

Design and Analysis of a Flipping Controller for RHex

Uluç Saranlı*, and Daniel E. Koditschek

Artificial Intelligence Laboratory
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, MI 48109-2110, USA
{ulucs, kod}@eecs.umich.edu

Abstract

We report on the design and analysis of a controller that can achieve dynamical self-righting of our hexapedal robot, RHex. We present an empirically tuned controller that works reasonably well on indoor surfaces, using a hybrid energy pumping strategy to overcome torque limitations of its actuators. Subsequent modeling and analysis yields a new controller with a much wider domain of success as well as a preliminary understanding of the hybrid control strategy. Simulation results demonstrate the superiority of the improved control strategy relative to the first generation empirically designed controller.

1 Introduction

RHex is an autonomous hexapod robot that negotiates badly irregular terrain at speeds better than one body length per second [8]. 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 “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. A comparative simulation study suggests that the new controller will extend significantly the terrain over which the self-righting maneuver succeeds.

Autonomy is an essential component for any robotic platform designed to operate in the real world. In addition to strict power and computational constraints it imposes on the design, it also requires basic self-manipulation capabilities for the survivability of the robot in the absence of a human operator. Even during teleoperation, where the computational demands on the platform are less stringent, the ability to recover from unexpected situations through self-manipulation is essential. Application such as planetary rovers and space missions are among the best examples where these requirements are most critical [1].

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 completely render a robot useless. Especially in outdoor environments with badly broken terrain and obstacles of various shapes and sizes, the debilitating effects of such accidents have been observed in the past [2].

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 [7]. However, various scenarios such as teleoperation and vision based navigation entail a nominal orientation as a result of the accompanying instrumentation and algorithms. Under these constraints, most existing robotic designs with self-righting capabilities incorporate special kinematic structures such as long extension arms or reconfigurable wheels [6, 9]. In consequence of weight and power limitations, RHex is not equipped with such structures and must rely on its existing morphology together with dynamic maneuvers to perform a flip-over.

Through an empirical design process, we have been able to construct a controller capable of inducing a dynamical back flip in RHex. The applicability of this controller, however, is limited to a small range of simple surfaces, such as linoleum, smooth concrete and carpet, outside which the robot cannot successfully flip over. We present in this paper a new model and a controller derived from it that promises to extend significantly the variability of terrain over which the desired behavior will be achieved. Beyond the existing behavior and the new multiple point collision/contact model, the main contributions of the paper are a torque control strategy that maximizes the energy injected into the system, validated for now by a comparative simulation study and a second empirical study, presently in progress.

2 Flipping RHex

2.1 Simple Open Loop Control

Our first generation flipping controller consists of a state machine (see Figure 1). Starting from a stationary position on the floor, the robot very quickly (in 0.2s) goes through two configurations (poses I and II in Figure 1), with front and middle legs successively leaving the ground. Depending on the frictional properties of the ground, these motions result 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. However, on most surfaces — gravel, grass and asphalt, but even some indoor settings such as carpet — this is not sufficient to flip the body over. Instead, the robot reaches some maximum pitch lying within the basin of the original configuration, and the robot falls back toward its initial state. Under these circumstances, the controller brings the legs back to Pose I of Figure 1 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 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. Thrusting is achieved by running a high gain proportional derivative control (PD)

concrete	linoleum	carpet
90%	100%	90%
asphalt	grass	gravel
100%	0%	0%

Table 1: Success rates of the first generation controller for 10 experiments each on different surfaces.

law around a judiciously selected constant velocity leg sweep motion. The maximum pitch attained by the body increases with each bounce up until the point where collision losses exceed the energy that can be imparted by the PD controller during the leg sweep phase interval. As Table 1 suggests, this pumping strategy works very reliably on a number of common surfaces such as linoleum, smooth concrete, carpet and asphalt.




Figure 1: Sequence of states for the flipping controller

However, on many surfaces — the outdoor environments most relevant to RHex’s presumed mission [1, 8] such as loose gravel, grass and soft ground — it does not perform nearly as well. To permit a reasonable degree of autonomous operation, we would like to improve on the range of conditions 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 hip motors results in the body lifting off the ground into stance mode, still in the wrong configuration. We require a strategy that can be tuned carefully enough 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.

2.2 Assumptions and Control Design Constraints

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

A 1 *The flipping behavior is primarily planar.*

The controller described in the previous section operates contralateral pairs of legs in synchrony. The robot’s response lies almost entirely in the sagittal plane, and departures are rare enough to be negligible.

A 2 *The leg masses are negligible in flight.*

We assume that the leg masses are sufficiently small so that their effect on the body dynamics when they are not in contact with the ground is negligible. However, we still need to consider the masses of the stance legs in formulating the dynamics in order to avoid peculiar singularities arising from the nonzero hip torque actuation.

A 3 *The tail of the body should keep contact with the ground throughout the flipping action.*

This assumption is motivated by a number of observations gathered during the empirical flipping experiments presented above. First, during the initial thrust phases, the front legs provide most of the torque. The configurations where the tail endpoint of the body is in contact with the ground yield the longest duration of contact for the front leg, harvesting the most possible benefit from the associated actuator until it leaves the ground.

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, 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.

In light of these assumptions, our controller designs face two major constraints: keeping the tail endpoint of the body on the ground and respecting the torque limitations of the actuators.

3 The Planar Flipping Model

3.1 A Generic Planar Model

Based on A3, our analysis will be largely confined to configurations where the tail of the body is constrained to lie on the ground. In this section, we describe a slightly more generic model to prepare a formal framework in which we define these constraints. Section 3.3 then presents the much simpler, single degree of freedom model that will be used in our algorithm design and subsequent analysis (presently in progress).




