stance. Images 1 and 5 correspond to the touchdown of a leg, images 2 and 6 to maximal compression, and 3 and 7 to liftoff. Image 4 corresponds to the apex of the flight phase when both legs are in swing mode.

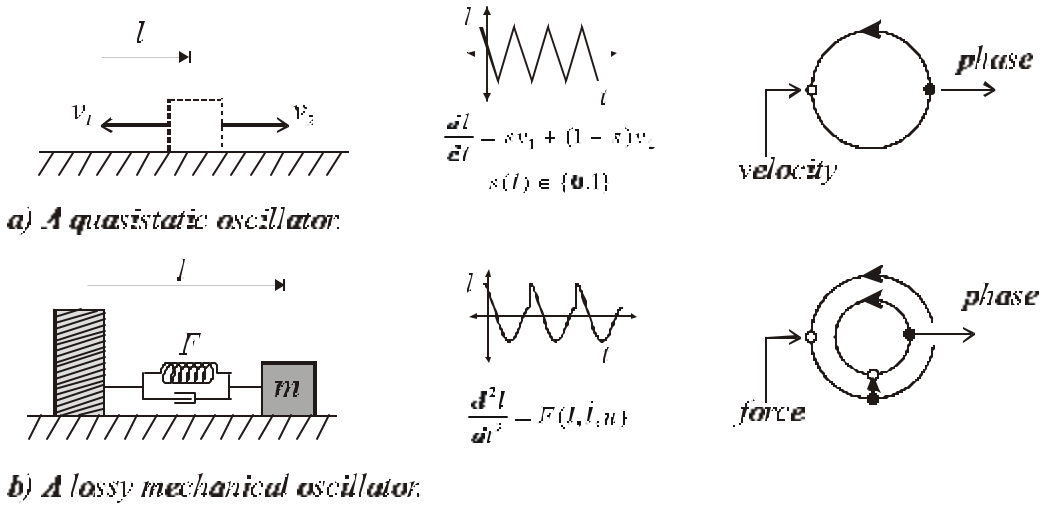


**Figure 5:** Two different types of oscillator with their corresponding equations and icons. (a) A First order system cannot oscillate without some switching controller. Because velocity is the control input to the system, we denote a first order oscillator by a single circle. (b) A second order oscillator has a phase mechanically determined by its energy. If it has damping, periodic control inputs (underscored in the figure by the discontinuities in the mass's trajectory) are needed to reintroduce energy. Because the control input to the system is force, which only indirectly affects phase, we denote a second order oscillator by two circles.

