9-1-2004

Model-Based Dynamic Self-Righting Maneuvers for a Hexapedal Robot

Uluc Saranli
Carnegie Mellon University

Alfred A. Rizzi
Carnegie Mellon University

Daniel E. Koditschek
University of Pennsylvania, kod@seas.upenn.edu

Reprinted from *The International Journal of Robotics Research*, Volume 23, Issue 9, September 2004, pages 903-918.
DOI: 10.1177/0278364904045594

NOTE: At the time of publication the author, Daniel Koditschek, was affiliated with the University of Michigan. Currently, he is a faculty member of the School of Engineering at the University of Pennsylvania.

Model-Based Dynamic Self-Righting Maneuvers for a Hexapedal Robot

Uluç Saranli
Alfred A. Rizzi

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15223, USA

Daniel E. Koditschek

Department of Electrical Engineering
and Computer Science
The University of Michigan
Ann Arbor, MI 48109-2110, USA

Abstract

We report on the design and analysis of a controller that can achieve dynamical self-righting of our hexapedal robot, RHex. Motivated by the initial success of an empirically tuned controller, we present a feedback controller based on a saggital plane model of the robot. We also extend this controller to develop a hybrid pumping strategy that overcomes actuator torque limitations, resulting in robust flipping behavior over a wide range of surfaces. We present simulations and experiments to validate the model and characterize the performance of the new controller.

KEY WORDS—legged robot, model-based control, contact modeling, dynamic manipulation, experimentation

1. Introduction

RHex (see Figure 1) is an autonomous hexapod robot that negotiates badly irregular terrain at speeds better than one body length per second (Saranli, Buehler, and Koditschek 2001). In this paper, we report on efforts to extend RHex's present capabilities with a self-righting controller. Motivated by the successes and limitations of an empirically developed largely open-loop "energy pumping" scheme, we introduce a careful multi-point contact and collision model so as to derive the maximum benefit of our robot's limited power budget. We present experiments and simulation results to demonstrate that the new controller yields significantly increased performance and extends on the range of surfaces over which the self-righting maneuver succeeds.

Physical autonomy—on-board power and computation—is essential for any robotic platform intended for operation in the real world. Beyond the strict power and computational constraints, unstructured environments demand some degree of behavioral autonomy as well, requiring at least basic self-manipulation capabilities for survivability in the absence (or inattention) of a human operator. Even during teleoperation, where the computational demands on the platform are less stringent, the ability to recover from unexpected adversity through self-manipulation is essential. Space applications such as planetary rovers and similar exploratory missions probably best exemplify settings where these requirements are most critical (Altendorfer et al. 2001).

Recovery of correct body orientation is among the simplest of self-manipulation tasks. In cases where it is impossible for a human operator to intervene, the inability to recover from a simple fall can render a robot completely useless and, indeed, the debilitating effects of such accidents in environments with badly broken terrain and variously shaped and sized obstacles have been reported in the literature (Bares and Wettergreen 1999).

RHex's morphology is roughly symmetric with respect to the horizontal plane, and allows nearly identical upside-down or right-side-up operation, a solution adopted by other mobile platforms (Matthies et al. 2000). However, many application scenarios such as teleoperation and vision-based navigation entail a nominal orientation arising from the accompanying instrumentation and algorithms. In such settings, designers typically incorporate special kinematic structures, e.g., long extension arms or reconfigurable wheels (Tunstel 1999; Hale et al. 2000; Fiorini and Burdick 2003), to secure such vital self-righting capabilities. In contrast, the imperatives of dynamical operation that underly RHex's design and confer its unusual mobility performance (Saranli, Buehler, and Koditschek 2001) preclude such structural appendages. RHex must rely on its existing morphology and dynamic maneuvers to achieve the necessary self-righting ability.

The International Journal of Robotics Research
Vol. 23, No. 9, September 2004, pp. 903-918,
DOI: 10.1177/0278364904045594

Fig. 1. RHex 1.5.

There is a significant body of literature in the control of locomotion addressing similar problems arising from both the dynamic and the hybrid nature of such systems. The work of Raibert (1986) on dynamically stable hopping robots was influential in the development of various other systems capable of performing dynamical maneuvers such as biped gymnastics (Hodgins and Raibert 1990) and brachiating robots (Nakanishi, Fukuda and Koditschek 2000). However, despite structural similarities, these methods are not directly applicable to our problem as they either aim to stabilize the system around neutral periodic orbits or concentrate on the control of non-holonomic flight dynamics.

Quasi-static posture control has been explored in the legged robotics literature (Waldron and Vohnout 1984; Nelson and Quinn 1999), but not the dynamical problem of present concern. In particular, the problem of dynamically righting a legged platform introduces the need to consider intermittent multiple contacts and collisions, while incurring constraints on feasible control strategies familiar within the legged robotics literature, arising from morphology, actuator and sensory limitations. Our recourse to an energy pumping control strategy is informed by earlier work on dynamically dexterous robotics such as the swing-up of a double pendulum (Spong 1995; Nakanishi, Fukuda, and Koditschek 1999; Yoo, Yang, and Hong 2001), which involves some of these constraints but, notably, does not require consideration of the hybrid nonlinearities that are inherent to our system (e.g., see Figure 4). Similarly, recent work on jumping using computational learning algorithms (Zhang et al. 1997) and simulation studies of ballistic flipping (Geng, Li, and Xu 2002; Geng et al. 2002) using Poincaré maps for the design of stable control policies for one-legged locomotion contend with aspects of dynamics relevant to self-righting, but consider neither multiple colliding contacts nor inherent or explicit constraints on feasible control inputs.

In this light, the central contributions of this paper include: (i) introducing a new multiple point collision/contact model that characterizes RHex's behavior during the flipping maneuver; (ii) the description of a new torque control strategy that uses the model to maximize the energy injected into the system in the face of these constraints (i.e., consistent with maintaining a set of postural invariants integral to the task at hand). We present experimental and simulation evidence to establish the validity of the model and demonstrate that the new controller significantly improves on the performance of our first generation open-loop controller.

2. Flipping with RHex

RHex's dynamic locomotion performance arises from our adoption of specific principles from biomechanics such as structural compliance in the legs and a sprawled posture (Altendorfer et al. 2001). Furthermore, its mechanical simplicity, with only one actuator per leg and minimal sensing, admits robust operation in outdoor settings over extended periods of time.

The rotation axes for RHex's actuators are all parallel and aligned with its transverse horizontal body axis. Consequently, the most natural backflip strategy for RHex pivots the body around one of its endpoints. Pitching the body in this manner, while keeping one of the body endpoints in contact with the ground, maximizes contact of the legs with the ground for the largest range of pitch angles and thus promises to yield the best utilization of available actuation. In contrast, flipping by producing a sideways rolling motion suffers from early liftoff of three legs on one side as well as the longer protrusion of the middle motor shafts.

For surfaces with sufficiently low lateral inclination, RHex's rectangular body and lateral symmetry restricts the motion described above to the saggital plane. When the tail or the nose of the body is fully in contact with the ground, the resulting support line provides static lateral stability as long as the gravity vector falls within the contact surface (see Section 3). As a result, a set of planar models suffices to analyze the flipping behavior within the acceptable range of inclinations.

Clearly, large slopes will invalidate this assumption and may lead to non-planar motion. However, we limit the scope of this paper to analysis on relatively flat terrain wherein the planar nature of the flipping motion remains valid. Before formally introducing the planar flipping models in Section 3, we will find it useful to describe the general structure of the flipping controller, as well as motivations and assumptions underlying its design.

2.1. Basic Controller Structure

All the flipping controllers presented in this paper share the same finite state machine structure, illustrated in Figure 2. Starting from a stationary position on the ground, the robot quickly thrusts itself upward while maintaining contact between the ground and the endpoint of its body (poses I and

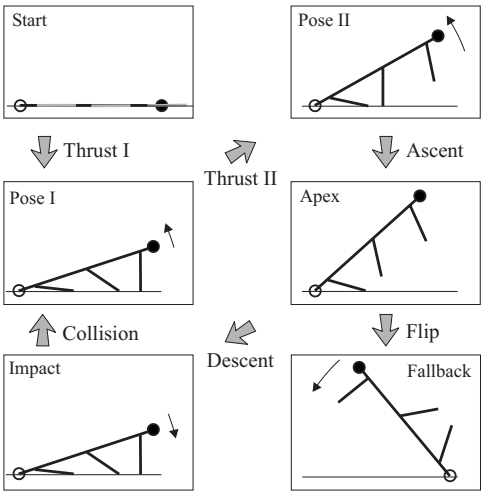

Fig. 2. Sequence of states for the flipping controller.

II in Figure 2) as the front and middle legs successively leave the ground. Depending on the frictional properties of the leg/ground contact, this thrust results in some initial kinetic energy of the body that may in some cases be sufficient to allow "escape" from the gravitational potential well of the initial configuration and fall into the other desired configuration. In cases where a single thrust is not sufficient to flip the body over, the robot reaches some maximum pitch lying within the basin of the original configuration, and falls back toward its initial state. Our controller then brings the legs back to Pose I of Figure 2 and waits for the impact of the front legs with the ground, avoiding negative work—a waste of battery energy given the familiar power-torque properties of RHex's conventional DC motors. The impact of the compliant front legs with the ground in their kinematically singular configuration recovers some of the body's kinetic energy, followed by additional thrust from the middle and back legs, during the period of decompression and flight of the front leg, i.e., during a phase interval when it is possible for the legs in contact to perform positive work on the robot's mass center. The maximum pitch attained by the body increases with each bounce up until the point where the robot flips or the energy that can be be imparted by the thrust phase balances collision losses at which point it must follow that flipping is not possible.