Figure 2: Simple rigid planar model of RHex

Figure 2 illustrates our unconstrained planar model. Three rigid legs with point masses m on the toes are attached to a rigid body with mass M and inertia I_b . The toe masses are only effective when the leg is touching the ground and are neglected when the leg is in flight. The attachment points of the legs are fixed, along a straight line through the center of mass. This line also defines the orientation of the body, θ , with respect to the horizontal. The body extends between the points N and T , defined to be the “nose” and the “tail”, respectively. The center of mass is midway between N and T , resulting in a body length of $2d$.

The nose, the tail and the toes cannot penetrate the ground. We assume that the body-ground friction is infinite thereby precluding any possibility of horizontal slip of the tail and the nose. In contrast, horizontal motion of the toes along the ground is central to the behavior of interest. We model the toe ground interaction as characterized by Coulomb friction with dynamic coefficient μ and viscous friction with damping constant k_d . Table 2 summarizes the notation used throughout the paper.

States and dependent variables	
$\theta \in \mathcal{Q}$	Body angle wrt horizontal (ccw)
$\mathbf{q} := [\theta, \dot{\theta}] \in T\mathcal{Q}$	State vector
b_x, b_z	Coordinates of the body COM
ϕ_i	Leg angle wrt the body vertical
γ_i	Leg angle wrt the horizontal (cw)
x_i	Horizontal position of the toe mass
Control inputs	
$\tau \in \mathbb{R}^3$	Hip torque control vector
$\mathcal{T}(\mathbf{q}, p)$	Set of allowable torque vectors
Model parameters	
d, d_i, l	Kinematic parameters
M, I_b, m	Body mass and inertia, toe mass
γ_c	Toe angle at leg-ground collision
k_r	Leg-ground coefficient of restitution
Various sets	
$\mathcal{H}_l, \mathcal{H}_b$	Leg and body contact state spaces

Table 2: Notation used throughout the paper.

3.2 Contact States and Constraints

In this section, we introduce a framework for modeling the interaction of the body and the legs with the ground. Five binary flags, — a pair for the body end points (the nose and the tail), and a triple for the legs, denoted, respectively, as s_n, s_t, s_1, s_2, s_3 — are sufficient to encode the contact configurations of the system within the contact state spaces \mathcal{H}_b and \mathcal{H}_l ,

$$\begin{aligned}
& [(s_n, s_t), (s_1, s_2, s_3)] \in \mathcal{H}_b \times \mathcal{H}_l \\
& \mathcal{H}_b := \{0, 1\} \times \{0, 1\}, \\
& \mathcal{H}_l := \{0, 1\} \times \{0, 1\} \times \{0, 1\}.
\end{aligned}$$

For any $s \in \mathcal{H}_l$, we use s_i to denote the corresponding contact state for the i^{th} leg for $i = 1, 2, 3$. In the sequel, we will find it useful to introduce a partial order on \mathcal{H}_l imposed by the following relation,

Definition 1 Let $p, r \in \mathcal{H}_l$. We denote by the symbol \geq , the following relation

$$(p \geq r) \Leftrightarrow ((\forall i \ p_i = 0 \rightarrow r_i = 0)) \quad (1)$$

Lemma 1 The relation \geq is a partial order on \mathcal{H}_l .

Proof: \geq is clearly reflexive and transitive. Also, if $p \geq r$ and $r \geq p$, then we have $\forall i, p_i = r_i$, which implies that $p = r$. Hence, \geq is also antisymmetric. \square

The morphology of our model requires that the endpoints of the body be above the ground,

$$b_z > d \sin |\theta|, \quad (2)$$

and that a leg must reach the ground

$$d_i \sin \theta > l - b_z \quad (3)$$

before it can apply any torque to the body.

d	d_1	d_2	d_3	l	M	I_b	m
(m)	(m)	(m)	(m)	(m)	(kg)	(kgm ²)	(kg)
0.28	-0.2	0	0.2	0.12	7.5	0.18	0.1

Table 3: RHex’s Kinematic and dynamic parameters.

Note that three instances of the leg contact constraint (3), partition the space of valid configurations into four different regions, (see Figure 3), within which certain legs are required to be in flight.

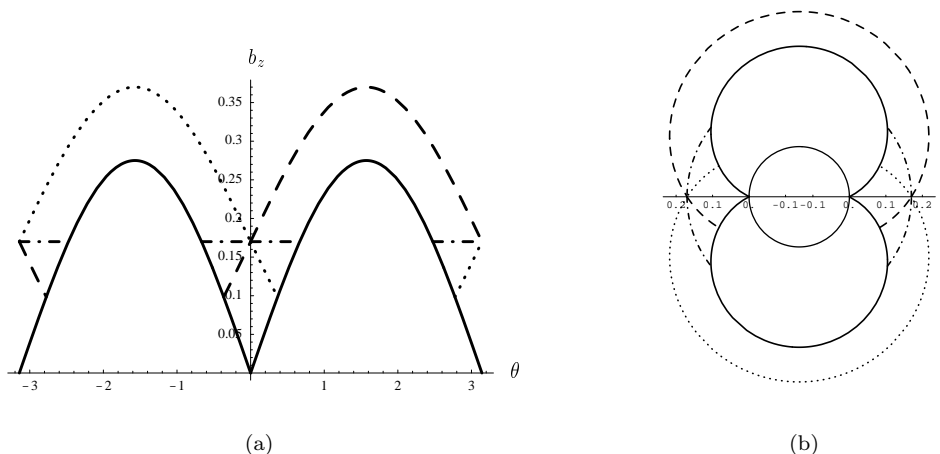


Figure 3: *Hybrid regions in the planar flipping model in (a) the $b_z - \theta$ plane (b) polar coordinates with angle θ and radius b_z using RHex’s morphology (see Table 3). Solid lines indicate body ground contact for the nose ($\theta < 0$) and the tail ($\theta > 0$). The liftoff transitions of the front, middle and back legs are represented with dotted, dash-dot and dashed lines, respectively.*

3.3 The 1DOF Planar Model

Our subsequent analysis mainly concerns configurations where the tail of the body is in contact with the ground and all the legs are positioned such that $0 < \gamma_i < \pi/2$. Our controller design in Section 4 respects these constraints by proper choice of control inputs, reducing the system to only one degree of freedom: the body angle θ with respect to the horizontal. As a convention, we assume that the tail coincides with the origin.