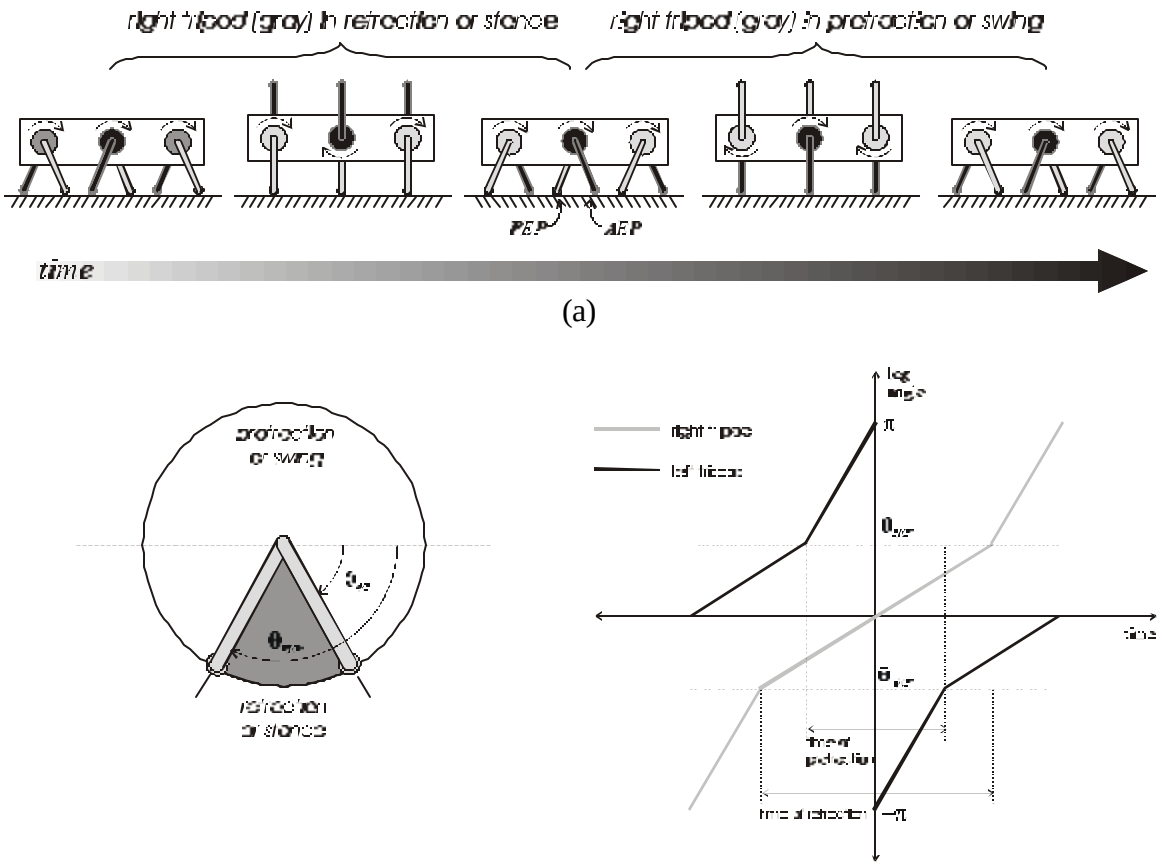


**Figure 6:** CPG control of RHex. (a) An animation sequence showing a full stride for the right tripod, shown in gray, of RHex. The counter motion of the left tripod is shown in black. (b) RHex's "clock" — the periodic feedforward reference signal presently used to drive an alternating tripod gait. Between the angles $q_{AEP}$ and $q_{PEP}$ the leg is in stance or retraction. The rest of the time it is in swing, or protraction. (c) The speed profile for the two tripods. During retraction, the leg moves slowly. While in protraction, it moves quickly, sweeping out a greater angle in the same amount of time.

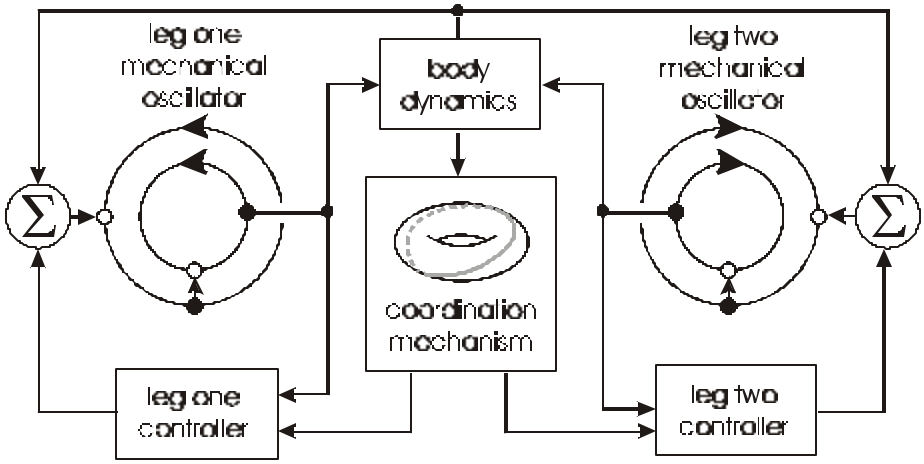


**Figure 7:** Block diagram suggesting the dynamics of the BBR model introduced in Figure 3 (see the Appendix for mathematical details). Each leg consists of an angular degree of freedom that periodically forces the body upon making ground contact, yielding, at each touchdown a second order mechanical oscillator. The resulting hybrid body dynamics (i.e., the evolution of the 1 degree of freedom center of mass position and velocity) together with the leg states determine whether one, two, or no legs are in contact. The coordination mechanism is allowed to set the posterior extreme position, *PEP*, of a leg. Reflex mechanisms adjust this input based on the state of the body and each leg; CPG mechanisms simply select the input in an unvaryingly periodic manner.