2.2. Observations and Motivation

The performance of the flipping controller is predominantly determined by the amount of energy that can be injected into the system through the "thrust" phase. In contrast, the feasibility of the hybrid pumping mechanism depends on the success of the thrust controller in maintaining body ground contact to ensure robust recovery of kinetic energy at impact. The main contribution of this paper is the design of effective thrust controllers and their analysis in conjunction with the hybrid pumping scheme to characterize the performance of flipping.

Our first generation flipping controller was primarily open loop at the task level, wherein we used high gain proportional derivative (PD) control to "track" judiciously selected constant velocity leg sweep motions (Saranli and Koditschek 2001, 2002). This scheme was motivated by its simplicity as well as the lack of adequate proprioceptive sensing capabilities in our experimental platform.

As reported in Saranli and Koditschek (2002), this simple strategy is capable of inducing a backflip of our earlier experimental platform (RHex version 0.5) for a variety of surfaces (see Extensions 1 and 2 for movies). However, it does so with relatively low efficiency (in terms of the number of required bounces) and low reliability. It shows very poor performance and reliability on softer surfaces such as grass and dirt—outdoor environments most relevant to RHex's presumed mission (Altendorfer et al. 2001; Saranli, Buehler, and Koditschek 2001). Furthermore, as we report in this paper, it fails altogether on newer versions of RHex which are slightly larger and heavier. To permit a reasonable degree of autonomous operation, we would like to improve on the range of conditions under which flipping can function. This requires a more aggressive torque generation strategy for the middle and rear legs. However, empirically, we find that driving all available legs with the maximum torque allowed by the motors usually results in either the body lifting off the ground into a standing posture, or unpredictable roll and yaw motions eliminating any chance for subsequent thrust phases. Rather, we seek a strategy that can be tuned to produce larger torques aimed specifically at pitching the body over. This requires a detailed model of the manner in which the robot can elicit ground reaction forces in consequence of hip torques operating at different body states and assuming varying leg contact configurations.

3. Planar Flipping Models

In this section, we present a number of planar models, starting with a generic model in Section 3.2, followed by various constrained versions in Sections 3.5 and 3.6. In each case, we derive the corresponding equations of motion, based on the common framework of Section 3.3.

3.1. Assumptions and Constraints

Several assumptions constitute the basis for our modeling and analysis of the flipping behavior.

ASSUMPTION 1. The flipping behavior is primarily planar.

The controller structure described in Section 2.1 operates contralateral pairs of legs in synchrony. On flat terrain, the robot's response lies almost entirely in the saggital plane and

departures are rare enough to be negligible. Our models and analysis will hence be constrained to the saggital plane.

Even though the scope of the present paper does not address in detail the flipping behavior on sloped surfaces, this assumption can be intuitively justified by the observation that the full contact of one of the body endpoints with the ground, if successfully enforced by the controller, yields lateral static stability by canceling the lateral moment induced by the action of gravity on the body. The largest moment is produced when the body is standing vertically on one of the endpoints, and can be counteracted for slopes of up to $\text{atan}(w/l)$ where w is the body width and l is the body length. Even though we do not present systematic experiments to verify this observation, this simple model suggests the potential validity of our planar analysis for a considerable range of lateral slopes as well.

ASSUMPTION 2. The leg masses are negligible relative to the body mass.

We assume that the leg masses are sufficiently small so that their effect on the body dynamics is limited to the transmission of the ground reaction forces at the toes to the body when they are in contact with the ground. This assumption is a fairly accurate approximation as a result of the very light fiber-glass legs on our experimental platform.

ASSUMPTION 3. The tail of the body should maintain contact with the ground throughout the flipping action.

This assumption is motivated by a number of observations gathered during our empirical flipping experiments. First, during the initial thrust phases, the front and middle legs provide most of the torque. Configurations where the tail endpoint of the body is in contact with the ground yield the longest duration of contact for these legs, harvesting greatest possible benefit from the associated actuators.

Furthermore, collisions of the body with the ground, which introduce significant losses due to the high damping in the body structure designed to absorb environmental shocks, can be avoided by preserving contact with the ground throughout the flipping action. It is also clear that one would not want to go through the vertical configuration of the body when the tail endpoint is not in contact with the ground as such configurations require overcoming a higher potential energy barrier and would be less likely to succeed.

Finally, the body ground contact is essential for maintaining the planar nature of the behavior and eliminating body roll. This is especially important for repeated thrust attempts of the hybrid energy pumping scheme, which rely on the robot body being properly aligned with as much of the impact kinetic energy recovered as possible.

In light of these assumptions, the design of thrust controllers has to satisfy two major constraints: keeping the tail endpoint of the body on the ground and respecting the torque limitations of the actuators.

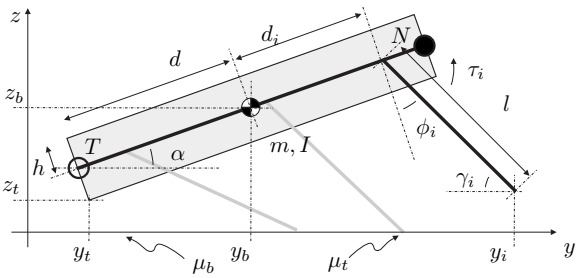

Fig. 3. Generic three-degrees-of-freedom (3DOF) planar flipping model.

3.2. The Generic Model

Even though our analysis will be largely confined to control strategies that enforce configurations where the tail of the body remains on the ground, we will find it useful to introduce a more general model to prepare a formal framework in which we will define various constraints.

Figure 3 illustrates the generic planar flipping model. Three massless rigid legs—each representing a pair of RHex's legs—are attached to a rectangular rigid body with mass m and inertia I. The attachment points of the legs are fixed at d_i, along the mid-line of the rectangular body. This line also defines the orientation of the body, α, with respect to the horizontal. The center of mass (COM) is midway between the points N and T, defined to be the "nose" and the "tail", respectively. The body length and height are $2d$ and $2h$, respectively. Finally, we assume that the body–ground and toe–ground contacts experience Coulomb friction with coefficients μ_b and μ_t, respectively. Table 1 summarizes the notation used throughout the paper.

Neither the rectangular body nor the toes can penetrate the ground. Our model hence requires that the endpoints of the body be above the ground

$$z_b > \begin{cases} d\sin|\alpha| + h\cos\alpha & \text{if } |\alpha| < \pi/2 \\ d\sin|\alpha| - h\cos\alpha & \text{otherwise} \end{cases}, \qquad (1)$$

and that a leg must reach the ground

$$z_b > l - d_i \sin\alpha \qquad (2)$$

before it can apply any torque to the body. As a result, the configuration space[1] (α, z_b) is partitioned into various regions, each with different kinematic and dynamic structure as illustrated in Figure 4. In the figure, the solid line corresponds to configurations where one of the body endpoints is in contact with the ground, determined by eq. (1). All the configurations below this line (white region) are inaccessible as they would require the body to penetrate the ground. Similarly, different

1. Contact constraints are invariant with respect to horizontal translation, allowing for the elimination of y_b.

Table 1. Notation Used Throughout the Paper

States and dependent variables	
$\mathbf{c} \in \mathcal{X}$	System configuration vector
$\mathbf{q} := [\mathbf{c}, \quad \dot{\mathbf{c}}]^T$	System state vector
y_b, z_b	Body COM coordinates
α	Body pitch
y_t, z_t	Coordinates of the tail endpoint
ϕ_i, γ_i	Hip and toe angles for ith leg
$y_i, \dot{y}_i$	Position and velocity of the ith toe
Contact forces	
F_i^y, F_i	GRF components on ith toe
F_c^y, F_c^z	GRF components on the tail
Control inputs	
$\boldsymbol{\tau} \in \mathbb{R}^3$	Hip torque control vector
$\mathcal{T}(\mathbf{q}) \subseteq \mathbb{R}^3$	Set of allowable torque vectors
Planar model parameters	
d, h	Body length and height
d_i, l	Leg attachment and length
μ_t, μ_b	Coulomb coefficient for toes and body
m, I	Body mass and inertia
k_r	Coefficient of restitution for rebound
Motor model parameters	
v_s	Power supply voltage
r_d, r_a	Motor drive and armature resistances
K_s, K_τ	Motor speed and torque constants
m_g, h_g	Motor gear ratio and efficiency

Table 2. RHex's Kinematic and Dynamic Parameters

d	0.25 m	h	0.05 m
d_1	–0.19 m	m	8.5 kg
d_2	0.015 m	I	0.144 kg m^2
d_3	0.22 m	l	0.17 m

shades of gray in Figure 4 represent the number of legs that can reach the ground for a given configuration, with the boundaries determined by eq. (2). All legs can reach the ground for configurations shaded with the darkest gray whereas all legs must be in flight for those configurations shaded with the lightest gray. The shaded regions also extend naturally to configurations with body ground contact.