The foot position and leg orientation can then be expressed as functions of θ ,

$$\gamma_i = \text{asin} \left[\frac{d_i + d}{l} \sin \theta \right] \quad (4)$$

$$x_i = (d_i + d) \cos \theta + l \cos \gamma_i . \quad (5)$$

We will also find it convenient to write the leg contact constraints of (3) in functional form for this low degree of freedom model, $s_c : \mathcal{Q} \rightarrow \mathcal{H}_l$, with the component for the i^{th} leg specified as

$$s_c(\theta)_i = \begin{cases} 1 & \text{if } \sin \theta \leq l/(d + d_i) \\ 0 & \text{otherwise} \end{cases} .$$

In the context of this low degree of freedom model, we will refer to the dynamical state of the system, \mathbf{q} , as the *body state* and the discrete leg touchdown configuration as the *contact state*.

3.4 Continuous Dynamics

In this section, we derive the contact constraint forces and the vector field for the constrained model of Section 3.3, for a particular choice of contact state, assuming that the legs that are touching the ground as well as the tail of the body are vertically constrained in both directions (i.e. the ground reaction force can be negative as well as positive). We then present the final form of the equations of motion using the actual contact state $s_m(\mathbf{q}, \tau) \leq s_c(\theta)$, defined in Section 3.5 to yield the continuous dynamics for our model.




Figure 4: Free body diagrams for the body and one of the legs

Figure 4 illustrates the free body diagrams for the body link as well as one of the legs in contact with the ground. Solving the force and moment balance equations for leg i , we have,

$$(l \cos \gamma_i + l \bar{\mu}_i \sin \gamma_i) F_i = l m a_i^x \sin \gamma_i - l m b_i^x \ddot{\theta} \sin \gamma_i - \tau_i + k_d \dot{x}_i \quad (6)$$

where $\bar{\mu}_i := -\mu \text{sign}(\dot{x}_i)$ is defined as the effective Coulomb friction coefficient, k_d is the frictional damping constant and $\ddot{x}_i = a_i^x - b_i^x \ddot{\theta}$ is obtained by differentiating (5) twice.

Combined with the moment balance for the body link around the tail, instances of (6) for each leg that can reach the ground result in a linear set of equations whose solution yields the dynamics.

The number of these equations, however, varies based on the value of $s_c(\theta)$. The following presentation assumes that all the legs can reach the ground, i.e. $s_c(\theta) = [1, 1, 1]$, but the readers should note that there are $2^3 = 8$ different cases for different contact states. According to the free body diagram, we may write

$$\mathbf{A}_p(\mathbf{q}) \mathbf{v} = \mathbf{b}(\mathbf{q}, \tau) \quad (7)$$

where $q \in T\mathcal{Q}$, $p = [p_1, p_2, p_3] \in \mathcal{H}_l$ is an arbitrary contact state such that $p \leq s_c(\theta)$ and the arrays are specified as follows:

$$\mathbf{A}_p(\mathbf{q}) := \begin{bmatrix} f_1 & 0 & 0 & l m b_1^x \sin \gamma_1 \\ 0 & f_2 & 0 & l m b_2^x \sin \gamma_2 \\ 0 & 0 & f_3 & l m b_3^x \sin \gamma_3 \\ p_1 x_1 & p_2 x_2 & p_3 x_3 & -(I_b + M d^2) \end{bmatrix} \quad (8)$$

$$\mathbf{b}(\mathbf{q}, \tau) := \begin{bmatrix} l m a_1^x \sin \gamma_1 + k_d \dot{x}_1 - \tau_1 \\ l m a_2^x \sin \gamma_2 + k_d \dot{x}_2 - \tau_2 \\ l m a_3^x \sin \gamma_3 + k_d \dot{x}_3 - \tau_3 \\ M g d \cos \theta \end{bmatrix}$$

$$\mathbf{v} := [F_1 \quad F_2 \quad F_3 \quad \ddot{\theta}]^T$$

$$f_i := l \cos \gamma_i + l \bar{\mu}_i \sin \gamma_i . \quad (9)$$

Lemma 2 Suppose $0 < \theta < \pi/2$, $\dot{\theta} > 0$ and $0 < \gamma_i < \pi/2$ (the normal range of operation for our thrust controller). Then the matrix $\mathbf{A}_p(\mathbf{q})$ is full rank.

Proof: We first rewrite $\mathbf{A}_p(\mathbf{q})$ in the form of four column vectors

$$\mathbf{A}_p(\mathbf{q}) = [w_1 \quad w_2 \quad w_3 \quad w_4]$$

For the given range of states, the foot locations x_i are clearly positive. Moreover, as a consequence of (9), together with the assumption that $\dot{\theta} < 0$ we have $f_i > 0$. Finally, differentiation of (5) also yields $b_i^x > 0$.

