

Jorge G. Cham

Jonathan K. Karpick

Mark R. Cutkosky

Center for Design Research
Stanford University
Stanford, CA 94305-2232, USA
contact: jgcham@cdr.stanford.edu

Stride Period Adaptation for a Biomimetic Running Hexapod

Abstract

We demonstrate an adaptation strategy for adjusting the stride period in a hexapedal running robot. The robot is inspired by discoveries about the self-stabilizing properties of insects and uses a sprawled posture, a bouncing alternating-tripod gait, and passive compliance and damping in the limbs to achieve fast (over 4 body-lengths per second), stable locomotion. The robot is controlled by an open-loop motor pattern that activates the legs at fixed intervals. For maximum speed and efficiency, the stride period of the pattern should be adjusted to match changes in terrain (e.g. slopes) or loading conditions (e.g. carrying an object). An ideal adaptation strategy will complement the design philosophy behind the robot and take advantage of the self-stabilizing role of the mechanical system. In this paper we describe an adaptation scheme based on measurements of ground contact timing obtained from binary sensors on the robot's feet. We discuss the motivation for the approach, putting it in the context of previous research on the dynamic properties of running machines and bouncing multi-legged animals, and show results of experiments.

1. Introduction

Resonance Tuning

We have built a family of hexapedal robots that are inspired by discoveries concerning the locomotion of insects and, in particular, of the cockroach. These animals run rapidly (between 10 and 50 body-lengths/second depending on the species) and over rough terrain using a combination of open-loop muscle activation patterns and "preflexes," that is, passive mechanisms that stabilize the animals' motion in response to perturbations (Ahn and Full, 1997; Full et al., 1998; Kubow and Full, 1999; Meijer and Full, in press). Like the insects that inspired them, the robots employ passive mechanical properties that enable them to run quickly

(over 4 body-lengths per second) and over hip-height obstacles (see Fig. 1) without closed-loop control (Cham et al., 2002; see also multimedia Extensions 1 and 2). Although this approach works well in the laboratory, there is a question about its versatility. How effectively can a particular open-loop control and set of mechanical properties function over a range of conditions that may include variations in ground slope and hardness and changes in loading? Furthermore, the animals or robots themselves may change over time. A limb may become damaged or the mechanical properties may vary with temperature. A way to address this problem is adaptation, in which the parameters of the open-loop control are automatically tuned to optimize performance as conditions change.

Figure 2 illustrates an approach in which adaptation is combined with reflexes for stable running. An open-loop, feed-forward, motor controller generates the pattern of actuator commands to achieve a steady alternating-tripod gait. The kinematic arrangement and passive compliance and damping of the limbs achieve the locomotion and provide stable response to perturbations. Sensory information is used at a slower rate to adapt, or tune, the motor pattern in response to changing conditions. In running insects, an important reason for relying on reflexes in combination with slow adaptation is that neural conduction speeds are too slow for feedback control to act effectively within each stride period. In robots, of course, the same limitation does not necessarily apply. However, for small and inexpensive robots like our prototype, "Sprawlita," the same approach allows the use of simple sensors without concerns that actuator delays, sensor noise or even failures will jeopardize short-term performance. This is an important consideration because many sensors become noisy when mounted on a small hexapod running at 7-10 Hz.

This approach has been hypothesized as the basis for the generation of rhythmic movements such as locomotion in

animals (Orlovsky, et al., 1999; Full and Koditshek, 2000). The presence of neural circuits called Central Pattern Generators (CPG) that can generate, in the absence of feedback, efferent motor patterns similar to those seen during locomotion is well established in both vertebrates and invertebrates. During normal movements, however, it has been found that sensory feedback from receptors in the skin, joints and muscles plays an important role in modulating the frequency and amplitude of the output of these pattern generators (Rosignol, et al., 1988; Orlovsky, et al., 1999). Moreover, it has been hypothesized that this interaction between sensory feedback and the pattern generator is designed to exploit the dynamics of the physical system by tuning the motor commands to the system's resonant frequency (Hatsopoulos, 1996), thereby maximizing performance while minimizing the amount of work utilized. The exact manner in which sensory feedback might be used to accomplish this "resonance tuning," however, is still unknown.

Adaptation for a Biomimetic Running Hexapod

The basic design of the Sprawl family of hexapods consists of a body and legs built up in layers using a rapid-prototyping process called Shape Deposition Manufacturing. The design and manufacturing are detailed in Cham et al. (2002) and in Bailey et al. (1999). Each leg has two degrees of freedom but only the thrust direction is actuated, using pneumatic pistons embedded in the legs. When running, hip rotations are passive and are accomplished by flexures of visco-elastic material. This design is inspired by the trochanter-femur joint in cockroaches, which is believed to be mostly passive in the sagittal plane. A servomotor at each hip is used only to establish the equilibrium position of the




Figure 1. The hexapodal robot has a body and legs fabricated by Shape Deposition Manufacturing (Cham et al., 2002) and features embedded actuators and compliant legs. Here the robot is crossing a hip-height obstacle without using sensory feedback and without significantly slowing down or being knocked off course.

hip joint. Binary contact sensors are attached to the feet. Pneumatic solenoid valves regulate air-flow into the leg pistons from a high-pressure source. The original Sprawlita design uses two valves, one for each tripod of legs, embedded in the body of the robot. These on/off valves are activated according to an open-loop binary motor pattern. A newer design has a valve embedded in each leg, which results in faster actuator dynamics and more control over the timing of the thrust force at each foot. As will be seen in the following sections, these are important considerations. Depending on configuration, the robots weigh between 0.25-0.33 Kg and have a length of 10-15cm. Maximum speeds range from 0.5-0.8m/sec with preferred stride frequencies of 7-10 Hz.

The operational parameters that can be varied are the stride period (length of time between activation of each tripod) and the duty cycle (length of time that the valves are kept open during each stride) of the motor pattern and the equilibrium positions of the compliant hip joints. All of these parameters contribute to running performance and could be subject to adaptation. In this paper, we focus on stride period and duty cycle.

Figure 3 shows the robot's speed as a function of ground slope for two different stride frequencies, and illustrates the importance of adjusting the stride frequency to changing conditions. On level ground the fastest locomotion is

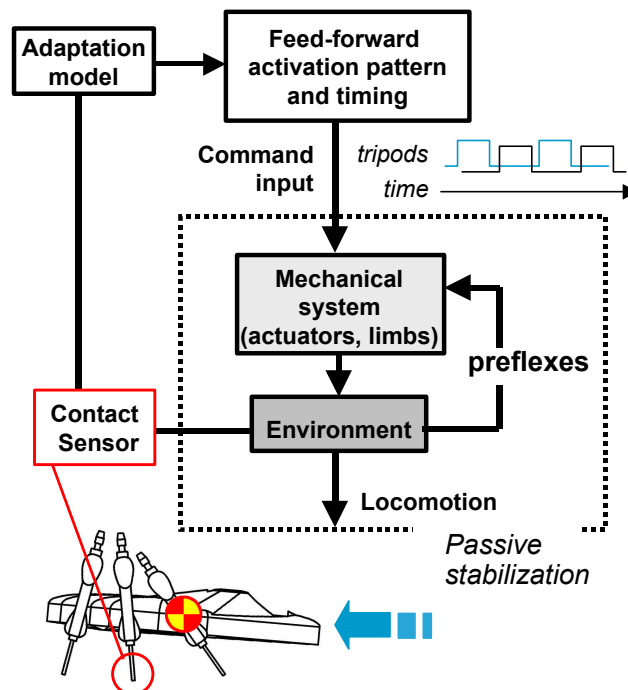


Figure 2. A combination of stabilizing passive mechanisms, or "preflexes," and sensor-based adaptation of an open-loop feed-forward controller provides insects and small robots with a robust, stable and versatile approach to running over rough terrain.

obtained with a frequency of approximately 10 Hz. But on a 20 degree slope, locomotion is considerably faster at 5 Hz than at 10 Hz. The optimal frequency also varies somewhat from one robot to the next due to manufacturing tolerances and variations in the materials properties of the legs. Consequently, there is a motivation to make the robots "self-tuning" over a range of operating conditions. Ideally, we would like an adaptation strategy that does not require adding expensive or complicated sensors to the robot.