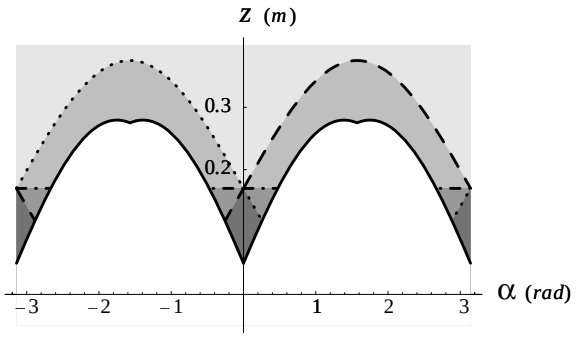

Fig. 4. Hybrid regions in the planar flipping model based on RHex's morphology (see Table 2). Solid lines indicate body ground contact for the nose ($\alpha < 0$) and the tail ($\alpha > 0$). The liftoff transitions of the front, middle and back legs are represented by dotted, dash-dot and dashed lines, respectively. Lighter shades of gray indicate that fewer legs can reach the ground.

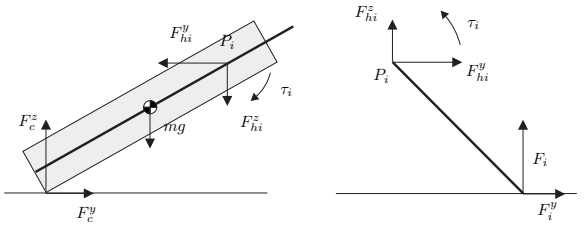

Fig. 5. Free body diagrams for the body and one of the legs.

3.3. Framework and Definitions

In deriving the equations of motion for all constrained models in this paper, we use a Newton–Euler formulation, presented in this section so as to unify the free-body diagram analysis of all three models.

Figure 5 illustrates the generic free-body diagrams for the body link and one of the leg links. Based on whether a link is in flight, in fixed contact with the ground or sliding on the ground, the associated force and moment balances yield linear equations in the unknown forces and accelerations, taking the form

$$\mathbf{A}(\mathbf{c})\mathbf{v} = \mathbf{b}(\mathbf{c}, \dot{\mathbf{c}}) + \mathbf{D}(\mathbf{c})\ \boldsymbol{\tau}\ . \tag{3}$$

Here, $\boldsymbol{\tau} := [\tau_1, \quad \tau_2, \quad \tau_3]^T$ is the torque actuation vector, $\mathbf{c}$ is the configuration vector, and $\mathbf{v}$ is the vector of unknown forces and accelerations. The definitions of both $\mathbf{c}$ and $\mathbf{v}$, as well as the matrices $\mathbf{A}(\mathbf{c})$, $\mathbf{b}(\mathbf{c}, \dot{\mathbf{c}})$ and $\mathbf{D}(\mathbf{c})$ are dependent on the particular contact configuration and will be made explicit in subsequent sections.

3.4. Unconstrained Dynamics with No Body Contact

Ideally, our flipping controllers will attempt to maintain contact between the body and the ground. However, part of our analysis requires the investigation of the unconstrained dynamics.

For this general case, no ground reaction forces act on the body link and the tail end of the body is free to move. Furthermore, assuming that all legs are in sliding contact with the ground, the friction forces take the form $F_i^y = -\mu_t \ F_i \ \text{sign}(\dot{y}_i)$, where $\dot{y}_i$ represents the translational velocity of the ith foot. In this case, the vector of unknowns and the system state are defined as

$$\begin{aligned} \mathbf{v} &:= [F_1, \ F_2, \ F_3, \ \ddot{\alpha}, \ \ddot{y}_t, \ \ddot{z}_t]^T \\ \mathbf{c} &:= [\alpha, \ y_b, \ z_b]^T \ . \end{aligned} \tag{4}$$

For each leg, we can write the moment balance equations as

$$(l \cos\gamma_i + l\bar{\mu}_i \sin\gamma_i) F_i = -\tau_i \ , \tag{5}$$

where $\bar{\mu}_i := -\mu_t \ \text{sign}(\dot{y}_i)$ is the effective Coulomb friction coefficient and γ_i corresponds to the toe angle as shown in Figure 3. In the operational range of the flipping controller, these equations are solvable. However, there are interesting "jamming" singularities in the remaining parts of the state space, which we investigate in Section 3.7.

Similarly, force and moment balances for the body link yield

$$\begin{aligned} \bar{\mu}_1 F_1 + \bar{\mu}_2 F_2 + \bar{\mu}_3 F_3 - m\ddot{y}_b &= 0 \\ F_1 + F_2 + F_3 - m\ddot{z}_b &= mg \\ \sum_{i=1}^{3} (d_i \cos\alpha - d_i \bar{\mu}_i \sin\alpha) F_i - I\ddot{\alpha} &= \sum_{i=1}^{3} \tau_i \end{aligned} \tag{6}$$

where $\ddot{y}_b$ and $\ddot{z}_b$ are components of the body acceleration and can be written as affine functions of $\ddot{\alpha}$, $\ddot{y}_t$ and $\ddot{z}_t$ by simple differentiation of the kinematics. The combination of eqs. (5) and (6) yields the matrices $\mathbf{A}(\mathbf{c})$, $\mathbf{b}(\mathbf{c}, \dot{\mathbf{c}})$ and $\mathbf{D}(\mathbf{c})$.

3.5. Dynamics with Sliding Body, Sliding Toe Contacts

In general, we observe that throughout the execution of our flipping behaviors, both the leg and body contacts slide on the ground. As a consequence, we can rewrite the horizontal components of ground reaction forces in terms of their vertical components using Coulomb's friction law. Here, the vector of unknown quantities becomes

$$\begin{aligned} \mathbf{v} &:= [F_1, \ F_2, \ F_3, \ \ddot{\alpha}, \ F_c^z, \ \ddot{y}_t]^T \\ \mathbf{c} &:= [\alpha, \ y_b]^T \ , \end{aligned} \tag{7}$$

yielding a system with two degrees of freedom: the body pitch and the horizontal position of the tail.

In this case, the moment balance for each leg remains the same as eq. (5) and the body balance equations become

$$\begin{aligned} \bar{\mu}_1 F_1 + \bar{\mu}_2 F_2 + \bar{\mu}_3 F_3 - \bar{\mu}_b \ F_c^z - m\ddot{y}_b &= 0 \\ F_1 + F_2 + F_3 + F_c^z - m\ddot{z}_b &= mg \\ \sum_{i=1}^{3} (d_i \cos\alpha - d_i \bar{\mu}_i \sin\alpha) F_i + [(h + \bar{\mu}_b d)\sin\alpha & \\ + (\bar{\mu}_b h - d)\cos\alpha] F_c^z - I\ddot{\alpha} &= \sum_{i=1}^{3} \tau_i \ , \end{aligned} \tag{8}$$

where, once again, system kinematics yields the body accelerations $\ddot{y}_b$ and $\ddot{z}_b$ as functions of $\ddot{\alpha}$ and $\ddot{y}_t$. As before, the combination of eqs. (5) and (8) yields the matrices $\mathbf{A}(\mathbf{c})$, $\mathbf{b}(\mathbf{c}, \dot{\mathbf{c}})$ and $\mathbf{D}(\mathbf{c})$.

3.6. Dynamics with Sliding Body, Fixed Rear Toe Contact

The third and final contact configuration we consider corresponds to cases where the rear toe is stationary under the influence of stiction. This model is primarily motivated by the observed behavior of various flipping controllers, where the rear toe stops sliding following the liftoff of the front and middle pairs of legs. Consequently, we incorporate this model into our feedback controller to be activated when the measured (or estimated) system state indicates that the rear toe is indeed stationary. Here, the vector of unknown quantities is

$$\begin{aligned} \mathbf{v} &:= [F_1 \ \ F_2 \ \ F_3 \ \ \ddot{\alpha} \ \ F_c^z \ \ F_1^y] \\ \mathbf{c} &:= \alpha \ , \end{aligned} \tag{9}$$

leaving a system with a single degree of freedom: the body pitch α. In this case, however, the moment balance for the rear leg is slightly different and includes the unknown horizontal ground reaction force, yielding

$$l\cos\gamma_1 F_1 + l\sin\gamma_1 F_1^y = -\tau_1, \tag{10}$$

while the moment balance equations for the middle and front legs remain the same as eq. (5). Finally, the balance equations for the body link now take the form

$$\begin{aligned} F_1^y + \bar{\mu}_2 F_2 + \bar{\mu}_3 F_3 - \bar{\mu}_b \ F_c^z - m\ddot{y}_b &= 0 \\ F_1 + F_2 + F_3 + F_c^z - m\ddot{z}_b &= mg \\ \sum_{i=1}^{3} (d_i \cos\alpha) F_i - \sum_{i=2}^{3} (d_i \bar{\mu}_i \sin\alpha) F_i & \\ - (d_1 \sin\alpha) F_1^y + [(h + \bar{\mu}_b d)\sin\alpha & \\ + (\bar{\mu}_b h - d)\cos\alpha] F_c^z - I\ddot{\alpha} &= \sum_{i=1}^{3} \tau_i . \end{aligned} \tag{11}$$

Similar to the previous two models, system kinematics yields the body accelerations $\ddot{y}_b$ and $\ddot{z}_b$ as functions of $\ddot{\alpha}$ we use eqs. (10) and (11) to compute the matrices $\mathbf{A}(\mathbf{c})$, $\mathbf{b}(\mathbf{c}, \dot{\mathbf{c}})$ and $\mathbf{D}(\mathbf{c})$.