Now, suppose $\sum_{i=1}^{i=4} c_i w_i = 0$ for some $c_i \in \mathbb{R}$. The first three rows of this vector sum yield

$$c_i = -\frac{c_4 l m b_i^x \sin \gamma_i}{f_i} \quad i = 1, 2, 3.$$

Substitution in the fourth row of the vector sum reads

$$-c_4 \left(l m \sum_{i=1}^{i=4} \frac{b_i^x \sin \gamma_i}{f_i} p_i x_i + I_b + M d^2 \right) = 0.$$

Clearly, the sum within the parentheses is strictly positive, yielding $c_i = 0$ for $i = 1, \dots, 4$. Consequently, the column vectors w_i are linearly independent and the matrix $\mathbf{A}_p(\mathbf{q})$ has full rank. \square

As a consequence of Lemma 2, the matrix $\mathbf{A}_p(\mathbf{q})$ is always invertible in the range of operation for our controller. Consequently, the solution to (7) yields the ground reaction forces on the legs as well as the vector field for the particular contact state choice p ,

$$\mathbf{v}_p(\mathbf{q}, \tau) := \mathbf{A}_p(\mathbf{q})^{-1} \mathbf{b}(\mathbf{q}, \tau) . \quad (10)$$

The final form of the equations of motion uses the actual contact state, $p = s_m(\mathbf{q}, \tau)$, defined in Section 3.5,

$$\ddot{\theta}(\mathbf{q}, \tau) = [0 \ 0 \ 0 \ 1] \mathbf{A}_{s_m(\mathbf{q}, \tau)}(\mathbf{q})^{-1} \mathbf{b}(\mathbf{q}, \tau) \quad (11)$$

and only depends on the current body state and the torque input vector, the value of the contact state already being determined.

3.5 Hybrid Leg Contacts

Given the current state \mathbf{q} , we can “read off” from $s_c(\theta)$ the number of kinematically possible leg contacts so as to determine the dimension of the square array $\mathbf{A}_p(\mathbf{q})$ in (7). However, only when a specific set of torques, $\tau \in \mathbb{R}^3$, is imposed at a specified body state, $\mathbf{q} \in T\mathcal{Q}$, can we determine the actual leg contact state according to the function, $s_m : T\mathcal{Q} \times \mathbb{R}^3 \rightarrow \mathcal{H}_l$, and complete the specification of the dynamics in (11). The following definitions will be important in the computation of s_m and subsequent discussions:

Definition 2 $p \in \mathcal{H}_l$ is called consistent at a particular state \mathbf{q} and for a given control input vector τ , denoted $\text{cons}_{[\mathbf{q}, \tau]}(p)$, if and only if

$$\forall i (p_i = 1) \rightarrow F_i(q, \tau, p) > 0$$

Definition 3 Let $p \in \mathcal{H}_l$ be a contact state. p is maximal at \mathbf{q} and τ , denoted $\text{maximal}_{[\mathbf{q}, \tau]}(p)$, if and only if

$$\text{cons}_{[\mathbf{q}, \tau]}(p) \rightarrow (\forall r \in \mathcal{H}_l \text{ cons}_{[\mathbf{q}, \tau]}(r) \rightarrow (p \geq r)) \quad (12)$$

We prove the existence and uniqueness of the maximal contact state in Appendix A. The following Theorem states a very important consequence of maximality.

Theorem 1 *If $p \in \mathcal{H}_l$ is the maximal contact state at \mathbf{q} for a given τ , then*

$$\forall r \in \mathcal{H}_l, (r \leq p) \rightarrow (\ddot{\theta}_r(\mathbf{q}, \tau) \leq \ddot{\theta}_p(\mathbf{q}, \tau)).$$

Proof: This proof is a continuation of the existence and uniqueness proof for the maximal contact state (see Appendix A). It follows trivially from the last row of the vector equality (16). \square

Finally, the following assumption constitutes the basis of our hybrid leg contact model.

A 4 *The contact state of the system is the maximal contact state for its current body state \mathbf{q} and the control torque vector τ .*

Based on this assumption, the following algorithm computes the maximal and hence the actual contact state for a given state \mathbf{q} and torque vector τ .

Algorithm 2 (*Definition of $s_m : T\mathcal{Q} \times \mathbb{R}^3 \rightarrow \mathcal{H}_l$*)

For a given state $\mathbf{q} \in T\mathcal{Q}$ and control inputs $\tau \in \mathbb{R}^3$, this iterative algorithm determines a consistent contact state assignment which is also maximal.

1. *Start with an initial assumption for the leg contact states based on the kinematic constraints.*

$$p^0 = s_c(\theta)$$

2. *Using (7), compute the ground reaction forces $F_i(p^k)$ resulting from the current leg contact state assignment p^k .*
3. *If $\forall i, F_i(p^k) > 0$, p_k is the actual touchdown state, stop the iterations. Otherwise, proceed with the next step.*
4. *Choose the next leg touchdown states to be considered as follows.*

$$p_i^{k+1} = \begin{cases} p_i^k & \text{if } F_i(p^k) > 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{for } i = 1, 2, 3$$

5. *Go to step 2 with $k \leftarrow k + 1$*

3.6 Leg-Ground Collisions

The flipping behavior described in Section 2.1 involves collisions of the front legs with the ground. In order to recover as much of the impact kinetic energy as possible before each thrust cycle, our controllers position the front leg vertically prior to impact, resulting in the radial compliance of the leg to do most of the work. The vertical placement also avoids slippage of the leg as well as friction losses and 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. We will find it convenient to summarize the results of the entire front leg stance phase via an effective “coefficient of restitution” model whose nature we now detail.

In order to derive an accurate model of the collision, it would be possible to extend the continuous dynamics of Section 3.4 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 impulse models can be found in the literature [5]. However, the accuracy of such models is still hostage to the difficulty of determining the dynamic properties of materials as well as other unmodeled effects [3].

In consequence, we have chosen to incorporate a purely algebraic collision law in our model, where a single coefficient of restitution summarizes — as a function of the configuration, θ and the “average torque magnitude” during the collision stance phase — the incremental effects of leg compression/decompression and additional thrust from the middle and back legs.

The following assumptions underlie the construction of our collision law.

A 5 *If a leg is in flight, its angular velocity relative to the body is always zero ($\dot{\phi}_i = 0$), but its position can be arbitrarily specified under the constraint that the toe cannot penetrate the ground.*

A 6 *A collision occurs whenever the body is falling ($\dot{\theta} < 0$), and a leg previously in flight comes into contact with the ground either by explicit positioning through ϕ_i or by the virtue of changes in the configuration θ as a result of the dynamics.*

Accurate modeling of multiple simultaneous collisions is a very fragile and somewhat ill-posed problem [4]. Accordingly, by design, our flipping controller, due to the very particular sequence of leg placements that it enforces, never encounters multiple simultaneous collisions, motivating the following assumption.