For these reasons, we examined the relationship between ground contact times (obtained from binary sensors in the robots' feet) and the timing parameters (frequency and duty cycle) of the open-loop motor control. As will be discussed in Section 3, of the various timing quantities that we can examine, the interval between end of thrust (closing of the pneumatic valve) and lift-off of the feet is a good indicator for adjusting the stride frequency. To better understand why this approach works, we first examine a simplified one-legged vertical hopper model. Although the six-legged robots are considerably more complex, their motion in response to the firing of each tripod is qualitatively similar to that of a single mass and foot.

2. Simplified Model for Open-Loop Locomotion and Adaptation

To understand how monitoring the ground contact time can provide information to the robot about the effectiveness of its current motor pattern in running, we start with a simple vertical hopping model that has been frequently utilized in the literature. This model consists of a point mass constrained to move in the vertical direction with a telescoping leg that comes into intermittent contact with the ground. Compliance and damping as well as the ability to generate thrust are commonly modeled in the leg. Although this model cannot tell us about the coupling between vertical and

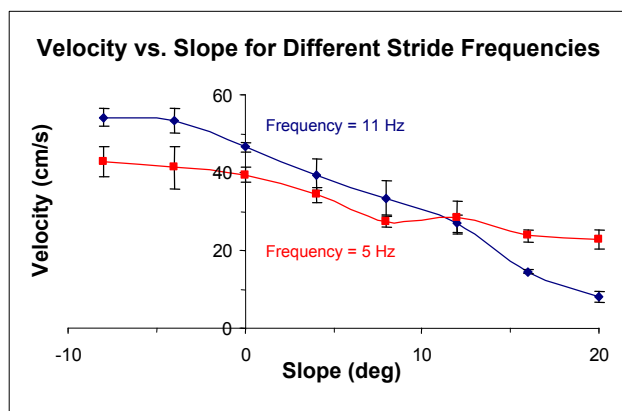


Figure 3. Robot ground speed versus terrain slope for two different stride frequencies. As shown, the optimal stride frequency for maximum speed depends on the slope, which illustrates the need for adaptation.

horizontal motion, an important factor in the dynamics of the Sprawl robot family, it does shed light on the relationship between system energy and actuator timing. It provides insight into the circumstances under which a stable, steady-state hopping cycle is achieved with an open-loop control scheme and it helps us determine how to get the most work out of our actuators using only simple sensors.

Variations on this basic model have been examined by several investigators. Despite its apparent simplicity, the one-legged vertical hopper exhibits a rich set of dynamic behaviors including stable and unstable periodic motion. Raibert (1986), Koditschek and Buehler (1991), and Vakakis, et al. (1991) analyze one-legged vertical hopping models in which thrust is activated in a closed-loop manner when the leg reaches its maximum compression. Stable open-loop models of vertical hopping were demonstrated and analyzed by Ringrose (1997), Berkemeier and Desai (1998) and Komsuoglu and Koditschek (2000). The effects of varying the stride period of the open-loop motor control pattern were considered by Ringrose (1997) and Berkemeier and Desai (2000). Both observe that maximum hopping height is achieved when thrust is initiated when the leg reaches maximum compression, and that increasing the stride period further results in unstable hopping. Using this result, Berkemeier and Desai (2000) proposed an adaptation law that senses the velocity at thrust activation and increases the stride period until this velocity is zero, which occurs at the leg's maximum compression.

The physical implementation of the Sprawl robots requires that minimal damping assumptions made in the previous models mentioned be reconsidered. In insects, and in the Sprawl family of robots, viscoelastic materials dissipate substantial amounts of energy per cycle. Typical dimension-

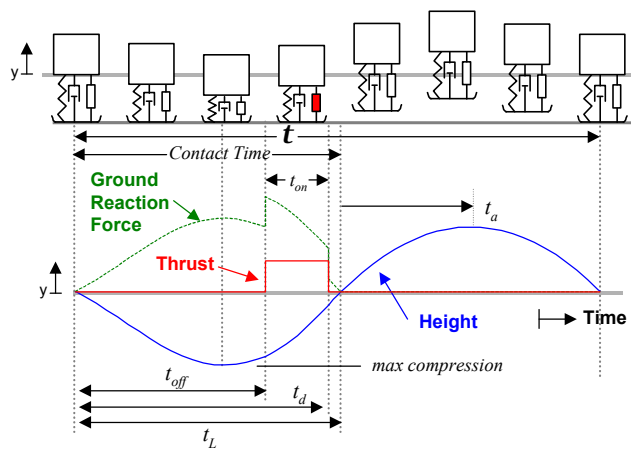


Figure 4. Time history of a single degree-of-freedom vertical hopper. The mass is attached to a massless leg with stiffness k and damping b . After some time t_{off} a thruster in parallel applies a thrust force, f . At some time t_L , the mass lifts off the ground and travels ballistically in the air.

less damping ratios are on the order of $\zeta = 0.3$ (Garcia et al., 2000). Additionally, a model was needed to determine the practicality of using a simple binary switch for feedback instead of more complicated and noise-susceptible velocity sensors.

We therefore consider a single-legged vertical hopper that includes substantial viscous damping. Figure 4 depicts a sample time history of this hopping model. The massless leg has stiffness, k , damping, b , and an actuator that is able to provide a thrust force $f(t)$ that is initiated after some time t_{off} and terminated after a fixed duration, t_{on} , or by lift-off, whichever occurs first. The stride period begins at $t=0$ as the robot touches down, which occurs when $y=0$, the spring's neutral length. During the ground contact phase, the ground reaction force is given by:

$$GRF = ky + b\dot{y} - f(t) \quad (1)$$

The equation of motion for the mass is:

$$m\ddot{y} = -ky - b\dot{y} - g + f(t) \quad (2)$$

Lift-off occurs when the ground reaction force is equal to zero, and the hopper transitions to an airborne phase, where

it travels ballistically. In the next section, we consider the conditions for optimal hopping height of this model.