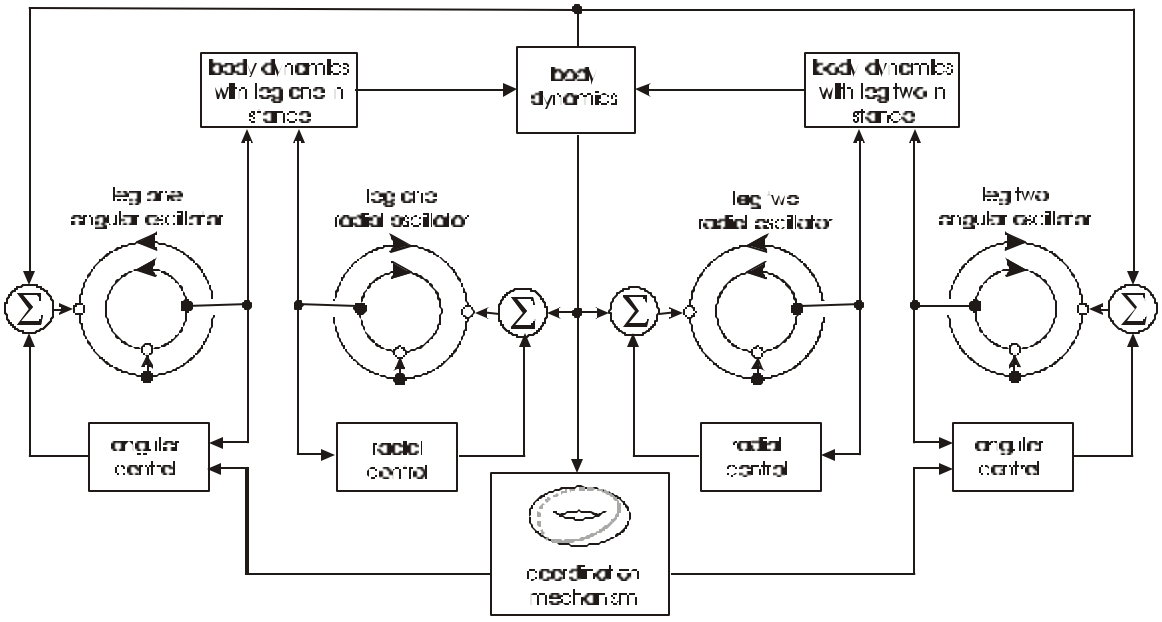


**Figure 8:** Block diagram suggesting the dynamics of the BSLIP model introduced in Figure 4 (see the Appendix for mathematical details). Each leg consists of a radial and an angular degree of freedom that periodically force the body upon making ground contact, yielding, each at touchdown a parallel composition of two second order mechanical oscillators. The resulting hybrid body dynamics (i.e., the evolution of the two degree of freedom center of mass position and velocity) together with the leg states determine whether one, two, or no legs are in contact. The coordination mechanism is allowed to set either the posterior extreme position, *PEP*, or the swing velocity of the leg. Reflex mechanisms adjust this input based on the state of the body and each leg; clock mechanisms simply select the input in an unvaryingly periodic manner.