A 7 *Multiple simultaneous collisions are not allowed.*

In RHex, the rebounding of the body as a result of leg ground collision mainly involves the radial compression and decompression of the leg.

A 8 *During the collision, we assume that $\tau_i = 0$ and all the contact forces are due to the radial compression of the leg. As a consequence, the impulsive force at the contact point acts along the leg.*

When the radial leg direction lies outside the friction cone, the required frictional impulse would be greater than the Coulomb friction force, violating one of the basic constraints of algebraic collision laws [3]. Even though the flipping controller always chooses leg angles prior to collision to satisfy this constraint, for completeness, we augment our model to use an incremental approach through the dynamics of (11) for collisions outside the domain of the algebraic law.

A 9 *If the leg touches the ground outside the friction cone (i.e. $|\tan(\theta + \phi_i)| > \mu$), then the leg immediately starts slipping and transitions into stance without any impulsive collisions. The system velocities remain continuous ($\theta^+ = \theta^-$).*

Under these assumptions, our collision law models the damping losses arising from the compression and decompression of the front leg as well as the additional thrust provided by the middle and back legs. We assume that these losses can be lumped into a single coefficient of restitution $-1 \leq k_r(\gamma_c)$ as a function of the toe angle at the onset of collision, γ_c .

$$\dot{\theta}^+ = -k_r(\gamma_c) \dot{\theta}^- \tag{13}$$

Note that as a result of Assumption 9, $k_r(\gamma_c) = -1$ whenever $|\cot(\gamma_c)| > \mu$.

3.7 Body Contact Forces During Collision

The collision model described in Section 3.6 implicitly assumes that the impulsive forces necessary to yield the discontinuous change in the body velocity are applied both at the toe and the body-ground contact point. However, the constraints imposed by these contact points are unilateral and only allow ground reaction forces with positive vertical components. If the necessary impulsive forces violate this constraint, the collision model becomes invalid and some unmodeled behavior (such as the tip of the body leaving the ground) will arise. In this section, we derive the ground reaction forces arising from the collision and investigate the validity of the model as a function of kinematic robot design as well as the leg configuration at the time of collision.

The algebraic law of (13) can only be realized through impulsive forces both at the toe and at the body-ground contact point. Figure 5 illustrates the impulsive contact forces effective during the collision. Note that P acts along the leg as a result of Assumption 8.

Moment balance around the point O , yields the leg contact impulse P ,

$$P = \frac{(I_b + Md^2)(\dot{\theta}^+ - \dot{\theta}^-)}{(d + d_i) \cos \phi_i}.$$

The step change in the translational body velocities yield the body ground contact impulses.




Figure 5: Impulsive contact forces during leg-ground collision

$$\begin{aligned} P_c^x &= -Md \sin \theta (\dot{\theta}^+ - \dot{\theta}^-) + P \cos \gamma_i \\ P_c^z &= Md \cos \theta (\dot{\theta}^+ - \dot{\theta}^-) - P \sin \gamma_i \end{aligned}$$

Substituting the algebraic collision model of (13), we obtain

$$\begin{aligned} P_c^x &= \left(-Md \sin \theta + \frac{(I_b + Md^2)}{(d + d_i) \cos \phi_i} \cos \gamma_i \right) (1 + k_r(\gamma_{ci})) \dot{\theta}^- \\ P_c^z &= \left(Md \cos \theta - \frac{(I_b + Md^2)}{(d + d_i) \cos \phi_i} \sin \gamma_i \right) (1 + k_r(\gamma_c)) \dot{\theta}^- \end{aligned}$$

As a result of the unilateral contact constraint, our collision model is only valid when $P_c^z > 0$. Hence, the regions where this collision model is valid are given by the following inequality.

$$Md \cos \theta - \frac{(I_b + Md^2)}{(d + d_i) \cos \phi_i} \sin \gamma_i > 0$$

We can write this inequality as a function of only the configuration of the system at the time of collision and the kinematic robot parameters, yielding

$$\frac{d_i + d}{l} \cos^2 \theta + \sqrt{1 - \frac{(d + d_i)^2}{l^2} \sin^2 \theta} \cos \theta - \frac{I_b + Md^2}{Mld} > 0. \quad (14)$$

Figure 6.a illustrates regions in robot configuration and design space where our algebraic collision model remains valid. Note that with $d_i = 0.2$ for the front leg, and the collision occurring with the leg vertical (boundary of the white and grey regions), the model is valid for RHex's flipping controller. Furthermore, there is also some remaining freedom in choosing different leg angles at impact, possibly yielding better restitution coefficients and recovery of more energy.

This is much more clearly illustrated in Figure 6.b, where model validity is portrayed in a slightly different coordinate system, with θ slaved to the angle of the hip at the time of collision ϕ_c .

4 An Improved Controller

4.1 Constraints on the Control Inputs

Given a particular contact state p , ground reaction forces on the toes can be determined using (6). Similarly, we can compute the contact force on the tail,




Figure 6: Dependence of collision model validity on model parameter d_i and (a) the configuration variable θ , (b) the hip angle ϕ_i . The gray region is where the leg cannot reach the ground. The white region is where the ground reaction force on the body contact has a positive component (i.e. (14) holds), whereas the black region is where the algebraic collision model becomes invalid. Note that in the second graph, a larger range for ϕ_i is covered and θ is slaved to the hip angle through a kinematic touchdown condition.

$$F_c^z = [-p^T \quad Md \cos \theta] \mathbf{v}_p(\mathbf{q}, \tau) + Mg - Md \sin \theta \dot{\theta}^2.$$

To preserve consistency with the assumed contact state in a physically realistic way, all of these ground reaction forces must be positive, inducing limitations on the set of input torque vectors. We capture these constraints and the practical limits on the magnitude of the input torques with the following definition.