2.1 Optimal hopping height

In steady-state, the landing velocity of one cycle must equal the landing velocity of the previous cycle. Simulations and analytical and numerical calculations, detailed in Appendix A, were performed to determine which values of force application delay, t_{off} , and force duration, t_{on} , meet steady state conditions for given values of m , k , b , g , and force f . Each solution pair (t_{off}, t_{on}) corresponds to a steady-state hopping cycle with a given steady-state hopping height. Figure 5 shows the steady-state hopping height of these solutions plotted against the velocity at which thrust is initiated for $m=g=k=1$, $b=0.2$, and $f=4$ and for a range of thrust durations. The velocity at thrust activation on the horizontal axis of the figure is normalized by the velocity at take-off for each steady-state solution. This normalized velocity is zero if thrust is initiated at maximum compression, approaches -1 if thrust is initiated at landing, and approaches +1 if thrust is initiated near take-off. Each line represents the set of solutions for a given thrust duration t_{on} , here specified as a percentage of the natural period $(2\pi)/\omega_n$.

Figure 5 shows that for thrust durations greater than 20% of the natural period, peak hopping height is achieved when the thrust is applied near the point of maximum spring compression, that is, when the velocity at activation is nearly zero. For shorter thrust durations, however, optimal steady-state heights occur when thrust is initiated *after* maximum spring compression (velocity at force application is positive).

In evaluating the conditions that determined the maximum hopping heights in Figure 5, it is seen that, for a given force level and duration, hopping height is maximized at the steady-state solution in which the net positive work performed by the actuator within a stride is maximized:

$$\int_{t_{off}}^{(t_{off}+t_{on})} f(t) \cdot \dot{y}(t) dt = WorkInput \quad (3)$$

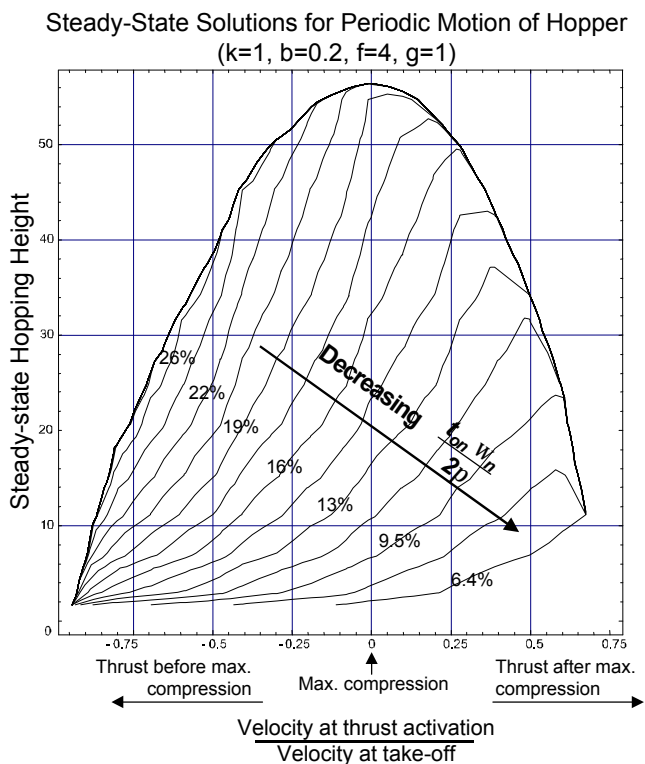


Figure 5. Hopping height as a function of normalized velocity at thrust application for the one-DOF hopper. Each line represents the solutions to the steady-state constraint equations for a given set of operating parameters and thrust duration.




Figure 6. Work performed by the actuator depends on the timing of thrust activation and deactivation. In the case of a constant thrust force, the work performed is proportional to the net change in leg length during thrust.

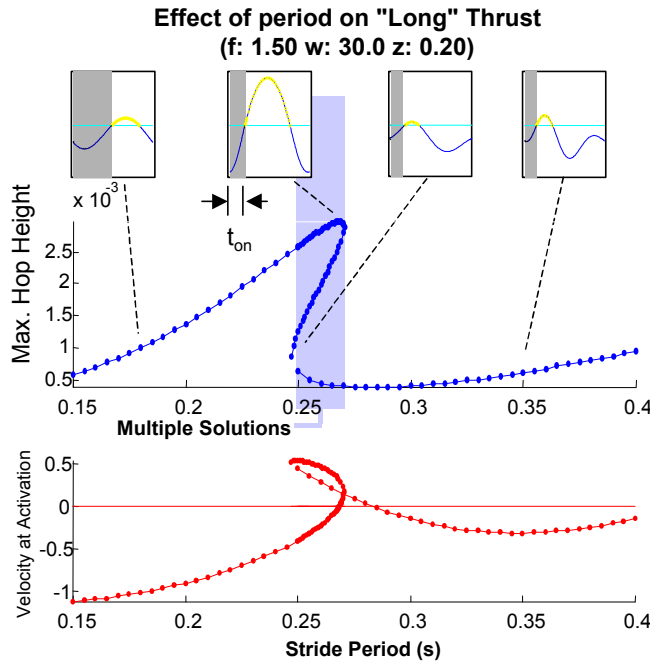


Figure 7. Effects of changing the stride period on steady-state motion and stability for the one-DOF hopper case in which thrust application ends at or after lift-off.

If $f(t)$ is constant, then this integral is proportional to the change in leg length between thrust activation and deactivation, as illustrated in Figure 6. For the case in which thrust is applied until the end of the ground contact phase, here called “Long Thrust,” this integral is maximized when thrust activation coincides with maximum compression, when the leg velocity is zero. However, for the case in which thrust ends before the end of the ground contact phase, here called “Short Thrust,” the conditions for the maximization of this integral are more complex but can be shown to roughly coincide with maximizing the upward velocity at thrust activation (while still achieving end-of-thrust before lift-off).

2.2 Effects of varying stride period

Figure 7 shows a typical example of the effects of changing the stride period of the open-loop motor pattern for the “Long Thrust” case with a given thrust magnitude, f , natural frequency, w , and damping ratio, ζ . For short stride periods, hopping height starts out very small, as shown in the top plot. At these periods, thrust application starts well before maximum compression, given by the negative velocity at thrust application (in the figure, this velocity is normalized by the magnitude of the take-off velocity). These solutions are termed “Regular Hopping” as they represent a desired mode of hopping behavior. As the stride period is increased, hopping height increases, and velocity at application approaches zero. Finally, at a certain period (near 275ms

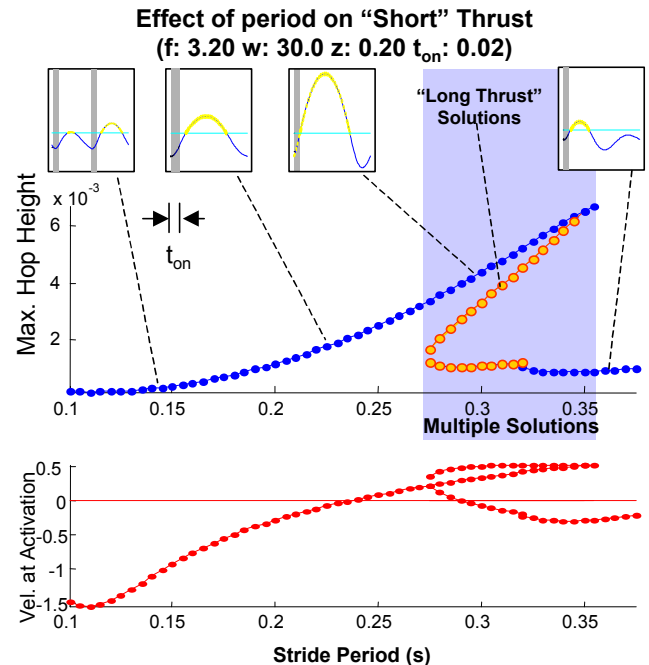


Figure 8. Effects of changing the stride period on steady-state motion and stability for the one-DOF hopper case in which thrust application ends before lift-off.

period), height is maximized when velocity at application is nearly zero, as predicted. Simulations of the hopper, though, never reach this point. As shown in the figure, other solutions to the state-steady conditions become available at a period near 250ms as the continuum of solutions folds back with respect to stride period. Of the two new sets of solutions available in this range of stride periods, one of them involves activating thrust after maximum compression and is generally found to be unstable. The solutions in the second set are termed “Hop-settle-fire” as the mass has started to settle before thrust is applied. The hopping heights for these solutions are much lower, but they are generally more stable, and the simulations converge to these solutions.