3.7. Existence of Solutions and Leg Jamming

In the preceding sections, we presented a number of constrained models with their associated equations in the unknown forces and accelerations. However, the equations by themselves do not ensure the existence of solutions. In this section, we present conditions sufficient for these model to admit solutions, and show that the flipping controller operates within the resulting consistent regions in the state space.

In this context, a major singularity arises in computing the ground reaction forces on sliding legs using the moment balance equation (5). To illustrate the inconsistency, suppose that leg i is sliding forward with $\dot{y}_i > 0$ and the leg is within the friction cone with $\cot \gamma_i < \mu_t$. When $\tau_i < 0$, the massless legs in our model require a positive vertical component for the ground reaction force, $F_i > 0$. However, the solution of the leg moment balance equation yields

$$F_i = -\frac{\tau_i}{l \cos \gamma_i - l \mu_t \sin \gamma_i} < 0 \ , \tag{12}$$

resulting in an inconsistency. Consequently, when the leg is sliding forward and is inside the friction cone, there are no consistent solutions for the unknown forces and accelerations.

It turns out that this problem is a special case of the well-known Painlevé problem of a rigid rod sliding on a frictional surface (Painlevé 1895; Mason 2001). For certain parameter and state combinations, it is impossible to find any consistent set of finite forces and accelerations and one needs to seek impulsive solutions for the unknown quantities. This problem and its variations stimulated a large body of work in frictional collisions (Baraff 1991; Wang and Mason 1992; Stewart and Trinkle 1997; Stewart 1998), which hypothesize that the rigid rod would "jam" in such cases and start pivoting around its contact point.

Even though such impulsive force based approaches are extremely useful in evaluating the equations of motion for simulation purposes, their utility diminishes significantly when our goal is the design of a feedback controller. Even very small parametric errors or sensor noise could result in the measured state becoming inconsistent, putting the system outside the domain of the model-based controller. Unlike simulated systems, we do not have the luxury of applying impulsive forces to a physical robot through its actuators to bring it to a state where consistent solutions exist.

Fortunately, empirical evidence accumulated over months of physical experiments with the robot reveals that, in the absence of dramatic external disturbances, RHex operates in regions of its state space away from these singularities. Starting from a stationary position, the front four legs always slide backward, which guarantees a solution for the associated ground reaction forces. Furthermore, even though the rear legs usually slide forward, RHex's kinematics ensure that the orientation of the rear two legs is always outside the friction cone, yielding a consistent solution for the associated reaction forces. Finally, the body link always slides forward and admits a consistent solution once the toe reaction forces are identified.

4. Model-Based Control of Flipping

We have presented, in the previous sections, the equations of motion for a variety of planar flipping models that are constrained versions of the generic model described in Section 3.2. In this section, we use these models to design a controller that is capable of performing dynamic back flips with our hexapod platform.

In particular, our controller attempts to maximize the acceleration of the body pitch, while maintaining contact of the body endpoint with the ground and respecting torque constraints of the motors. Depending on the current measured (or estimated) state of the rear toe, the appropriate model is chosen among those presented in Sections 3.5 and 3.6 in formulating the maximization problem. The resulting feedback controller implicitly defines a switching law based on the physical state of the rear toe, with no explicit discrete internal states. On RHex, direct measurement of toe stiction is not possible and we instead use an empirically designed estimator, described in Section 5.3.

For both planar models, when the system is far from singular regions described in Section 3.7, the unknown forces and accelerations can be computed by directly solving eq. (3), yielding

$$\mathbf{v} = \mathbf{A}^{-1}(\mathbf{c})\mathbf{b}(\mathbf{c}, \dot{\mathbf{c}}) + \mathbf{A}^{-1}(\mathbf{c})\mathbf{D}(\mathbf{c})\ \boldsymbol{\tau} \ . \tag{13}$$

Both constrained systems are underactuated and direct inversion of these dynamics to obtain torque solutions is generally not possible. Furthermore, our task is not specified in terms of particular choices of ground reaction forces and accelerations. Rather, we are interested in the (in)stability properties of particular degrees of freedom in the system, particularly the body pitch, as well as various constraints arising from our assumptions in Section 3.1. As a consequence, our controller is based on a constrained optimization formulation informed by the underlying dynamics.

4.1. Constraints on Control Inputs

The first set of constraints in solving eq. (13) arises from physical limitations of the actuators in RHex. Torque limitations for the simplest, resistive model of a geared DC motor arise from the interaction between the back EMF voltage, the maximum available supply voltage and the armature resistance. Our model-based controller is designed to respect constraints based on this simple model for each motor, yielding decoupled torque limits in the form

$$\frac{K_\tau h_g(-v_s - \frac{\dot{\phi}_i}{m_g K_s})}{m_g(r_a + r_d)} < \tau_i < \frac{K_\tau h_g(v_s - \frac{\dot{\phi}_i}{m_g K_s})}{m_g(r_a + r_d)} \ , \tag{14}$$

where v_s is the supply voltage, r_a and r_d are the armature and drive resistances, K_s and K_τ are the speed and torque constants and finally, m_g and h_g are the gear ratio and efficiency. These limits clearly depend on the system state through the motor shaft velocities $\dot{\phi}_i$.

We introduce a second constraint on the control inputs to ensure that Assumption 3 holds. Our controllers must explicitly enforce body–ground contact throughout the progression of the remaining degrees of freedom. Fortunately, this requirement is easily captured through the constraint

$$F_c^z > 0 \,, \tag{15}$$

an inequality that is linear in the input torques as can be seen from the corresponding component of eq. (13).

DEFINITION 1. For a particular state $\mathbf{q} \in \mathcal{Q}$, we define the corresponding set of allowable torques, $\mathcal{T}(\mathbf{q})$ as the set of all torque input vectors $\boldsymbol{\tau} \in \mathbb{R}^3$ such that

$$\begin{aligned} & F_c^z(\mathbf{q}, \boldsymbol{\tau}) \geq 0 \\ \forall i, \quad & F_i(\mathbf{q}, \boldsymbol{\tau}) \geq 0 \\ \forall i, \quad & \tau_i^{min}(\mathbf{q}) \leq \tau_i \leq \tau_i^{max}(\mathbf{q}) \,. \end{aligned}$$

4.2. Maximal Thrust Control

For both models of Sections 3.5 and 3.6, the solutions for $\ddot{\alpha}$ and $\ddot{y}_t$ are continuous functions of the input torques. For any given state, this functional relationship is defined through our hybrid toe contact model and the solutions for the ground reaction forces, subject to the constraints described in the previous section. As a consequence, the problem of choosing hip controls to maximize thrust becomes a constrained optimization problem over the allowable input torque space.

DEFINITION 2. Given the current state $\mathbf{q} \in \mathcal{Q}$, we define the maximal torque input $\boldsymbol{\tau}^*$ as the torque vector that yields the maximum pitch thrust:

$$\boldsymbol{\tau}^*(\mathbf{q}) := \underset{\tau \in \mathcal{T}(\mathbf{q})}{\operatorname{argmax}}(\ddot{\alpha}(\mathbf{q}, \tau)).$$

Fortunately, the solutions of eq. (13) depend linearly on the input torques. Consequently, the constraints in Definition 1 as well as the objective function, $\ddot{\alpha}(\mathbf{q}, \boldsymbol{\tau})$, are linear in the input torques as well. As a result, standard linear programming techniques can be employed to identify efficiently the maximal torque solution $\boldsymbol{\tau}^*$. In particular, we use a simple geometric solution that exploits the low dimension and the largely decoupled structure of the constraints (Saranli 2002). Specifically, the motor torque limits of eq. (14) can be expressed as an axis-aligned constraint cube in 3-space, which we then intersect with the plane defined by the inequality constraint of eq. (15). The optimal solution can be obtained by simply evaluating the objective function on the vertices of the resulting intersection polygon as well as the corners of the cube.

In summary, we start by computing the unknown forces and accelerations for the current system state as affine functions of the torque inputs using eq. (13). It is then straightforward to construct the linear constraints of Definition 1. Finally, a geometric, computationally efficient algorithm is used to find the exact solution to the resulting linear programming problem, maximizing the thrust to the pitch degree of freedom while maintaining body–ground contact and respecting the limitations of the actuators. It is important to note that these computations are sufficiently simple as to be implemented in real time ($\sim$500 Hz) on the 300 MHz Pentium class processor used in RHex's control system.

4.3. Hybrid Energy Pumping

Depending on the frictional properties of the surface, our maximal thrust controller may or may not inject sufficient energy to complete the flip. In cases where it fails to achieve the sufficient energy level in the first attempt, our controller uses the same strategy as the first generation controller presented in Section 2.1. Once the body starts falling, local PD loops servo all legs to predetermined angles and wait until the collision of the front legs with the ground.

In order to recover as much of the impact kinetic energy as possible before the next thrust cycle, our controllers position the front leg vertically prior to impact, exposing the (passive) radial compliance of the leg to the bulk of the work performed. The vertical placement also avoids slippage of the leg as well as friction losses and, as noted above, eliminates the need for the motor to apply any torque during the collision due to the kinematically singular configuration. Moreover, during the decompression of the front leg, the middle and back legs can still apply additional thrust to inject energy even during the collision.