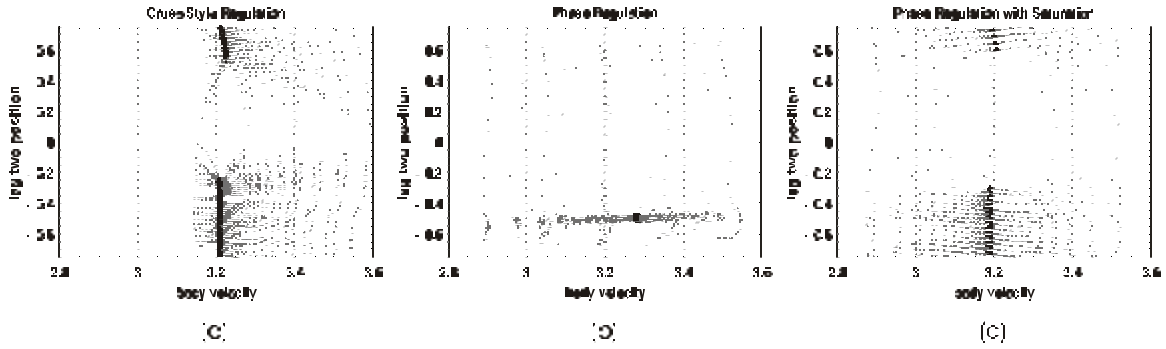


**Figure 9:** Comparing the "return maps" generated by two styles of reflex coordination mechanism applied to the BBR model. Each figure exhibits the same BBR "Poincare section" wherein the horizontal axis represents body velocity at the instant leg one touches down while the vertical axis represents the position of leg two at that same instant. The dots represent samples from several hundred runs, each run having initiated at a distinct different state and sampled forward in time until steady state (depicted by a darker dot) is achieved. (a) Quasi-static style reflex coordination achieves a neutrally stable region (in dark dots). (b) Coupled Oscillators style reflex coordination stabilizes at an exact, alternating gait – in other words, the stable region is a point. (c) The Coupled Oscillators style reflex coordination mechanism with a saturated energy function reproduces the effect in (a).

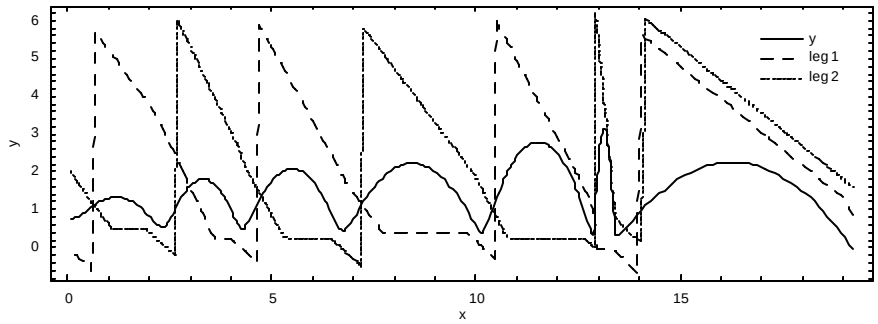


(a)

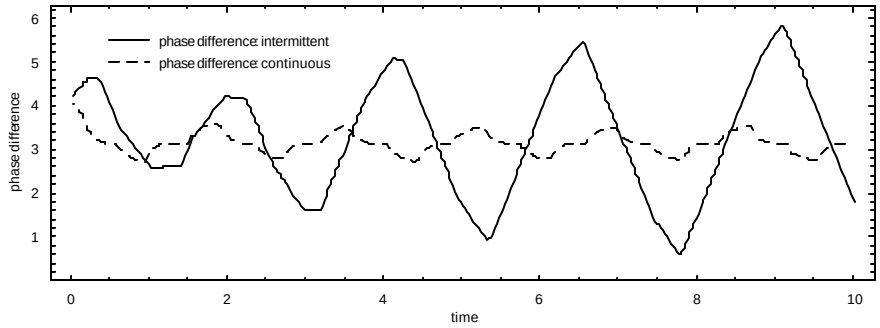


(b)

**Figure 10:** Two characteristic failure modes of the Quasi-static style reflex control mechanism applied to the BSLIP model. (a) Using a higher synchronization gain and low swing velocity. The faltering trajectory of the center of mass height and the increasingly uncoordinated leg angles are plotted as functions of center of mass horizontal position. (b) Using a lower synchronization gain and higher swing velocity. The phase difference between the two legs, plotted as a function of time (solid line), fluctuates with increasing magnitude as each leg "waits" a longer time to touchdown (because the swing velocity is set high) until, eventually, the body falls to the ground. Compare this to the neutrally stable phase difference trajectory (dashed line) resulting from the Coupled Oscillator style reflex control mechanism.