Figure 8 shows a typical example of the effects of changing the open-loop stride period for the “Short Thrust” case. For short periods, the solutions start out as “Regular Hopping.” As the period is further increased, the velocity at thrust activation increases, and changes from negative (thrusting before maximum compression) to positive (thrusting after maximum compression), while still maintaining stable hopping. Maximum hopping height is also increased with period and keeps increasing until the continuum of “Short Thrust” solutions ends as thrust deactivation starts to occur after the leg leaves the ground. Thus, hopping height is maximized when thrust is deactivated just as the leg loses contact with the ground. Near 275ms period another continuum of valid steady-state “Short Thrust” solutions begins.

This is again the “Hop-settle-fire” solution, for which hopping height is lower. The stability of these steady-state solutions is addressed in more detail in Cham (2003).

In general, the onset of the “Hop-settle-fire” solutions is determined by the system’s natural frequency. These solutions become available when the period of thrust application is long enough that the system is allowed to settle according to its natural period. The addition of damping also makes these solutions possible, since without damping the system would not settle.

2.3 One-DOF model conclusions

From the analysis above, we draw the following observations which we postulate generalize to similar hopping systems with clock-driven activation pattern:

1) For a wide range of activation periods, there exists one or more solutions to the steady-state conditions. These solutions may or may not be locally period-1 stable. If there are multiple solutions for a given period, the system will converge to the most stable solution, or may vacillate between equally stable solutions.

2) A given solution is such that the total amount of energy does not change within the stride period (the total amount of energy injected by the forcing function equals the total amount of energy passively dissipated). Since the magnitude and duration of the forcing function are determined by the stride period and duty cycle of the open-loop activation pattern, a given solution will entail a phase difference between the forcing function and the motion of the system such that the forcing function may perform both positive and negative work.

3) The total amount of energy within a stride is maximized when the forcing function performs the most positive work, given by the force-velocity integral in Equation 3. With a fixed thrust magnitude, this integral depends on the velocity at the start and end of thrust relative to the point of maximum compression such that it is proportional to the net change in leg length during thrust application. Activating thrust at the leg’s maximum compression may not be optimal in terms of maximizing the amount of work performed by the actuator if thrust duration is limited.

4) In general, increasing the stride period tends to increase the velocity at both the start and end of activation and maximize the work input integral. However, as shown, increasing the stride period will eventually result in unstable behavior (as in the “Long Thrust” case) or in “Hop-settle-fire” behavior, where the system settles to equilibrium between thrust periods.

These observations suggest that a simple way to infer how effectively the actuator is being utilized is to monitor the start and end of thrust relative to the motion of the body. Since we are interested in using only simple sensors such as binary contact switches, we pay particular attention to the relationship between end of thrust activation and the end of

the ground contact phase, or lift-off. The one-DOF analysis suggests that steady-state solutions in which the end of thrust occurs well before or after lift-off can be suboptimal in terms of the work input integral within one stride. This simple heuristic is explored and validated with experimental results of the multi-DOF hexapedal robot in the following section.

3. Stride Period Adaptation

3.1 Robot performance tests

The one-DOF model provided insight into the basic behavior of an open-loop hopping system in terms of the work performed by the actuator and the resulting performance. In order to develop an adaptation law for the six-legged, multi-DOF robot, we must look at the factors that affect its performance and see whether the same basic mechanisms are evident. Figure 9 shows the performance results of the hexapedal robot as a function of open-loop stride period for three different cases. The dotted lines represent the results for a first prototype, here called robot 1, running on flat ground. The solid lines are for the same robot on flat ground, but with different actuators, here called robot 2. The new actuators are pneumatic pistons with faster air flow and less damping. Finally, the dashed lines are for robot 2 running on a 5 degree uphill slope.

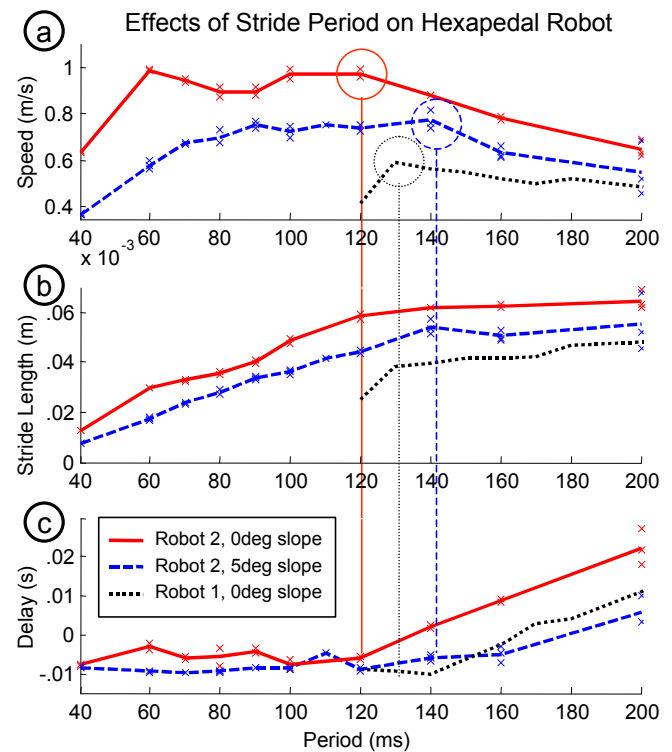


Figure 9. Performance tests for the hexapedal robot in flat ground and uphill terrain, as a function of open-loop stride period with constant duty cycle percentages.

As shown in 9a, speed is maximized at different stride periods for different conditions, again motivating the need for self-tuning or adaptation of the stride period. Starting at high values of the stride period, the motion of the robot was observed to be associated with “Hop-settle-fire” behavior similar to that observed in the one-DOF model, where the stride period is long enough that much of the energy from the previous hop is dissipated before thrust is initiated. As the period is decreased, speed increases and then levels off. This levelling off of the speed curve as stride period is decreased corresponds to the point in which stride length, plotted in Figure 9b, is no longer maximized. Although stride length decreases for shorter stride periods, speed continues to be maximized, due to the fact that speed is the product of stride length and the inverse of stride period. However, speed dramatically decreases below a certain period, due to the limitations in bandwidth of the pneumatic actuators, which limits the maximum speed attainable. Thus, speed is maximized over a range of periods bounded on one side by the bandwidth of the actuators, and on the other side by the stride period in which stride length is maximized.