It would be possible to extend the continuous dynamics of Section 3 to incorporate compliance and other dynamical reaction forces of the front leg so as to construct a "stance phase" model that might then be integrated to obtain a more accurate prediction of the body kinetic energy returned at the next leg liftoff event. Examples of such predictive models can be found in the literature (Goldsmith 1960). However, their accuracy is still hostage to the difficulty of determining the dynamic properties of materials as well as other unmodeled effects (Chatterjee 1997; Chatterjee and Ruina 1998).

In consequence, we chose to incorporate a purely algebraic collision law in our model, where a single coefficient of restitution $-1 \leq k_r \leq 1$ summarizes the incremental effects of leg compression/decompression and additional thrust contributed by the middle and back legs during the decompression of the front legs. In doing so, we assume that no torque is applied to the front legs during the collision, constraining impulsive forces to act along the leg. Furthermore, we require that the impact occurs while the leg is within the friction cone to avoid

toe slip. Finally, we assume that the tail of the body comes to rest ($\dot{y}_b = 0$) during the fallback of the body, leaving the system with only one degree of freedom—the pitch, α. In light of these assumptions, we use the algebraic law

$$\dot{\alpha}^+ = -k_r\, \dot{\alpha}^- \,, \tag{16}$$

relating the pitch prior to and following the collision ($\dot{\alpha}^-$ and $\dot{\alpha}^+$, respectively) to verify that the resulting impulsive forces on the body do not cause liftoff of the tail. This yields appropriate initial conditions for the subsequent thrust phase. Again, empirical evidence reveals that these simplifying assumptions approximate well the physical behavior that RHex exhibits in the vast majority of circumstances.[2] As our experimental platform has no means for detecting tail liftoff and subsequent compensation, we use a conservative choice for the front leg angle prior to impact to minimize chances for such an event.

5. Experimental Results

5.1. Experimental Platform

The most recent version of the robot, RHex 1.5, adopted for the present experiments, has a rigid body that measures $50 \times 20 \times 15$ cm^3, and houses all the computational and motor control hardware, including batteries and two PC104 stacks for control and vision tasks. Each leg is directly actuated by a Maxon RE118751 20W brushed DC motor combined with a Maxon 114473 two-stage 33 : 1 planetary gear (Interelectric AG 1997/98), delivering an intermittent stall torque of 6 Nm at 24 V. The total weight of the robot is roughly 8.5 kg.

In contrast to the earlier versions, RHex 1.5 incorporates a three-axis gyro for inertial sensing of the body orientation in addition to the motor encoders. Recently developed behaviors on RHex increasingly rely on accurate estimation of the spatial body orientation. As a consequence, we use a quaternion representation together with integration of gyro readings at 300 Hz to implement a singularity-free robust estimator, supporting flipping as well as other inertially guided locomotion primitives (Skaff et al. 2003). Furthermore, for the flipping behavior in particular, we minimize gyro drift by resetting orientation estimates following each collision when it is possible to compute the robot orientation through kinematics.

The legs on the current RHex are monolithic pieces of compliant fiber glass, attached to motor shafts through aluminum hip fixtures. As the third design iteration on possible leg materials and morphologies, they exhibit significantly improved reliability and compliance characteristics (Moore et al. 2002). Each of the legs in the set used for the experiments in this paper is roughly 16.5 cm long, weighs 80 g and has a radial compliance[3] of 1900 N m^{-1}.

2. Due to the lack of sensing of the translational body coordinates in RHex 1.5, our only evidence for this observation comes from qualitative analysis of video footage from flipping experiments.

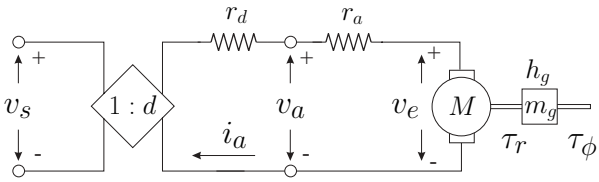

Fig. 6. Simple motor model for RHex.

5.2. Approximate Model-Based Torque Control

RHex 1.5 does not have hardware support for controlling the hip actuators in torque mode. Due to space and power limitations, the current design operates in voltage mode, where commands from the on-board computer drive an H-bridge amplifier, whose output is then fed through a low-pass filter and connected directly to the armature of the DC motors at each hip. In order to implement our flipping controller, we use an inverse model of the motor amplifier—a slight variant of the model described in McMordie, Prahacs, and Buehler (2003), with less than 6% prediction error—to achieve reasonably accurate control of the hip torque.

Figure 6 illustrates our simple model for the amplifier and motor stages. We assume that the combination of the PWM amplifier with the LC filter can be approximated with a passive resistor on the load side in series with an ideal transformer, whose duty factor $d \in [-1, 1]$ can be arbitrarily commanded. Furthermore, we assume a simple resistive model for the motor, followed by a gear head with a reduction ratio of m_g and efficiency of h_g.

In order to obtain the desired torque on the output shaft, the commanded duty factor must be chosen to yield an appropriate armature current. A straightforward solution of the circuit in Figure 6 yields

$$d = \frac{m_g (r_a + r_d)\tau_\phi^*}{h_g K_\tau v_s} + \frac{\dot{\phi}_i}{m_g K_s v_s} \tag{17}$$

as the input command to the PWM amplifier.

Table 3 summarizes parameter values for RHex's motors. In all the experiments reported in this paper, we used a 1 kHz software loop to implement eq. (17), which yields a steady state RMS error of around 6% between the actual and desired hip torques (McMordie, Prahacs, and Buehler 2003).

3. Even though compliance is critical in RHex's dynamic locomotion performance, it is not nearly as dominant for the flipping behavior. Only small discrepancies are introduced in the leg length due to the radial compliance and the accuracy of the torque control suffers small delays due to the rotational compliance of RHex's legs.

Table 3. Parameters for RHex's Hip Motors

r_d	0.45 Ω	K_τ	0.01763 N m A^{-1}
r_a	1.65 Ω	m_g	0.03 leg/mtr
K_s	59.21 V s rad^{-1}	h_g	0.8 leg/mtr

5.3. Detection of Toe Stiction

Our switching controller design requires measurement of rear toe velocity to determine which of the models in Sections 3.5 and 3.6 is to be used for the maximal torque controller. However, RHex is not equipped with sufficiently accurate inertial sensors to estimate the translational velocity of the body. Furthermore, our crude model of the robot kinematics results in further errors in the transformation to the leg states, rendering detection of toe stiction through estimation of body velocity infeasible.

Nevertheless, high-speed video footage (see Extension 3) of flipping on various surfaces reveals that the behavior of the rear toe is very consistent and regular across different experiments. As the front and middle legs leave the ground, the rear toes start sliding forward. Briefly after the liftoff of the middle legs, the motion of the rear toes come to a stop. For the remainder of the thrust, they only move intermittently and exhibit a stick–slip style low-frequency chattering due to the passive compliance in the legs. This sticking of the rear toes is also consistently marked by a relatively sharp increase in the pitch acceleration.

Motivated by these observations, we devised a filter for the pitch acceleration measurements as a mechanism for detecting toe stiction. Beyond a certain pitch value where both the front and middle legs are in flight, we switch to the "stuck" toe model when there is a "sharp" change in the pitch acceleration. Our sharpness measure is based on manually tuned threshold parameters, which are specific to each surface. In the future, we plan to incorporate estimates of body velocity as well as contact sensors on the legs, which should eliminate the need for this filter and the associated manual tuning.

5.4. Thrust Phase Model Performance

In this section, we present experimental data to establish the baseline performance of our thrust controller on linoleum, a slippery surface with relatively consistent frictional properties.

For the experiments presented in this section, we fixed the friction coefficient for the body contact as $\mu_b = 0.4$, based on ranges indicated in Marks (1996) for plastic on linoleum type surfaces. This parameter has very little effect on controller performance due to the small associated ground reaction force enforced by the controller.

In contrast, in order to estimate the much more important toe friction coefficient and assess the corresponding model performance, we ran a number of experiments using approximate measurements of RHex's kinematic and dynamic parameters (see Table 2), for different settings of the toe friction coefficient μ_t. Four runs were recorded for 15 different settings in the range $\mu_t \in [0.1, 0.6]$ (see Extension 4 for experimental data).

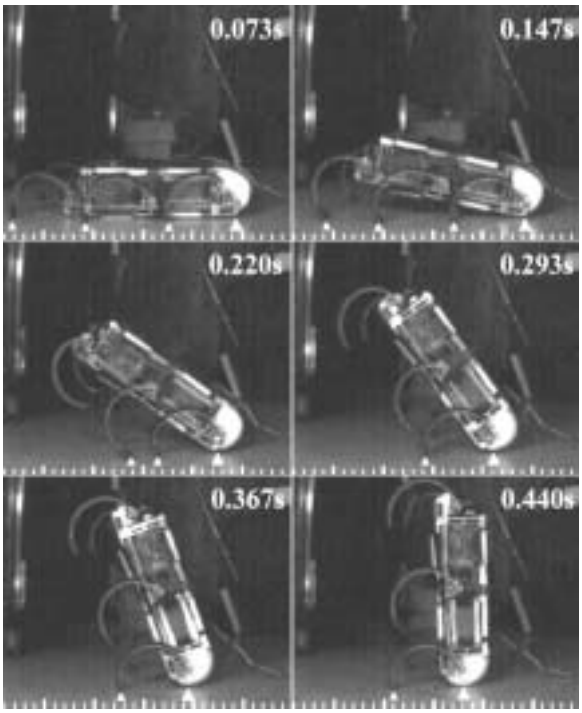

Fig. 7. Sequence of snapshots during flipping on linoleum with $\mu_t = 0.39$, for the period where the maximal thrust controller is active. Subsequent frames in which the robot falls back and recovers are not included. White arrows indicate contact points for the toes and the body.