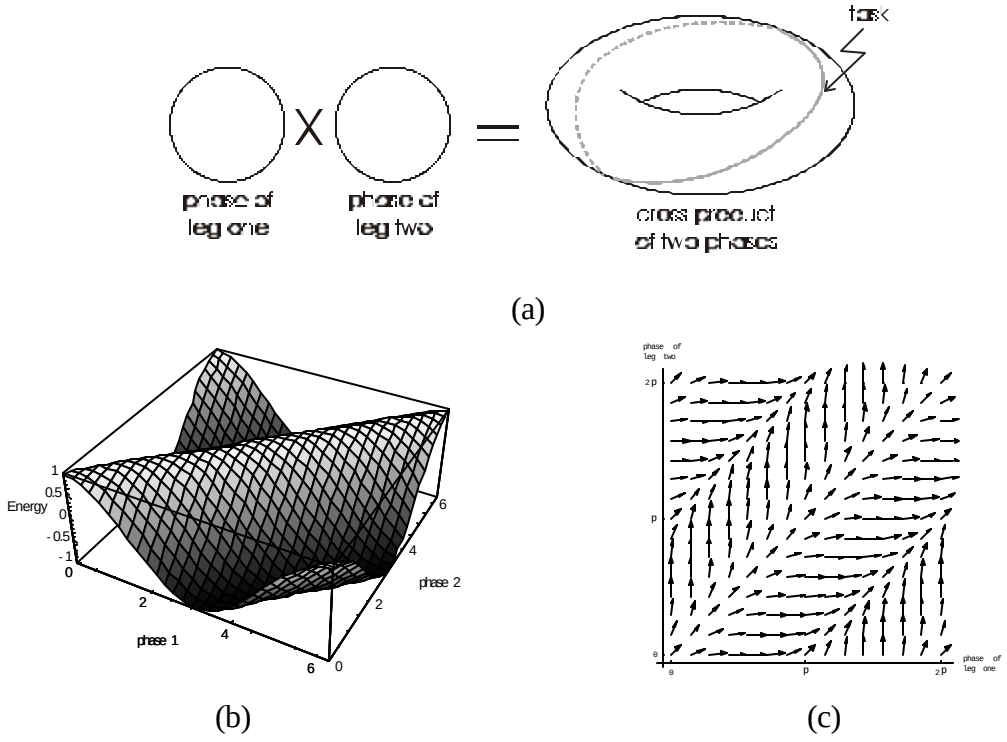


**Figure 11**: The Coupled Oscillator style reflex coordination mechanism results from correcting a measured pair of phase velocities at a measured pair of phases toward a desired pair at that measured phase. The desired pair of phase velocities is specified by a reference phase dynamical system constructed as follows. (a) The set of phase angle pairs (the "cross product" of two cyclic phase spaces) is a torus. The task of cycling out of phase is encoded as a circle embedded in the torus along with a desired instantaneous phase velocity for each point along the circle. (b) A "pseudo-energy" function $V(\boldsymbol{j}_1, \boldsymbol{j}_2) = \cos(\boldsymbol{j}_1 - \boldsymbol{j}_2)$ yields a reference field derived by taking the gradient (directional derivative) at each phase pair. The desired phase difference cycle lies along the "low energy trough." (c) The reference field $R(\boldsymbol{j}_1, \boldsymbol{j}_2) = k_1 - k_2 \nabla V(\boldsymbol{j}_1, \boldsymbol{j}_2)$ derived from *V*. The legs are driven (at a rate proportional by constant $k_2$ to the gradient magnitude) out of phase down onto the "trough" along which the phase velocities are pulled at the constant rate $k_1$.