Definition 4 For a particular state $\mathbf{q} \in T\mathcal{Q}$ and a contact state $p \in \mathcal{H}_l$, we define the corresponding set of allowable torques, $\mathcal{T}(\mathbf{q}, p)$ as the set of all torque input vectors $\tau \in \mathbb{R}^3$ such that

$$\begin{aligned} F_c^z(\mathbf{q}, \tau, p) &\geq 0 \\ \forall i, F_i(\mathbf{q}, \tau, p) &\geq 0 \\ \forall i, |\tau_i| &\leq \tau_{max} \end{aligned}$$

4.2 Maximal Thrust Control

In this section, we present a torque control strategy which maximizes the thrust while respecting the body contact constraint as well as limitations on the torque deliverable by the hip actuators.

The vector field (11) is a continuous function of the state and the input torques. As a consequence, the problem of choosing hip controls to maximize the thrust becomes a constrained optimization problem over the allowable input torque space, given the current system state and the torque limit constraints. However, different regions in the state space with different leg contact state assignments result in a vector field which is a nonlinear function of the input torques, yielding a computationally demanding nonlinear optimization problem.

Fortunately, in each of the distinct leg contact states, the optimization problem is linear in the control input torques. Hence, the problem decomposes into a small number of separate linear programming problems, from whose independent solutions may be derived a single correct torque value for the three hips. Namely, the optimal torque vector is computed for each distinct leg contact state vector compatible with the kinematic constraints. Then, putative motor torques determined for each kinematically possible leg contact, the unique physical contact state vector is selected by recourse to Algorithm 1, and the command torque vector to each motor derived. More formally, given a leg contact state vector, $p \in \mathcal{H}_l$, we pose the corresponding linear programming problem for that region of configuration state space:

$$\begin{aligned}\tau_M(\mathbf{q}, p) &:= \operatorname{argmax}_{\tau \in \mathcal{T}(\mathbf{q}, p)} \left(\ddot{\theta}(\mathbf{q}, \tau, p) \right) \\ \ddot{\theta}_M(\mathbf{q}, p) &:= \ddot{\theta}(\mathbf{q}, \tau_M(\mathbf{q}, p), p) .\end{aligned}$$

The set of contact state assignments that we need to consider is determined by the kinematic constraints, $\mathcal{P}_{\mathbf{q}} := \{p \in \mathcal{H}_l \mid s_c(\theta) \geq p\}$. The solution to the global problem then becomes,

$$\begin{aligned}\tau &= \tau(\mathbf{q}, p_{max}) \\ p_{max} &:= \operatorname{argmax}_{p \in \mathcal{P}_{\mathbf{q}}} \left(\ddot{\theta}_M(\mathbf{q}, p) \right) .\end{aligned}\tag{15}$$

Note that p_{max} is also maximal at the current body state and with the torque solution to the above optimization problem as a result of Claim 1. Consequently, the actual contact state determined by the algorithm of Section 3.5 will necessarily match p_{max} , that is

$$s_m(\mathbf{q}, \tau(\mathbf{q}, p_{max})) = p_{max} .$$

4.3 Hybrid Energy Pumping

Depending on the frictional properties of the surface, our maximal thrust controller may or may not be enough 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 the first generation controller presented in Section 2.1. Once the body starts falling, the new controller waits until the front legs collide with the ground and repeatedly applies maximal thrust following each collision.

Currently, we have very little analytical understanding of the behavior arising from this hybrid controller. As a consequence, we only explore in simulation the flipping behavior and its dependence on various surface parameters in the following sections.

5 Simulations

In this section, we use simulations based on the dynamical model of Sections 3.4 and 3.5 to evaluate the performance of the maximal thrust control approach combined with the hybrid energy pumping strategy.

5.1 Apex Return Maps

Figure 7 illustrates an example simulation of a successful flip with multiple hops. Note that in the last cycle, as the robot “escapes” from the basin of attraction of the initial configuration, the point of inflection at $t \approx 4.5$ corresponds to the configuration where the downward torque from gravity is just balanced by the rear leg thrusting torque.

In presenting the properties of the hybrid pumping strategy, we will find it most convenient to sample the θ trajectory at its highest point (apex) during each cycle. This will result in a one dimensional return map, characterizing the behavior of the energy pumping strategy under the maximal thrust actuation. This section explores on the basis of numerical simulation some of the properties of this return map and their dependence on various surface parameters and the coefficient of restitution. Formal analysis of this model is presently in progress.

Figure 8 illustrates different types of return maps resulting from different choices of the surface parameters μ , k_c and k_d . This collection of return maps appears to capture all the possible types of phenomena that arise from our hybrid controller.

The upper left case has low ground friction and hence the initial thrust is sufficient to flip the robot body over. Moreover, there is no fixed point, so even if the robot were to start off the ground, it would always be able to flip. In contrast, the upper right case has enough friction to make flipping in one thrust impossible, but still has no fixed point, yielding successful flipping after several hops.




Figure 7: An example simulation of (11) with RHex parameters (see Table 3) and $\mu = 0.8$, $k_d = 11$ and $k_c = 0.75$; using the maximal thrust feedback controller (15). Dotted and dashed lines indicate the kinematic liftoff constraints for the front and middle legs, respectively.