Figure 7 displays a sequence of snapshots for one of these experiments, extracted from the high-speed video footage of Extension 3. We only included the most relevant part of the experiment, which is the period where the maximal thrust controller was active. A number of important details are illustrated by these snapshots. First, the tail of the robot remains in contact with the ground throughout the whole run, which indicates the controller's success in maintaining its constraints. Secondly, throughout flipping, the front two legs slide backwards, whereas the rear legs and the body contact are sliding forward, which justifies our static assumptions for the directions of frictional forces to yield $\bar{\mu}_i$. Finally, even though it is hard to recognize in the snapshots alone, rear toe stiction and chattering in the second half of flipping is evident in the associated video footage of Extension 3.

For each experiment, we logged the pitch rate and angle measurements. We computed the model prediction for the

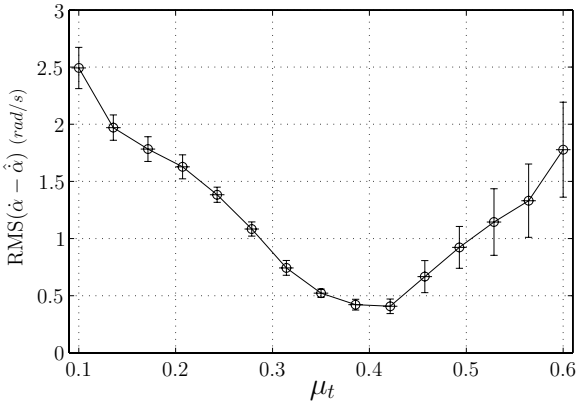

Fig. 8. Mean RMS error in estimated pitch velocity as a function of the toe friction coefficient. For each data point, vertical bars indicate standard deviation across four experiments. The script PlotRMSError.m of Extension 4 can be used to generate this figure.

pitch acceleration using eq. (13), subsequently integrated to obtain the model predictions for the pitch rate.[4] The results are then compared with the actual measurements to characterize the accuracy of the model predictions.

Figure 8 plots the RMS error between the measured and predicted pitch rates for different friction coefficients. For each setting, the mean and standard deviations across four repeated experiments are shown. The best model performance is obtained for $\mu_t \approxeq 0.4$. Considering various levels of approximations used in our model, including the inaccuracies in the kinematic and dynamic parameters, unmodeled compliance in RHex's legs and the approximate software torque control, the RMS error of 0.4 rad s^{-1} is surprisingly small—less than 10% of the maximum speed attained during the flip.

Figure 9 illustrates the best run with $\mu_t = 0.39$, resulting in an RMS error in the pitch rate prediction of 0.4 rad s^{-1}. Figure 10 portrays a similar comparison between the actual robot performance and a pure simulation with the same initial conditions and model parameters (see Extension 4 for data). The reader must bear in mind two qualifications in comparing these two figures. First, the simulation uses a fixed supply voltage at 23.5 V, roughly modeling the average battery voltage drop during the experiments. Secondly, for the simulated model, we replace the model switching logic of Section 5.3 with a direct measurement of the toe velocities, yielding proper detection of toe stiction.

These results suggest that our model provides an accurate quantitative representation of the thrust phase. Nevertheless, there are a number of inaccuracies in its prediction, mainly visible in the pitch acceleration plots. Most significantly, our model fails to predict the large overestimation of the initial acceleration and the subsequent, relatively large oscillations in the measured acceleration.

We believe that the origin of both discrepancies is the compliance in RHex's legs. The initially uncompressed legs introduce some delay in responding to the torque commands, resulting in a small delay in the measured pitch acceleration. A similar effect is visible subsequent to the liftoff of the middle legs which also causes oscillations due to the sudden loading of the rear legs.

Nevertheless, most of these differences do not significantly influence the average performance of the model prediction. In addition to the accurate prediction of the pitch velocity, the robot successfully keeps its tail on the ground and consistently performs successful flips on linoleum in a single thrust (see Figure 7 and Extension 3).

For a better understanding of the controller performance, it is also useful to look at the resulting motor commands. Figure 11 illustrates torque outputs of combined pairs of rear, middle and front legs estimated using the motor model of eq. (17). Throughout the first phase, when all the legs are in contact with the ground, the controller applies maximum available torque to all the legs, which decreases as the motor shafts rotate faster. Following the liftoff of the front legs, the body ground reaction force constraint becomes dominant and the torque output to the rear legs is constrained to avoid tail liftoff. This continues through the liftoff of the middle legs, until the detection of rear toe stiction (see Section 5.3). The underlying model is then switched to that of Section 3.6. Towards the end of the flip, the tail liftoff constraint becomes dominant once again due to centrifugal forces.

5.5. Multi-Thrust Flipping on Linoleum

In this section, we present experimental data on flipping through multiple thrust phases. As in the previous section, we use a linoleum surface for these experiments due to the consistency of its frictional properties.

On linoleum, our maximal thrust controller on RHex always performs a backflip in a single thrust phase. Consequently, we artificially scale the torque limits of eq. (14) to decrease the injected energy during each cycle necessitating multiple thrust phases to flip the robot over. In particular, we decrease the torque limits for the front, middle and rear leg pairs by 90%, 20% and 20%, respectively, from the actual constraint computed using eq. (14). The larger reduction of the front leg torque limits is intentional and significantly decreases the energy injected in the first thrust phase—the only time period in which front legs do any active work. As a result, the first thrust is forced to fail, necessitating additional cycles for a successful flip.

We ran 10 experiments on linoleum with these parameter settings, all of which successfully flipped RHex after three thrust phases (see Extension 4 for experimental data). A

4. Note that the integral of the predicted acceleration throughout the experiment is not the same as a pure simulation of the model dynamics.

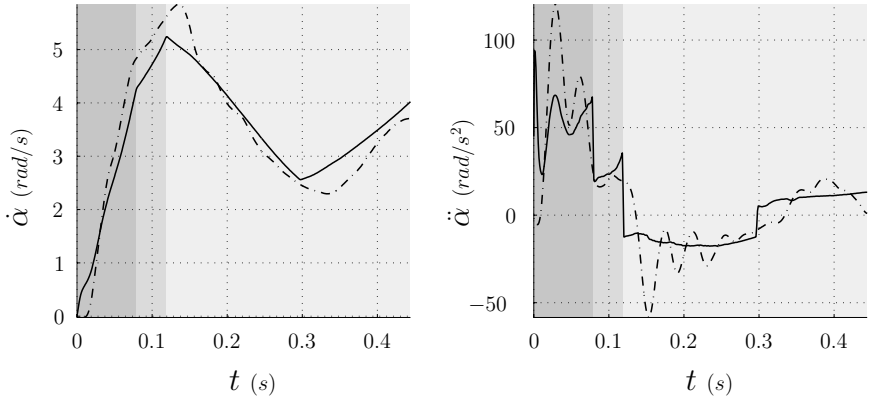

Fig. 9. Model predictions for the pitch velocity (left) and acceleration (right) compared to the experiment with best model performance ($\mu_t = 0.39$). Solid lines indicate model prediction whereas dashed lines show the actual measurements. Shaded regions indicate different number of legs in contact with the ground: dark (3), middle (2) and light (1).

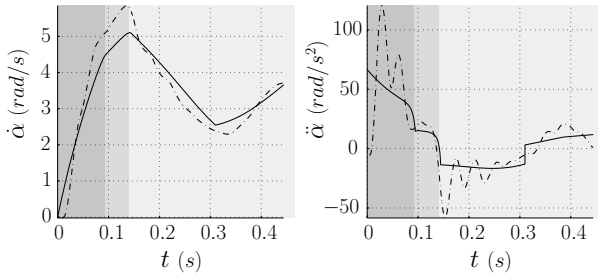

Fig. 10. Pure simulation predictions for pitch velocity (left) and acceleration (right) compared to the experiment with best model performance ($\mu_t = 0.39$). Solid lines indicate simulated model trajectories whereas dashed lines show the actual measurements. Shaded regions indicate different number of legs in contact with the ground: dark (3), middle (2) and light (1). The script PlotSimulationData.m of Extension 4 can be used to generate this figure.

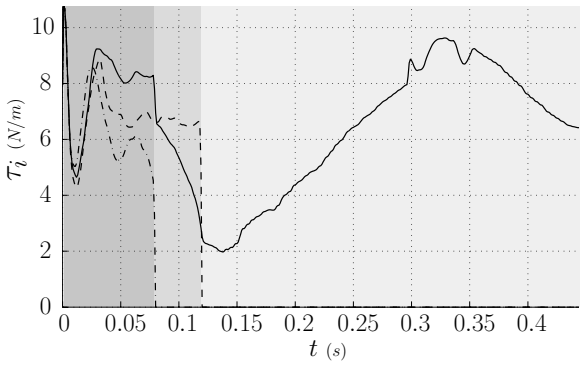

Fig. 11. Hip torque commands for the experiment with best model performance ($\mu_t = 0.39$). Solid, dashed and dash-dotted lines correspond to rear, middle and front leg pairs, respectively. Shaded regions indicate different number of legs in contact with the ground: dark (3), middle (2) and light (1).

representative experiment is illustrated in Figure 12, where the first thrust phase failed to flip the robot over and two more cycles were needed before enough energy was injected into the system. In the following paragraphs, we present a number of features of Figure 12 that were consistently observed in all 10 experiments.