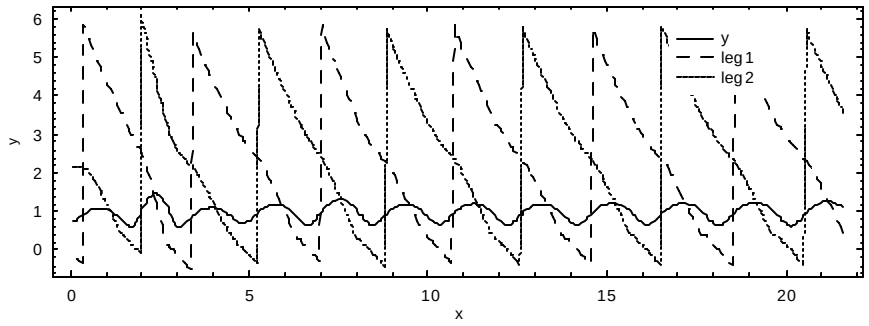


**Figure 12:** A simulation of the Coupled Oscillator style reflex coordination mechanism applied to the BSLIP model using the "pseudo-energy" function described in Figure 11. After an initial transient period, the system settles down into an alternating gait and the characteristic body motion of a spring loaded inverted pendulum. See Figure 10(a) for an explanation of the axes.

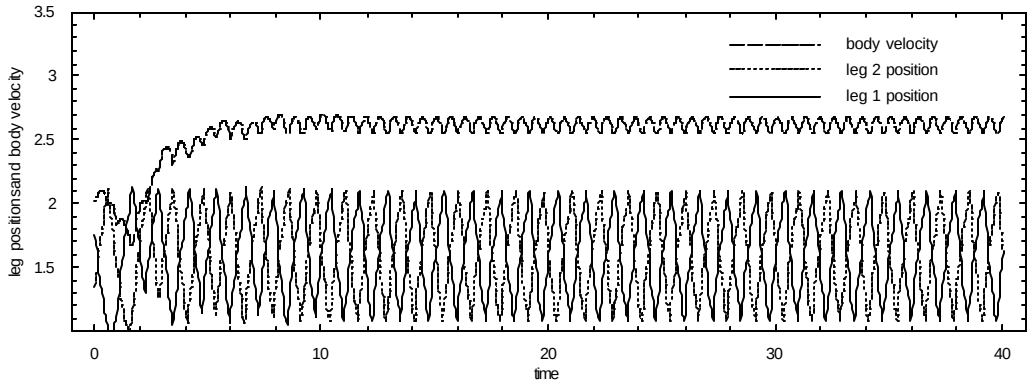


(a)

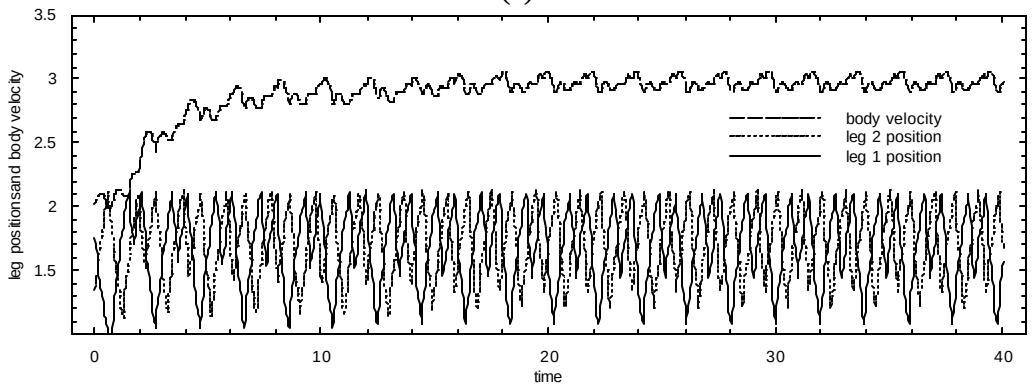


(b)

**Figure 13:** Simulations of the BBR model using a CPG for coordination as described in the text. The horizontal axis is time, the vertical axis contains the body's forward velocity and the positions of the two legs. (a) The CPG period is tuned to the natural frequency of the system as defined by the stiffness of the leg springs. Notice that the gait and the forward velocity stabilize. (b) The CPG is too fast: the legs cannot realize the signals it sends, resulting in an undesired, "limping" gait.