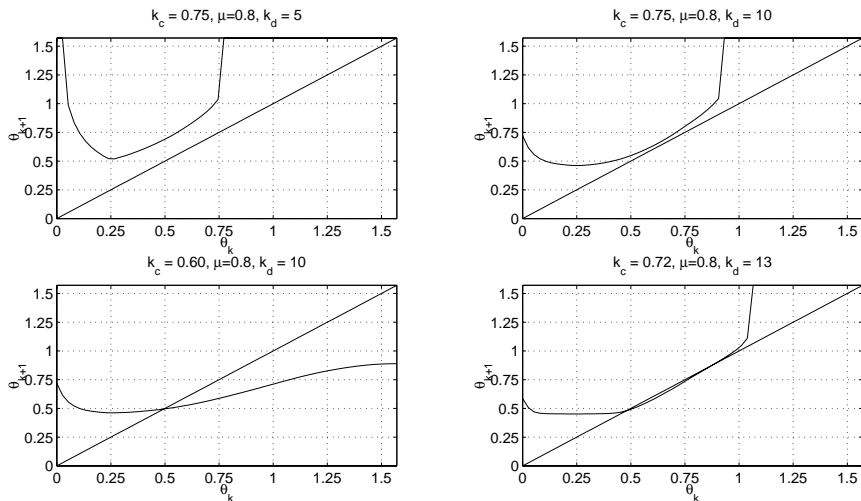


Figure 8: The predicted range of physical behaviors based upon numerical return maps computed for (11) with representative surface parameter settings using the maximal thrust feedback controller (15).

The bottom two cases have qualitatively different behavior. In contrast to the previous cases, they have stable fixed points, trapping the robot at a small angle. For the bottom left case, there is no other fixed point, which makes it impossible for the robot to flip. The bottom right case, however, also has another unstable fixed point, making a successful flip possible for initial conditions above the fixed point.

5.2 Maximal Thrust vs PD Control

Results of simulation runs for a range of surface friction parameters are illustrated in Figures 9 and 10, for $k_c = 0.75$ and $k_c = 0.9$, respectively. The simulations were run over a range of surface friction properties, until either the robot flipped over or at the end of 50 hops. A particular attempt was considered a successful flip if the body angle reached $\pi/2$ before the end of the simulation (labeled *single thrust* and *multiple hops* in the plots), or the sequence of apex heights kept increasing even in the last hops (indicating a weakly unstable fixed point in the return map, labeled *expected flip* in the plots). All other runs were considered failed flip attempts.

One of the reasons for the choice of such high coefficients of restitution is the active nature of the collisions we briefly describe in Section 3.6. In this context, higher coefficient of restitutions reflect the additional thrust occurring exerted during the decompression of the front leg. On RHex, we observed the

duration of the collision to be significant, increasing the effect of this active phase of the collision. Actual values of the coefficient of restitution still remain to be experimentally verified.




Figure 9: Outcomes for flipping attempts with lower coefficient of restitution, $k_c = 0.75$

These results demonstrate that maximal thrust control yields considerably better flipping performance than the PD control in all cases. For smaller k_c , where the “active” collision is not properly modeled, the PD controller never succeeds with multiple hops and only has a chance when the first thrust is sufficient. When the effects of the active collision are incorporated through the coefficient of restitution, the maximal thrust controller is still successful in a very large range of surface conditions and yields strictly better results than the PD control.

6 Conclusion and Future Work

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 — between three to five times the range (in regard to the effective viscous damping that can be overcome). Although the implementation of these improvements on our experimental platform awaits a more complete sensory suite, we believe the actual performance improvement on the robot will be comparable to what we have observed in simulation.

More formal analysis of the preliminary model we have described in this paper is also of great interest.




Figure 10: Outcomes for flipping attempts with higher coefficient of restitution, $k_c = 0.9$

Extensions of the flipping behavior such as uninterrupted rolling or handstands will require a much better analytical understanding of the model as well as modifications such as relaxing the friction constraint on the body. 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.

Acknowledgments

This work was supported in part by DARPA/ONR Grant N00014-98-1-0747.

A Existence and Uniqueness of the Maximal Contact State

First, we will need to state some of the properties of consistent contact state assignments and the partial order defined in Section 3.2.

Lemma 3 *At any given state \mathbf{q} and for any given control input vector τ , there is at least one consistent leg contact state assignment.*

Proof: $p \in \mathcal{H}_l$ such that $\forall i p_i = 0$ is always consistent. \square

Theorem 3 *Let $p, r \in \mathcal{H}_l$. The following statement holds*

$$((p \geq r) \wedge \text{cons}_{[\mathbf{q}, \tau]}(p)) \rightarrow \text{cons}_{[\mathbf{q}, \tau]}(r)$$

Proof: If $p = r$, then the statement holds trivially. So, suppose $p \geq r$, $p \neq r$ and that p is consistent. Let $\mathbf{v}_p := \mathbf{A}_p(\mathbf{q})^{-1}\mathbf{b}(\mathbf{q}, \tau)$ and $\mathbf{v}_r := \mathbf{A}_r(\mathbf{q})^{-1}\mathbf{b}(\mathbf{q}, \tau)$ denote the solutions of (7), for both leg touchdown state vectors p and r , respectively. Noting that $\mathbf{A}_r(\mathbf{q})\mathbf{v}_r = \mathbf{b}(\mathbf{q}, \tau) = \mathbf{A}_p(\mathbf{q})\mathbf{v}_p$, we have

$$\mathbf{A}_r(\mathbf{q})(\mathbf{v}_p - \mathbf{v}_r) = -(\mathbf{A}_p(\mathbf{q}) - \mathbf{A}_r(\mathbf{q}))\mathbf{v}_p.$$

The matrices $\mathbf{A}_p(\mathbf{q})$ and $\mathbf{A}_r(\mathbf{q})$ have the same elements except the last row, where the components of p and r differ. Let J be the set of indices for those legs j such that $p_j = 1$ and $r_j = 0$. J is always nonempty because $p \neq r$ and $p \geq r$. Consequently, we can write

$$-(\mathbf{A}_p(\mathbf{q}) - \mathbf{A}_r(\mathbf{q}))\mathbf{v}_p = [0, 0, 0, -\sum_{j \in J} p_j F_i(p) x_i]^T.$$