Our first observation is that the coefficient of restitution for the collision of the front legs with the ground is approximately 0.67. Even though this coefficient is primarily a function of the material properties of the front legs, there is also some contribution from the middle and rear legs. In particular, our approximate kinematic model results in premature contact of middle and rear legs with the ground during the compression of the front leg, resulting in undesired negative work. In contrast, the maximal thrust controller is engaged as the pitch velocity changes sign, yielding additional thrust from the middle legs prior to the end of the collision.

Two aspects of the velocity plot are also important to note. First, there is a significant deceleration at the beginning of the second and third thrust phases (indicated by horizontal arrows in Figure 12), which is not present in the very first thrust. This is once again primarily due to the compliance in the legs, which introduces a delay in the action of torque commands through the legs. As the middle and rear leg springs are very close to their rest positions following the collision, the body does not immediately feel the torque commands acted by the maximal thrust controller, resulting in a brief deceleration.

The second aspect concerns the sharp drops in the velocity following the end of each thrust cycle (indicated by the vertical arrows in Figure 12), as the robot body starts to fall towards

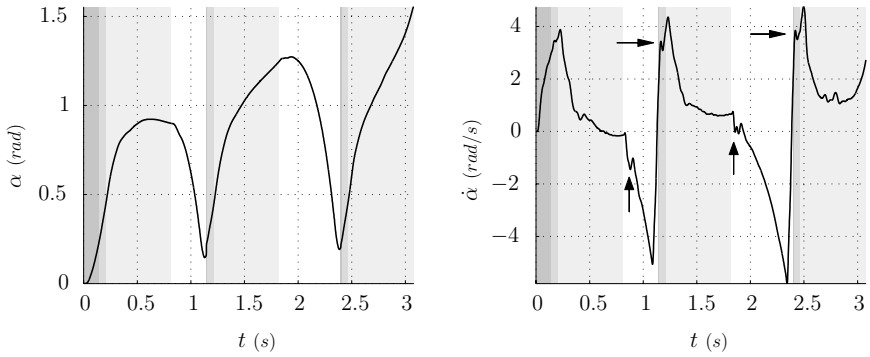

Fig. 12. Body pitch (left) and pitch velocity (right) for one of 10 flipping experiments on linoleum with torque limits for the front, middle and rear legs reduced by 90%, 20% and 20%, respectively. Three thrust phases are needed for a successful flip. Shaded regions indicate different number of legs in contact with the ground: dark (3), middle (2) and light (1). The thrust controller is inactive in white regions with the legs clear from the ground in preparation for the collision. The script PlotBounceData.m of Extension 4 can be used to generate this figure.

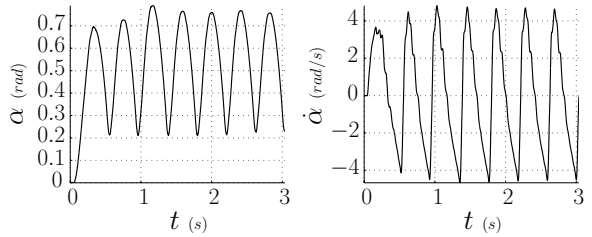

Fig. 13. Body pitch (left) and pitch velocity (right) for flipping on linoleum with the open-loop strategy. This controller always fails to flip RHex 1.5. The script PlotOpenloop.m of Extension 4 can be used to generate this figure.

the rebound. As described in Section 4.3, this is when local PD controllers for each leg are activated to bring them to fixed angles in preparation for the collision. Even though the legs have relatively small mass (0.48 kg total) compared to the robot body, their sudden movement effects the body velocity.

5.6. Performance of the Open-Loop Controller

As noted in Section 2.2, our first generation flipping controller was primarily open loop at the task level, with only local feedback at the hips for PD control of the motor angles and crude detection of maximum pitch at each thrust. With this simple initial design, we were able to achieve successful flipping maneuvers over a reasonable range of surfaces on an earlier version of RHex (Saranli and Koditschek 2001, 2002). However, with new sensors, improved computational hardware as well as structural ruggedization, RHex's newest version, 1.5, is heavier and slightly larger, resulting in consistent failure of the open-loop controller.

Figure 13 illustrates an example on linoleum, which is one of the least challenging surfaces for flipping with its low friction (see Extension 4 for experimental data). For this experiment, we manually tuned motor gains and leg trajectories in an attempt to gain as much thrust as possible while keeping the tail on the ground. Invariably, the unavailability of system state forced our tuning to be overly conservative, making it impossible to harvest maximum performance. As a consequence, our simple first generation controller consistently fails to flip over even after several thrust phases, unable to exceed a maximum pitch value of around 0.75 rad.

5.7. Flipping on Rugged Surfaces

The final set of experiments we present characterize the performance of model-based flipping on a number of indoor and outdoor surfaces. Figure 14 illustrates each of these surfaces with snapshots from associated experiments. The associated video footage is also included in Extensions 5, 6, 7, 8 and 9.

This section presents two families of experiments: flipping on hard surfaces with consistent frictional properties—carpet, asphalt and packed dirt—and flipping on soft outdoor surfaces—thin and thick grass—for which quantitive characterization is much less feasible due to the high level of inconsistency and variation in surface properties across experiments.

For the experiments on hard surfaces, we ran sets of three experiments for 10 different settings of the leg/ground coefficient of friction. As in Section 5.4, we compared model predictions to measured performance to identify the frictional properties of each surface. The first three rows of Table 4 summarize the coefficients of friction we identified for each hard surface, as well as the performance of model based flipping with the identified parameter in terms of the percentage of success and the number of required thrusts.

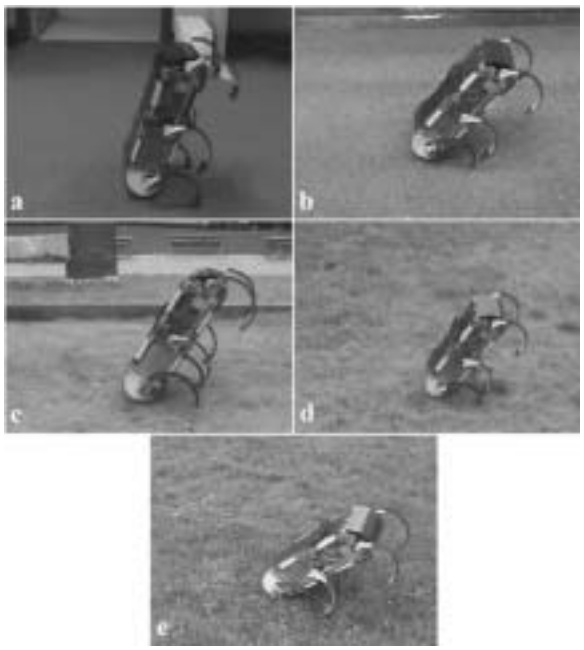

Fig. 14. Snapshots of flipping on various rough surfaces: (a) carpet; (b) asphalt; (c) packed dirt; (d) thin grass; (e) thick grass.

Table 4. Model-Based Flipping Performance on Hard and Soft Rough Surfaces

Terrain	μ_t	Success	Thrusts
Packed dirt	0.55	100%	1
Asphalt	0.6	100%	1
Carpet	0.7	100%	1
Thin grass	–	90%	1
Thick grass	–	0%	–

See Extensions 5, 6, 7, 8 and 9 for sample movies and Extension 4 for experiment data.

On all hard surfaces, the robot successfully flipped over for each attempt. In all cases, no more than a single thrust phase was required. For these experiments, we did not explicitly tune the detection algorithm for toe stiction, so there was significant variability in controller performance following the liftoff of the middle legs (see Extension 4). As a consequence, in identifying the friction coefficient for each surface, we only considered errors in the pitch rate estimation prior to the switching of the model. In practice, the effect of the model discrepancy on the flipping performance beyond the liftoff of the middle legs turns out to be not significant as the motor torque limits dominate over the remaining constraints.

Characterization of flipping performance on soft surfaces is much more challenging due to their inherent irregularity. Unlike the hard surfaces, more than two or three experiments on the same location result in the legs digging in the grass, changing the associated frictional properties. Furthermore, it is unreasonable to hope that the Coulomb friction model will be accurate in modeling the sliding of the legs in thick grass, which usually results in wedging and other unpredictable outcomes. Consequently, we only ran 10 experiments on both thin and thick grass and did not attempt to identify the frictional properties of these surfaces. The second part of Table 4 presents the success percentages of these experiments.

Not surprisingly, thick grass presents a significant challenge and the robot is not able to flip even with multiple thrust attempts (see Extension 8). On the other hand, thinner grass is much less demanding and usually admits flipping in a single thrust (see Extension 9). This is a significant improvement over the first generation open-loop controller, which was incapable of inducing a flip even with the lighter and smaller RHex 0.5.