As a goal of the adaptation, we choose the period in which stride length is maximized. Although speed continues to be maximized for shorter stride periods, the gain in speed is small relative to the increased energy consumption due to higher stride frequencies and lower efficiency. To find the stride period at which stride length is maximized, we investigate whether maximizing this measure of performance is related to maximizing the amount of work performed by the actuators, as suggested by the simplified model analysis.

Since work performed by the actuator is difficult to measure directly, we observe it indirectly through measurements of thrust timing relative to the motion of the system, again as suggested by the analysis. In the case of the Sprawl robots, the legs do not have a spring along the length of the leg with a nominal length that determines when ground contact occurs. Instead, for a given leg, contact with the ground occurs when the pistons extend, shortly after the valves are activated. Thus, thrust activation and landing occur a fixed time delay apart. Take-off, or loss of contact with the ground, on the other hand, occurs either when the leg reaches maximum extension, or when thrust is deactivated, which causes the leg to retract. As a result, we can indirectly measure the work performed by the actuator by looking at the time delay between the time that the valves are deactivated, and the time that take-off occurs. As indicated by the analysis of the simplified models, work is maximized when thrust deactivation occurs near the time that take-off occurs, given by the leg reaching maximum extension. Suboptimal work occurs when thrust is deactivated well before or well after maximum extension occurs.

The time difference between t_d , the time that thrust is deactivated, and t_T , the time that take-off occurs, ($t_d - t_T$), is

plotted in Figure 9c for the different cases examined. This time delay, ($t_d - t_T$), is positive for long stride periods, which indicates that thrust application ends after lift-off, here caused by end-of-stroke or full extension. This delay also monotonically increases for longer periods since the time of thrust application is set as a fixed percentage, or duty cycle, of the stride period due to valve and air flow limitations. Below a certain range of stride periods, however, the time delay is a nearly constant small negative value. In effect, deactivation of the tripod causes the spring-loaded leg pistons to retract and lose contact with the ground before full extension.

This change in the slope of the time delay ($t_d - t_T$) relative to the stride period occurs near the period for which stride length begins to decrease and ground speed starts to level off. Although the dynamics of the robot’s locomotion are affected by many factors, it is apparent that the stride period in which the amount of net positive work performed by the actuators is maximized, as indicated by the time delay ($t_d - t_T$), has a first-order correlation with the period in which stride length is maximized. This correlation is used as the basis for the simple adaptation law described in the following section.

3.2 Adaptation strategy

The results from the previous sections motivate the robot stride period adaptation strategy described here. As illustrated by the one-DOF model, it is advantageous to use a stride period that results in a steady-state cycle in which thrust is deactivated near the point where full piston extension occurs in order to maximize work input. Similar to the one-DOF model, lower stride periods result in sub-optimal work input as thrust is terminated before full extension. Moreover, like the one-DOF model, higher stride periods result in “Hop-settle-fire” behavior and sometimes in period-1 unstable oscillations. A prototype adaptation law for maximizing ground speed that takes these findings into consideration using foot contact information is as follows:

$$\tau_{n+1} = \tau_n - K_p(t_d - t_L - t_v) \quad (4)$$

Here, K_p is the adaptation gain, t_v is a constant offset parameter, t_d is the time at which the valve is deactivated and t_L is the measured lift-off time of the middle-foot. Figure 10 illustrates what these quantities represent, where time is measured with respect to the initiation of the gait cycle, which starts when the valve for one of the tripods is activated. In this case, the adaptation law is applied such that the stride period is updated at the end of every stride cycle that ground contact information is measured. Ground contact is measured by a binary switch attached to the middle foot of the same tripod. The deactivation time t_d is determined by the stride period, τ , and duty cycle, which in this

case is specified as a fixed percentage of the stride period. If there is no measured lift-off time, t_L , then the period is not modified.

Intuitively, this simple adaptation law can be described as trying to decrease the stride period as much as possible without exceeding the bandwidth of the actuators and without terminating the thrust application before full extension (to maximize available work). The stride period reaches an equilibrium value when $\Delta\tau$ is zero, which occurs when $(t_d - t_L)$ is equal to the offset value, t_v . Conditions for convergence of this adaptation law are derived in Appendix B, and include selection of an adaptation gain, K_p , such that:

$$K_p < 2/s \quad (5)$$

where s is the rate at which the time delay $(t_d - t_T)$ increases as a function of stride period for long stride periods. This condition is an absolute bound for stability, assuming that gait dynamics are negligible compared to the update rate of the adaptation law. In practice, if adaptation is needed to proceed at a comparable rate to the gait dynamics, then a lower value of K_p must be selected to ensure stability.

The offset value t_v should be selected near and slightly above the nearly constant baseline value of $(t_d - t_L)$ for short

stride periods, here called b . This ensures that the equilibrium stride period coincides with the change in slope (see Figure 10). In practice, considerable noise exists in the measurement of $(t_d - t_L)$, such that t_v cannot be selected too close to b , resulting in an offset error between the equilibrium stride period and the period at which the change in slope occurs.

3.3 Adaptation Results and Discussion

Figure 11 shows test results of the adaptation law implemented in the hexapedal robot 2 running on flat ground for several experiments in which the stride period was started at suboptimal values (see multimedia Extension 3). Figure 11a shows the ground speed of the robot as a function of time, and Figure 11b shows the stride period used to activate the tripods as it is changed by the adaptation law. The gain K_p was experimentally chosen to give the adaptation a fast learning rate while still achieving convergence. Note from Figure 11b that, although only a simple contact switch was used, the measured values of t_L are still prone to some noise, due to ground imperfections or disturbances to the robot, and adaptation does not necessarily proceed smoothly. This adaptation strategy was also shown to optimize speed in robot 1, with different pneumatic pistons, and for the case where the input actuator pressure was decreased in robot 2 by 13% (shown in Figure 12).

For an uphill ground slope of 5 degrees, the adaptation strategy also converges to an equilibrium stride period, as

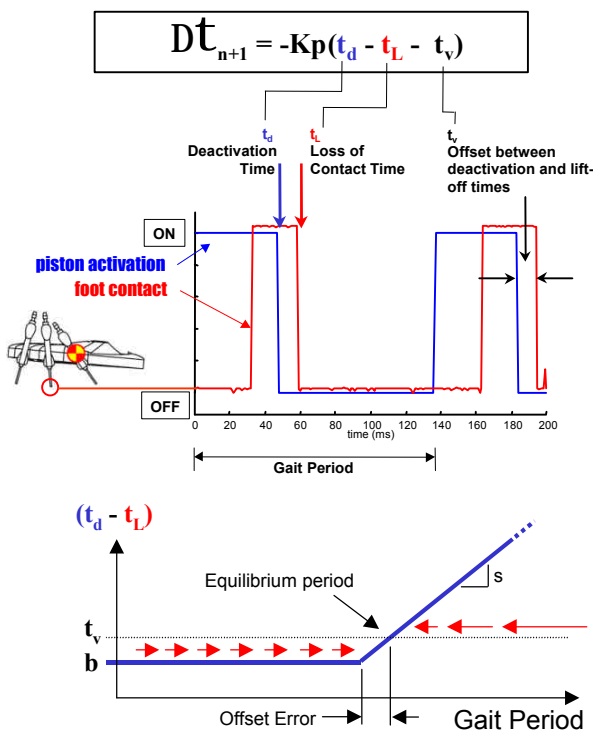


Figure 10. Simple gait period adaptation law based on the measured duration between foot contact events. A binary switch in the robot's middle foot provides contact information. This time duration is compared to the actuator valve deactivation time for the adaptation law.