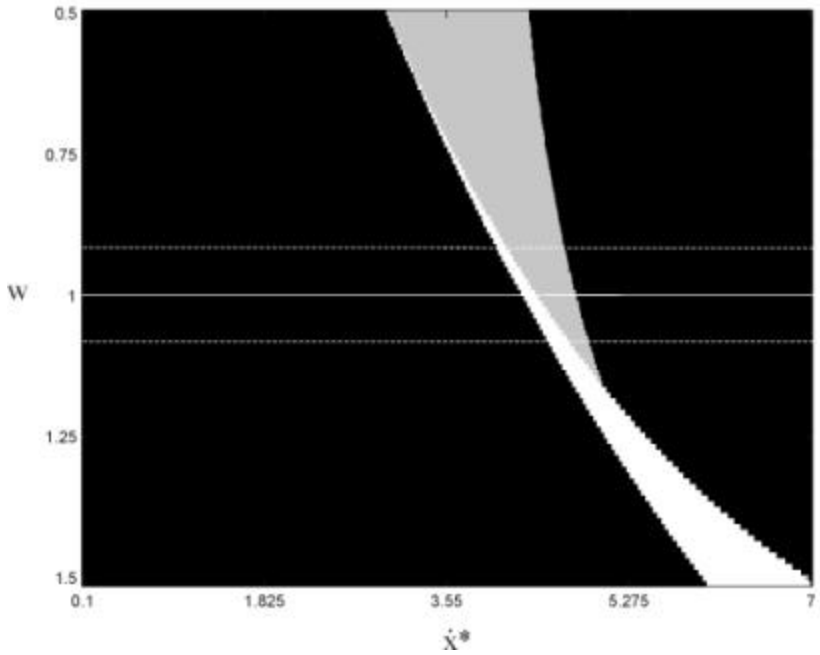


**Figure 14**: The relationship between the steady state behavior of the CPG coordinated BBR model and the environment into which it operates. The abscissarepresents the steady state velocity of the body, $\dot{x}^*$,sampled at the AEP event of a specified leg. The ordinate represents the distance, *w*, from the central rail to the outer guide-rails. The black regions indicate that there is no fixed point for the (*w*,*x**) pair. Gray regions indicate that there is a fixed point that is unstable. And, finally, the white region represents the stable fixed points.

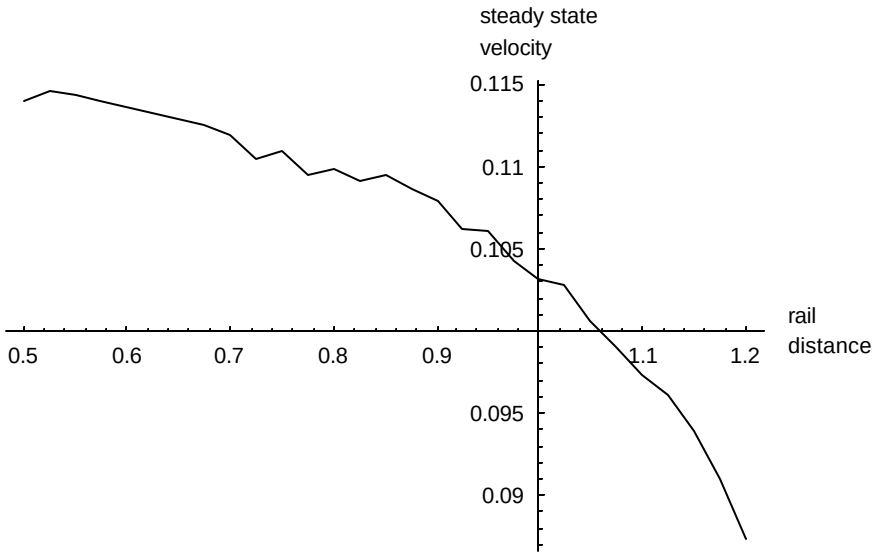


**Figure 15**: The function relating steady state behavior of the Reflex coordinated BBR model to the environment into which it operates. The abscissa and ordinate are the reverse from Figure 14. As expected, feedback confers significant robustness against environmental variation: there is a unique asymptotically stable body velocity for each environmental condition.