6. Conclusion

In robotic locomotion research, autonomy is likely to impose some of the most demanding constraints on design and limitations on behavior. It is very difficult, often impossible, to achieve in systems otherwise designed for non-autonomous operation. RHex, our hexapedal platform, demonstrated that autonomy as a design goal can achieve significant advances in real-world performance and robustness.

In this paper, we present a new controller to implement self-righting behavior on RHex, which is perhaps the simplest instance of self-manipulation other than locomotion itself. Our modeling and analysis yields significant improvements to the simple first generation controller, extending its domain of success to a wider range of terrain conditions. We present empirical evidence to verify the validity of our model and to document the performance of a new model based controller. We show that the maximal thrust controller we introduce performs successful flipping maneuvers on linoleum, carpet, packed dirt, asphalt and thin grass, usually with only a single thrust phase. We also demonstrate an "energy pumping" scheme, designed to handle disturbances or terrain conditions which may induce the failure of the first thrust attempt. In each case, we present empirical evidence to compare model predictions to actual measurements of robot performance.

The design of the new model-based controller makes a few simplifying assumptions to make a real-time implementation feasible. Relaxation of these assumptions through more formal analysis of the preliminary model we described in this paper is also of great interest. Extensions of the flipping behavior such as uninterrupted rolling or handstands will require a much better analytical understanding of the model.

We believe that such extensions to the behavioral suite of a morphology as limited as RHex is the best way to address the shortcomings of contemporary actuation and energy storage technology while continuing to press ahead in the development of practically useful robots.

Appendix: Index to Multimedia Extensions

The multimedia extension page is found at http://www.ijrr.org.

Table of Multimedia Extensions

Extension	Type	Description
1	Video	**OpenLoopSingleShot.mpg**. RHex 0.5 flipping on linoleum with the original open-loop controller.
2	Video	**OpenLoopCarpet.mpg**. RHex 0.5 flipping on carpet with the original open-loop controller. Multiple thrusts are required for a successful flip.
3	Video	**LinoleumSingleShot.mpg**. High-speed video (150fps) of RHex 1.5 flipping on linoleum with the model based controller.
4	Data	**data_scripts.tar.gz**. Data files and visualization scripts for all the experiments and simulations. Please see README.txt in this archive for details on the format of data files and the usage of associated scripts.
5	Video	**ModelBasedCarpet.mpg**. RHex 1.5 flipping on carpet with the model-based controller.
6	Video	**ModelBasedDirt.mpg**. RHex 1.5 flipping on packed dirt with the model-based controller.
7	Video	**ModelBasedAsphalt.mpg**. RHex 1.5 flipping on asphalt with the model-based controller.
8	Video	**ModelBasedThickGrass.mpg**. RHex 1.5 failing to flip on thick grass with the model-based controller.
9	Video	**ModelBasedThinGrass.mpg**. RHex 1.5 flipping on thin grass with the model-based controller.

Acknowledgments

We thank Professor Matt Mason for his insight on friction models and the connection to the Painlevé problem. This work was supported in part by DARPA/ONR Grant N00014-98-1-0747.

References

Altendorfer, R., Moore, N., Komsuoglu, H., Buehler, M., Brown H.B. Jr, McMordie, D., Saranli, U., Full, R.J., and Koditschek, D.E. 2001. RHex: a biologically inspired hexapod runner. *Autonomous Robots* 11:207–213.

Baraff, D. 1991. Coping with friction for non-penetrating rigid body simulation. *Computer Graphics* 25(4):31–40.

Bares, J.E., and Wettergreen, D.S. 1999. Dante II: technical description, results and lessons learned. *International Journal of Robotics Research* 18(7):1–29.

Chatterjee, A. 1997. *Rigid Body Collisions: Some General Considerations, New Collision Laws, and Some Experimental Data*. PhD thesis, Cornell University.

Chatterjee, A., and Ruina, A. 1998. Two interpretations of rigidity in rigid body collisions. *Journal of Applied Mechanics* 65(4):894–900.

Fiorini, P., and Burdick, J. 2003. The development of hopping capabilities for small robots. *Autonomous Robots* 14(2–3):239–254.

Geng, T., Li, X., and Xu, X. 2002. Control and simulation of a 3D one-legged robot. *Proceedings of the American Control Conference*, Anchorage, AK, May, Vol. 3, pp. 2499–2504.

Geng, T., Li, X., Shen, G., and Xu, X. 2002. Ballistic flip of a planar one-legged robot in simulation. *Proceedings of the American Control Conference*, Anchorage, AK, May, Vol. 3, pp. 2505–2509.

Goldsmith, W. 1960. *Impact: The Theory and Physical Behavior of Colliding Solids*. Edward Arnold, London.

Hale, E., Schara, N., Burdick, J.W., and Fiorini, P. 2000. A minimally actuated hopping rover for exploration of celestial bodies. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, April 24–28, pp. 420–427.

Hodgins, J., and Raibert, M.H. 1990. Biped gymnastics. *International Journal of Robotics Research* 9(2):115–132.

Interelectric AG. 1997/98. *Maxon Motor Catalog*, Sachseln, Switzerland (http://www.maxon.com).

McMordie, D., Prahacs, C., and Buehler, M. 2003. Towards a dynamic actuator model for a hexapod robot. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Taipei, Taiwan.

Marks, L.S. 1996. *Marks' Standard Handbook for Mechanical Engineers*. McGraw-Hill, New York.

Mason, M.T. 2001. *Mechanics of Robotic Manipulation*. MIT Press, Cambridge, MA.

Matthies, L., Xiong, Y., Hogg, R., Zhu, D., Rankin, A., Kennedy, B., Hebert, M., Maclachlan, R., Won, C., Frost, T., Sukhatme, G., McHenry, M., and Goldberg, S. 2000. A portable, autonomous, urban reconnaissance robot. *Proceedings of the 6th International Conference on Intelligent Autonomous Systems*, Venice, Italy, July.

Moore, E.Z., Campbell, D., Grimminger, F., and Buehler, M. 2002. Reliable stair climbing in the simple hexapod RHex. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Washington, DC, May 11–15, Vol. 3, pp. 2222–2227.

Nakanishi, J., Fukuda, T., and Koditschek, D.E. 1999. A hybrid swing up controller for a two-link brachiating robot. *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, Atlanta, GA, September, pp. 549–554.

Nakanishi, J., Fukuda, T., and Koditschek, D.E. 2000. A brachiating robot controller. *IEEE Transactions on Robotics and Automation* 16(2):109–123.

Nelson, G.M., and Quinn, R.D. 1999. Posture control of a cockroach-like robot. *IEEE Control Systems Magazine* 19(2):9–14.

Painlevé, P. 1895. Sur le lois de frottement de glissement. *Comptes Rendus Academie des Sciences Paris* 121:112–115.

Raibert, M. 1986. *Legged Robots That Balance*, MIT Press Series in Artificial Intelligence. MIT Press, Boston, MA.

Saranli, U. 2002. *Dynamic Locomotion with a Hexapod Robot*. PhD thesis, Computer Science and Engineering, University of Michigan.

Saranli, U., and Koditschek, D.E. 2001. Design and analysis of a flipping controller for RHex. Technical Report CSE-TR-452-02, UM, Ann Arbor, MI.

Saranli, U., and Koditschek, D.E. 2002. Back flips with a hexapedal robot. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Washington, DC, May 11–15, Vol. 3, pp. 128–134.

Saranli, U., Buehler, M., and Koditschek, D.E. 2001. RHex: a simple and highly mobile robot. *International Journal of Robotics Research* 20(7):616–631.

Skaff, S., Kantor, G., Maiwand, D., and Rizzi, A.A. 2003. Inertial navigation and visual line following for a dynamical hexapod robot. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, NV, October 27–31.

Spong, M.W. 1995. The swing up control problem for the acrobot. *IEEE Control Systems Magazine* 15(1):49–55.

Stewart, D.E. 1998. Convergence of a time-stepping scheme for rigid-body dynamics and resolution of Painlevé's problem. *Archive for Rational Mechanics and Analysis* 145(3):215–260.

Stewart, D.E., and Trinkle, J.C. 1997. Dynamics, friction, and complementarity problems. *Complementarity and Variational Problems*, M.C. Ferris and J.S. Pang, editors. SIAM, Philadelphia, pp. 425–439.

Tunstel, E. 1999. Evolution of autonomous self-righting behaviors for articulated nanorovers. *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Noordwijk, the Netherlands, June, pp. 341–346.

Waldron, K.J., and Vohnout, V.J. 1984. Configuration design of the adaptive suspension vehicle. *International Journal of Robotics Research* 3(2):37–48.

Wang, Y., and Mason, M. 1992. Two-dimensional rigid body collisions with friction. *Transactions of the ASME, Journal of Applied Mechanics* 59(3):635–642.

Yoo, K.J., Yang, D.H., and Hong, S.K. 2001. Swing up and stabilization control of the pendubot. *Proceedings of the International Conference on Control, Automation and Systems*, Jeju Island, Korea, October, pp. 2746–2749.

Zhang, W., Wang, G., Chambers, T., and Simon, W.E. 1997. Toward a folding-legged uniped that can learn to jump. *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, Orlando, FL, October, Vol. 5, pp. 4315–4319.