Figure 11. Adaptation results for flat terrain (dashed lines are approximate optimal values established empirically). The figures show the ground speed of the robot and the stride period as it is adapted from suboptimal starting conditions.

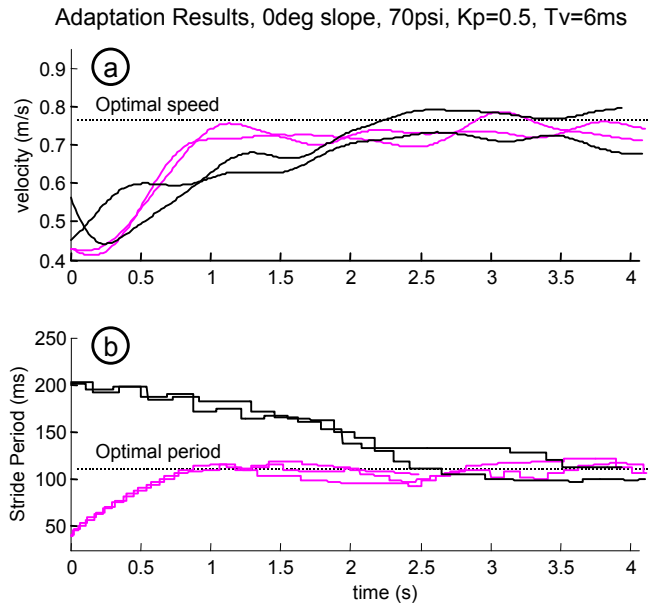


Figure 12. Adaptation results for flat terrain with a 13% decrease in pneumatic actuator input pressure. The adaptation optimizes ground speed by converging to a slightly higher stride period than in Figure 11.

shown in Figure 13 (see multimedia Extension 4). This new equilibrium period (~ 170 ms) is higher than the equilibrium period for flat ground running (~ 110 ms) and results in faster uphill running than with the optimal period for flat ground. However, the new equilibrium period is somewhat higher than the period found to be optimal at 5 deg. slopes (~ 140 ms). This indicates that, although it works to improve locomotion speed when transitioning to sloped terrain, the simple threshold-based adaptation law implemented here results in errors in the optimal equilibrium stride period for uphill running. This is mainly attributed to the gradual change in slope in the plot of $(t_d - t_l)$ for 5 deg uphill terrain compared to the prominent “kink” in the corresponding plot for flat terrain (see Figure 10). The use of a threshold to detect this change in slope results in equilibrium periods that are longer than optimal. Future work will experiment with more sophisticated ways to detect this change in slope. Furthermore, the factors that affect uphill running may need to be re-examined. For example, in seeking to increase the stride frequency for optimal ground speed in flat terrain, the prototype adaptation law presented here reduces oscillations in the direction normal to the ground, which we believe may have a significant role while climbing up-hill terrain.

4. Conclusions and Future Work

The analyses and experiments in the previous sections show that for an open-loop running robot, stride frequency and thrust duration are important parameters that govern hop height and forward speed. The single legged hopper model

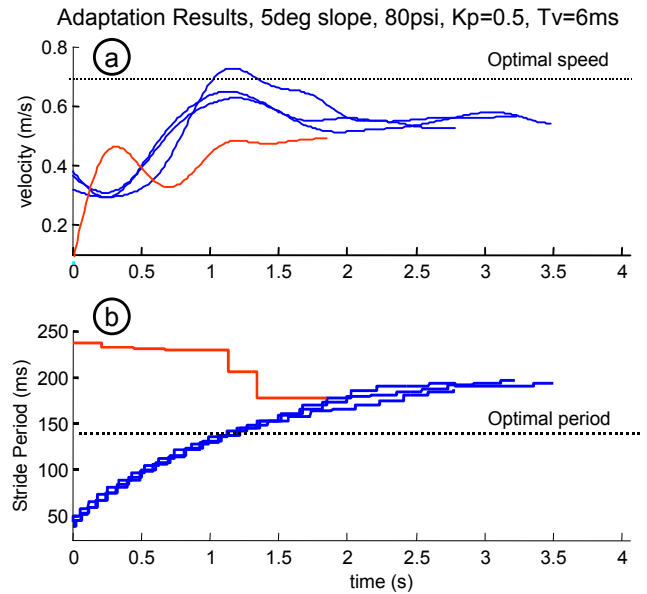


Figure 13. Adaptation results for an uphill slope of 5 degrees. The adaptation strategy improves the locomotion, but converges to a stride period slightly higher than the optimal stride period.

reveals that optimal hop height is obtained by maximizing the product of thrust force and velocity over the thrust duration. However, this product is subject to both dynamic constraints and hardware limitations. The dynamic constraints include the requirement of a stable, steady-state periodic solution to which the system will converge. Significant passive damping, as found in insects and in robots like *Sprawlita*, increases the regime of stable, periodic operation with open-loop forcing. The hardware limitations include the stroke length, the speed at which the piston can be filled and exhausted and the maximum thrust force available.

An adaptation strategy for the stride period that takes these limitations into account and tries to optimize ground speed was presented in this paper. The adaptation law seeks to obtain the most work from the actuators without exceeding their bandwidth. This adaptation law uses only the sensed duration of ground contact during each stride, and was shown to cause the stride period to converge to optimal values for a range of robot-to-robot variations and operating pressures. When making the transition from level to uphill running, the adaptation law improves locomotion, but converges to somewhat suboptimal values of stride period and velocity. The difficulty in this case is that the transition between optimal and over-long periods is less distinct and less easily identified with the simple threshold test used. More sophisticated detection of the transition is an area of ongoing work.

More generally, the adaptation scheme presented in this paper is an example of an approach that is particularly well suited for small, biomimetic robots by requiring no expen-

sive or sophisticated sensing or feedback. In this case, only simple binary switches are needed to provide an estimate of ground contact time. The adaptation scheme takes advantage of the passive properties of the robot that allow it to run stably over a range of open-loop stride frequencies and actuator duty cycles. In the event of sensor failure, the performance of the robot degrades only to that of the open-loop system without adaptation. This approach allows the robots to remain simple, inexpensive and robust while also being able to "tune" themselves to accommodate individual variabilities and changes in operating conditions.

Future work will build upon the simple adaptation law tested in this paper to incorporate other simple sensor information (e.g. tilt sensor, contact switches in other feet) in order to increase performance and adaptability. As discussed previously, further understanding of the robot's dynamic interaction with different types of terrain such as sloped or compliant surfaces will allow us to increase the adaptation's versatility. Finally, future work will study the effects of such an adaptation law on other types of behavior, such as rapid turning and navigation.