Rewriting $\mathbf{A}_r(\mathbf{q})$ as a block matrix, we have

$$\mathbf{A}_r(\mathbf{q})^{-1} = \begin{bmatrix} \Delta F_{(3x3)} & B_{(3x1)} \\ X_{(1x3)} & -D_{(1x1)} \end{bmatrix}^{-1} = \begin{bmatrix} \Delta G_{(3x3)} & W_{(3x1)} \\ Y_{(1x3)} & -Z_{(1x1)} \end{bmatrix}$$

where $\Delta G := (\Delta F + BD^{-1}X)^{-1}$, $Z := (D + X\Delta F^{-1}B)^{-1}$, $W := \Delta F^{-1}BZ$, $Y := D^{-1}X\Delta G$. In Lemma 2, we have shown that for the normal operating domain of our model, $f_i > 0$. Inspection of (8) then shows that the elements of the diagonal matrix ΔF and its inverse as well as B , X are positive. Moreover, $D > 0$, yielding $D^{-1} > 0$ and $Z > 0$ as well. Finally, being defined as products of matrices with positive elements, W and Z also have all positive elements. As a consequence, we obtain the following difference in the solution vectors for p and r

$$\mathbf{v}_p - \mathbf{v}_r = -\mathbf{A}_r(\mathbf{q})^{-1}(\mathbf{A}_p(\mathbf{q}) - \mathbf{A}_r(\mathbf{q}))\mathbf{v}_p = \begin{bmatrix} -W_{(3x1)} \\ Z_{(1x1)} \end{bmatrix} \sum_{j \in J} p_j F_i(p) x_i. \quad (16)$$

By definition of consistency, we have $F_i(p) > 0$. Hence, (16) yields $\forall i, F_i(p) - F_i(r) < 0$. We now have

$$\forall i \ r_i = 1 \rightarrow p_i = 1 \rightarrow F_i(p) > 0 \rightarrow F_i(r) > 0$$

which, by definition, implies the consistency of r . \square

Theorem 4 *Let $p, r, s \in \mathcal{H}_l$. The following statement holds*

$$\text{cons}_{[\mathbf{q}, \tau]}(p) \wedge \text{cons}_{[\mathbf{q}, \tau]}(r) \rightarrow (p \geq r) \vee (r \geq p) \vee (\exists s \in \mathcal{H}_l \text{ cons}_{[\mathbf{q}, \tau]}(s) \wedge (s \geq p) \wedge (s \geq r))$$

Theorem 5 *(Existence and uniqueness of a maximal contact state)*

In any given state \mathbf{q} and for a given control input vector τ , there is always a unique consistent leg contact state assignment which is also maximal.

Proof: Suppose that there are no maximal leg contact state assignments. As a consequence, for all leg contact assignments, the opposite of (12) must hold, namely

$$\forall p \in \mathcal{H}_l \ \neg \text{cons}_{[\mathbf{q}, \tau]}(p) \vee (\exists r \in \mathcal{H}_l \text{ cons}_{[\mathbf{q}, \tau]}(r) \wedge \neg(p \geq r)) \quad (17)$$

Now, consider the following inductive construction. Pick p_0 to be a consistent leg contact state assignment. By Theorem 3, such an assignment can always be found. For the inductive step, suppose $\text{cons}_{[\mathbf{q}, \tau]}(p_k)$ holds. By (17), we can always find $r \in \mathcal{H}_l$, which is consistent such that either $r > p_k$ holds, or neither $p \geq r$ nor $r \geq p$ hold. If $r > p_k$, then choose $p_{k+1} = r$. Otherwise, using Theorem 4, we can always find $s \in \mathcal{H}_l$ such that s is consistent, and it is greater than both p and r . Note that $(s \neq p) \wedge (s \neq r)$ because any other case would imply that r and p could be ordered. We can now choose $p_{k+1} = s$.

This inductive construction results in an infinite ordered sequence $\{p_k\}$ such that $\forall k \in \mathbb{N}$, $(p_{k+1} \geq p_k) \wedge (p_{k+1} \neq p_k)$. However, we know that \mathcal{H}_l is finite, therefore, there exists at least one maximal leg contact state assignment for any given state and control input vector. \square

References

- [1] R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, H. B. Brown Jr., D. McMordie, U. Saranli, R. J. Full, and D. E. Koditschek. RHex: A Biologically Inspired Hexapod Runner. *Autonomous Robots*, 11:207–213, 2001.
- [2] J. E. Bares and D. S. Wettergreen. Dante II: Technical Description, Results and Lessons Learned. *International Journal of Robotics Research*, 18(7):1–29, July 1999.
- [3] A. Chatterjee. *Rigid Body Collisions: Some General Considerations, New Collision Laws, and Some Experimental Data*. PhD thesis, Cornell University, 1997.
- [4] A. Chatterjee and A. Ruina. A new algebraic rigid body collision law based on impulse space considerations. *Journal of Applied Mechanics*, 65(4):894–900, December 1998.
- [5] W. Goldsmith. *Impact: The theory and physical behavior of colliding solids*. Edward Arnold, Ltd., London, 1960.
- [6] E. Hale, N. Schara, J. W. Burdick, and P. Fiorini. A Minimally Actuated Hopping Rover for Exploration of Celestial Bodies. In *Proceedings of the IEEE International Conference On Robotics and Automation*, pages 420–7, San Francisco, CA, April 2000.
- [7] L. Matthies et al. A Portable, Autonomous, Urban Reconnaissance Robot. In *Proceedings of The 6th International Conference on Intelligent Autonomous Systems*, Venice, Italy, July 2000.
- [8] U. Saranli, M. Buehler, and D. E. Koditschek. RHex: A Simple and Highly Mobile Robot. *International Journal of Robotics Research*, 20(7):616–631, July 2001.
- [9] E. Tunstel. Evolution of Autonomous Self-Righting Behaviors for Articulated Nanorovers. In *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, pages 341–6, Noordwijk, The Netherlands, June 1999.