5. Acknowledgments

We have been fortunate to collaborate with Professor Robert J. Full at the U.C. Berkeley Polypedal Laboratory throughout this work and to benefit from his prodigious knowledge and understanding of the mechanisms underlying insect locomotion. We are also grateful for lengthy discussions on adaptation with Dr. Goran Djordjevic at the Johns Hopkins' Laboratory for Human Motor Learning. Thanks also to Jonathan Clark, Sean Bailey, Edward Froehlich and the other members of the Biomimetics team at CDR. Funding was provided by the Office of Naval Research under grant N00014-98-1-0669. Jonathan Karpick is supported by an Alex and Brit d'Arbeloff Stanford Graduate Fellowship.

6. References

- Ahn, A N; Full, R J. 1997. "A motor and a brake: Similar EMGs in two adjacent leg extensor muscles result in completely different function." *American Zoologist*. 37:107A.
- Bailey, S. A., Cham, J. G., Cutkosky, M. R., and Full, R. J., "Biomimetic Mechanisms via Shape Deposition Manufacturing," in *Robotics Research: the 9th International Symposium*, J. Hollerbach and D. Koditschek (Eds), Springer-Verlag, London, 1999.
- Berkemeier, M. D. and Desai, K.V., "A Comparison of Three Approaches for the Control of Hopping Height in Legged Robots." Submitted to the *International Journal of Robotics Research*, 1998.
- Cham, J. G., "On Performance and Stability in Open Loop Running," Ph. D. Thesis, Stanford University, 2003.
- Cham, J. G., Bailey, S. A., Clark, J. E., Full, R. J. and Cutkosky, M. R., "Fast and Robust: Hexapedal Robots via Shape Deposition Manufacturing," *Intl. Journal of Robotics Research* Vol 21., No. 10, 2002 pp. 869-883.
- DeCarlo, R., Branicky, M., Pettersson, S., and Lennartson B., "Perspectives and Results on the Stability and Stabilizability of Hybrid Systems" *Proceedings of the IEEE*, July 2000, pp. 1069-1082.
- Full, R.J., Autumn, K., Chung, J.I., Ahn, A., "Rapid negotiation of rough terrain by the death-head cockroach." *American Zoologist*. 38:81A. 1998.
- Garcia, M., Kuo, A., Peattie, A. M., Wang, P. C. and Full, R. J., "Damping and Size: Insights and Biological Inspiration" in *Proc. of the Intl. Symp. on Adaptive Motion of Animals and Machines*, Montreal, Canada, August 2000.
- Hatsopoulos, N. G., "Coupling the Neural and Physical Dynamics in Rhythmic Movements," in *Neural Computation*, 8, 567-581, 1996.
- Koditschek, D. E. and Buhler, M., "Analysis of a Simplified Hopping Robot," *International Journal of Robotics Research*, 10:6, 587-605. 1991.
- Komsuoglu, H. and Koditschek, D. E., "Preliminary Analysis of a Biologically Inspired 1-DOF 'Clock' Stabilized Hopper", *Proc. of World Multiconference on Systemics, Cybernetics and Informatics (SCI2000)*, Orlando, USA, July 23-26, 2000, Vol IX, pp 670-675.
- Kubow, T.M., Full, R.J., "The role of the mechanical system in control: A hypothesis of self-stabilization in hexapedal runners." *Phil. Trans. Roy. Soc. London B*. 354, 849-862. 1999.
- Meijer, K., Full, R.J., "Stabilizing properties of invertebrate skeletal muscle." *American Zoologist*. In press.
- Orlovsky, G. N., Deliagnia, T. G. and Grilner, S., "Neuronal Control of Locomotion," Oxford University Press, New York, 1999.
- Raibert, M. H., "Legged robots that balance." MIT Press, Cambridge, MA, 1986.
- Ringrose, R., "Self-stabilizing running," *IEEE ICRA Proceedings*, Albuquerque, NM, 1997.
- Rossignol, S., Lund, J. P. and Drew, T., "The Role of Sensory Inputs in Regulating Patterns of Rhythmical Movements in Higher Vertebrates," in "Neural Control of Rhythmic Movements in Vertebrates" edited by Cohen, A. H., Rossignol, S. and Grilner S., 1988.
- Sastry, S. "Nonlinear Systems: Analysis, Stability and Control". Springer Verlag. 1999.
- Vakakis, A. F., Burdick, J. W., and Caughey, T. K., "An Interesting Strange Attractor in the Dynamics of a Hopping Robot," *Intl. J. of Robotics Res.*, 10:606-618.

Appendix A - Vertical Hopper Model

Equations of Motion

The equations of motion for the one-DOF hopper in Figure 4 can be written, in normalized coordinates, as:

$$\dot{X} = AX + B \quad (\text{A.1})$$

where A and B are defined as:

$$A = \begin{bmatrix} 0 & 1 \\ -w^2 & -2\zeta w \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ f(t) - 1 \end{bmatrix} \quad \left. \vphantom{\begin{matrix} A \\ B \end{matrix}} \right\} \quad (x \leq 0) \quad (\text{A.2})$$

during the stance, or ground-contact phase, and:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ -1 \end{bmatrix} \quad \left. \vphantom{\begin{matrix} A \\ B \end{matrix}} \right\} \quad (x > 0) \quad (\text{A.3})$$

during the airborne, or ballistic, phase. w is the natural frequency and ζ is the damping ratio of the mass-spring-damper system. The thrust force $f(t)$ is determined by the open-loop motor control pattern:

$$f(t) = \begin{cases} T_n & t_{off} < t \leq t_{off} + t_{on} \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.4})$$

where T_n is the normalized thrust magnitude. Here, t is reset to $t=0$ when t reaches τ . This system is treated as a piecewise affine linear hybrid dynamic system with continuity of state at the mode transitions (Branicky, et al., 2000). The three modes are termed "AIR" (airborne phase), "ON" (stance phase with active thrust) and "OFF" (stance phase with zero thrust). The time solutions of the state vector $X(t)$ for the three modes are:

$$\text{AIR} \quad X(t) = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} X_0 - \begin{bmatrix} t^2/2 \\ t \end{bmatrix} \quad (\text{A.5})$$

$$\text{ON} \quad X(t) = e^{At}(X_0 - X_{eon}) + X_{eon}$$

$$\text{OFF} \quad X(t) = e^{At}(X_0 - X_{eoff}) + X_{eoff}$$

Here, X_0 is the state at the beginning of each mode and X_{eon} and X_{eoff} are the equilibrium states for each of the stance modes:

$$X_{eon} = \begin{bmatrix} \frac{T_n - 1}{w^2} \\ 0 \end{bmatrix} \quad X_{eoff} = \begin{bmatrix} \frac{1}{w^2} \\ 0 \end{bmatrix} \quad (\text{A.6})$$

Return Maps

In order to study the steady-state motion and local stability of the hopper, we define a return map $F(X_n)$ based on the state at thrust application, X_n (Sastry, 1999). Since mode switches are both a function of the state and of the open-loop motor pattern, the system trajectory can undergo an indeterminate number of sequences of mode changes. In this analysis, we consider the two hopping behaviors characterized as "Long Thrust" and "Short Thrust." In "Long Thrust," we assume that the hopper lands and activates thrust during stance, and that the thrust application duration is long enough to continue until or past lift-off, such that the mode sequence is ON-AIR-OFF. In "Short Thrust," we assume that the hopper lands and also activates thrust during stance, but that thrust application ends before lift-off, such that the mode sequence is ON-OFF-AIR-OFF.

For "Long Thrust," we introduce the two timing variables t_{on} (duration of active thrust application) and t_a (1/2 the duration of the airborne phase). To derive the return map, we take advantage of the facts that the take-off velocity (velocity at the ON-AIR mode transition) is the negative of the landing velocity (velocity at the AIR-OFF mode transition) and that this velocity is, in normalized coordinates, equal to t_a . We also take advantage of the fact that the total duration of the modes must equal τ . The return map can then be found by nesting the time solutions for the individual modes in the ON-AIR-OFF sequence:

$$X_{n+1} = F(X_n) = X_{eoff} - e^{A(\tau - 2t_a - t_{on})}(X_{eon} + X_{eoff}) \dots - e^{A(\tau - 2t_a)}(X_n - X_{eon}) \quad (\text{A.7})$$

In order to find the steady-state solutions, or "fixed points," we impose the constraint $X_{takeoff} = -X_{landing} = [0 \ t_a]^T$ and seek an expression with only two unknowns:

$$X_{takeoff} = e^{A(t_{on} + t_{off})}(X_{touchdown} - X_{eoff}) + \dots e^{At_{on}}(X_{eoff} - X_{eon}) + X_{eon} = -X_{touchdown} \quad (\text{A.8})$$

which implies:

$$X_{takeoff} + e^{A(t_{on} + t_{off})}(X_{takeoff} + X_{eoff}) + \dots e^{At_{on}}(X_{eon} - X_{eoff}) - X_{eon} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (\text{A.9})$$

This set of equations is solved numerically, where our solution vector is $[t_{on} \ t_a]$. If a solution exists, it may be unique or there may be multiple solutions. Once found, the solution $[t_{on}^* \ t_a^*]$ can be used to find X^* , which satisfies:

$$X^* = F(X^*) \quad (\text{A.10})$$

The return map for the ‘‘Short Thrust’’ mode sequence can be similarly found by nesting the time solutions in the ON-OFF-AIR-OFF sequence (see Cham, 2003, for more details).

Appendix B - Stability of the Adaptation Law

In this section we address the stability of the adaptation law presented in this paper. As stated, the adaptation law uses a simple threshold-based feedback law, presented in Equation 9. For simplicity in the derivation, we will rename some of the terms, such that the adaptation law becomes:

$$T_{n+1} = T_n - K(d - t) \quad (\text{B.1})$$

where K is the adaptation gain, t is the constant offset parameter, d is the delay between the time the valve is deactivated and the time that loss of contact is measured in the middle-foot. T_n and T_{n+1} are the current and the updated stride periods respectively.

To derive conditions on t and K for which the adaptation law will converge, we first assume that the time delay between deactivation and lift-off can be modeled as a direct function of the stride period, $d(T)$. Thus, we assume that the gait adjusts to changes in the stride period much faster than the rate at which the stride period is updated.

As described in section 3.1, the time delay $d(T)$ is approximately constant for short stride periods (when deactivation causes loss of contact with the ground), and monotonically increases at a nearly constant rate for longer stride periods (when maximum extension of the leg causes loss of contact). Thus, the time delay $d(T)$ can be well modeled as a piece-wise linear function:

$$d(T) = \begin{cases} b & T < T_c \\ s(T - T_c) + b & T \geq T_c \end{cases} \quad (\text{B.2})$$

where b is the baseline value of the time delay for short stride periods, s is the slope of the time delay function for higher stride periods, and T_c is the stride period at which this change in slope occurs. As stated, the adaptation law seeks to converge the current stride period to T_c , which indicates that maximum work is being performed by the actuators but which may vary according to changing conditions.

Ideally, t is chosen to be slightly greater than b , such that the adaptation law converges to an equilibrium when $d(T) = t$, which occurs when $T = T_c + (t - b)/s$. If t is chosen appropriately close to b , this will occur near T_c , as illustrated in Figure 10. We first show that the following condition:

$$K < 2/s \quad (\text{B.3})$$

is a necessary and sufficient condition for convergence of the adaptation law, assuming that $t > b$ and $s > 0$. To do so, we first consider starting conditions of the stride period, T , that lie within the range $T_c > T > T_c + 2(t - b)/s$, that is, a range

whose center is the point where $d(T) = t$ and bounded such that $d(T)$ is a smooth sloped line in this range. For this range, combining Equations B.1 and B.2 results in:

$$T_{n+1} = T_n - K(s(T_n - T_c) + b - t) \quad (\text{B.4})$$

$$T_{n+1} = (1 - Ks)T_n + K(t - b - sT_c) \quad (\text{B.5})$$

This one-dimensional discrete system is asymptotically stable if:

$$|1 - Ks| < 1 \quad (\text{B.6})$$

Assuming both K and s are always positive, this condition becomes:

$$K \cdot s < 2 \quad (\text{B.7})$$

Thus, any starting point of T within this range will converge asymptotically to a single value, as long as K is chosen such that $K < 2/s$. To demonstrate global convergence, we now show that any starting point outside this range eventually becomes a value inside the range.

For $T < T_c$, the function $d(T)$ is a constant value b , and Equation B.1 becomes:

$$T_{n+1} = T_n - K(b - t) \quad (\text{B.8})$$

Since $t > b$ and if $K < 2/s$, it can be shown that for $T < T_c$, T will increase in fixed step sizes until it falls within the range $T_c > T > T_c + 2(t - b)/s$, thus leading to convergence as shown above.

For $T > T_c + 2(t - b)/s$, Equation B.1 becomes Equation B.4, and since $d(T) > t$ in this range, it can be shown that T will decrease with finite step sizes until it falls either in the range $T < T_c$ or $T_c < T < T_c + 2(t - b)/s$, again leading to convergence.

For $K > 2/s$, the equilibrium point $T = T_c + (t - b)/s$ becomes unstable, such that any starting point within the range $T_c > T > T_c + 2(t - b)/s$ eventually leaves the range. In this case, it can be shown that a limit cycle exists in which T alternates between the following two values:

$$T_{LC1} = (t - b)/(2s) + T_c \quad (\text{B.9})$$

$$T_{LC2} = T_{LC1} - K(t - b) \quad (\text{B.10})$$

However, this limit cycle can be shown to be locally unstable for $K < 2/s$, which is a condition on its existence. Although the limit cycle is unstable, it can be shown that the stride period will oscillate about this limit cycle, but stay bounded in steady-state by the following two values:

$$T_{min} = T_c + K(t - b) \quad (\text{B.11})$$

$$T_{max} = T_c + K(t - b)(2 - Ks) \quad (\text{B.12})$$

Index to Multimedia Extensions

Extension	Media Type	Description
1	video	Video showing "Sprawlita" running despite large disturbances. This disturbance rejection is accomplished without sensory feedback through the robot's passive properties and open-loop control.
2	video	Video showing "Sprawlita" overcoming hip-height obstacles.
3	video	Video showing sample results of implementing the adaptation strategy based on binary contact information from a switch in one the robot's feet. The video shows experiments on flat ground in which the robot's stride period was started at suboptimal values.
4	video	Video sample results of the adaptation strategy on an uphill slope of 5 degrees. The video shows experiments in which the robot's stride period was started at suboptimal values.