RABBIT: A Testbed for Advanced Control Theory

C. Chevallereau, G. Abba, Y. Aoustin, F. Plestan, E.R. Westervelt,
C. Canudas-de-Wit, and J.W. Grizzle

Introduction

RABBIT is a bipedal robot specifically designed to advance the fundamental understanding of controlled legged locomotion. The motivation for studying walking robots arises from diverse sociological and commercial interests, ranging from the desire to replace humans in hazardous occupations (demining, nuclear power plant inspection, military interventions, etc.), to the restoration of motion in the disabled (dynamically-controlled lower-limb prostheses, and FNS/FES), and the appeal of machines that operate in anthropomorphic or animal-like ways (well-known biped and quadruped toys). From a control design perspective, the challenges in legged robots arise from the many degrees of freedom in the mechanisms, the intermittent nature of the contact conditions with the environment, and underactuation. This article describes the development of our walking robot prototype, some key aspects of its mathematical model, our progress in developing novel control strategies for this interesting hybrid system, and some initial experiments. Supplemental material related to the article—such as photographs of the mechanism, detailed model equations, and videos of the experiments—are available on-line at [1].

Most robot control strategies are built around trajectory tracking, with the trajectories being generated either off-line during a path-planning phase, or on-line through a high-level motion planner. Bipedal robots should clearly be an exception to this as stability is an overriding concern. The combination of highly interactive dynamics, intermittent ground contact, and underactuation makes the planning of asymptotically stabilizable, dynamic motions extremely difficult.

An impressive amount of technology has been amassed and specifically developed to build walking robot prototypes. A quick search of the literature, see for example [2], reveals over a hundred walking mechanisms built by public research laboratories, universities and major companies. Nevertheless, conceptual control breakthroughs have not been keeping pace with the technological developments. There is an amazing deficit in fundamental concepts in comparison to the amount of accumulated technology. The result is a heavy reliance on heuristics, such as the zero moment point (ZMP) principle [3], [4]. This heuristic oversimplifies the stability of the robot's dynamics to a set of steady-state force-balance equations, thereby reducing the notion of a stable motion to the naive and limited view of a frozen dynamics. As a result, only slow motions may be achieved in a stable manner. Truly dynamic motions, such as balancing, running or fast walking, are clearly excluded with this approach [5]. Indeed, for a robot like RABBIT without feet, the ZMP criterion says that the robot is never stable when

moving!

A canonical problem in bipedal robots is how to design a controller that generates closed-loop motions, such as walking, running, or balancing, that are periodic and stable (i.e., limit cycles). Due to the inherent underactuation and the changing contact conditions with the ground, this task is far from being solved through existing control methods. New paradigms, concepts and control analysis are thus needed to deal with this problem. There is also a need to build testbeds with a reasonable level of complexity (degrees of freedom) so that these new concepts can be experimentally verified and refined.

The RABBIT testbed shown in Fig. 1 is the result of a joint effort by several French research laboratories, encompassing mechanical engineering, automatic control, and robotics [6]; the University of Michigan joined in the control effort in late 1998, as the result of a sabbatical in Strasbourg. The project is funded by the CNRS and the National Research Council. Initiated in 1997, its central mission is to build a prototype for studying truly dynamic motion control. In particular, the mechanism was designed to allow for high speed walking and running. RABBIT's lateral stabilization is assured by a rotating bar, and thus only 2-D motion in the sagittal plane is considered (loosely speaking, this is the plane of forward motion, or said another way, it is the plane that divides the body into left and right halves). Except for this limitation, the prototype captures the main difficulties inherent in this type of nonlinear system: underactuation (no feet), variable structure (the state dimension varies as a function of the motion phase), and state jumps (sudden state variations resulting from impacts with the ground). To some extent, this mechanism can be seen as a hybrid system (or a system with a switching structure) that is not everywhere locally controllable. Asymptotically stable motion is thus only achievable through a detailed study of the full dynamics, including impact phases.

To-date, two major achievements of the project are the concept of virtual constraints (VC), and the concept of hybrid zero dynamics (HZD). Virtual constraints are relations among the links of the mechanism that are dynamically imposed through feedback control. Their function is to coordinate the evolution of the various links throughout a step, which is another way of saying that they reduce the degrees of freedom, with the goal of achieving a closed-loop mechanism that naturally gives rise to a desired periodic motion. The mathematical model that describes the reduced-order mechanism is called the hybrid zero dynamics; its hybrid nature is due to the same impacts experienced by the full-order mechanism during locomotion. For a biped robot such as RABBIT that has one degree of underactuation (difference between the number of degrees of freedom and the number of independent actuators), the hybrid zero dynamics evolve on a two-dimensional invariant surface. The result is that the stability of the full-order closed-loop system can be studied on the basis of a nonlinear, two-state differential equation with jumps.

The concepts of virtual constraints and hybrid zero dynamics form a powerful analytical basis for

designing periodic walking motions. These concepts can also be used in more complex tasks where the walking speed varies during the robot's motion, in balancing phases where the robot is dynamically standing on one leg, and, probably, in the near future, in running as well. Moreover, these concepts can be applied to a wide variety of robot designs and structures when the degree of underactuation is equal to one; extensions to higher degrees of underactuation seem possible. This article is devoted to describing the process we followed to design, build and control the biped robot, RABBIT. In particular, we intend to underline the advances in control theory that this testbed has inspired.

Description of the Mechanism

RABBIT was conceived to be the simplest mechanical structure that is still representative of human walking. The requirement of mechanical simplicity naturally led to restricting its motion to the sagittal plane, with lateral stabilization being achieved by external means. However, many of the other design decisions that went into the prototype are less obvious, involving numerous tradeoffs to achieve dynamic performance, scientific objectives, simplicity, and robustness at a cost compatible with a university budget. This section gives an overview of the key design decisions that went into the conception and construction of RABBIT. Photographs of the mechanism are available at [1]. Some of the components are specified in Table I.

How RABBIT came to have five links and no feet. Work conducted in recent years on passive bipedal walking has shown that it is possible to design three-dimensional, anthropomorphic robots that can walk stably down a sloped surface without any actuation whatsoever [7], [8]! One must therefore reflect on the essential role of each link in the design of a walking mechanism, and, in particular, one must ask the question of whether a given joint needs to be actuated or not.

Numerous studies on controlled biped robots have shown that actuation of the hips and knees is essential for providing locomotive power to the robot for walking on a flat or upwardly sloped surface, and for ensuring clearance of the swing leg during a step. However, the case for including actuation at the ankles is less clear. From the start of the RABBIT project, we had the goal of demonstrating that actuated ankles are *not* absolutely necessary for the existence of asymptotically stable locomotion, and thus RABBIT has no feet. Without actuated ankles, lighter feet can be designed, which is more efficient for walking and running. If the robot can still be shown to achieve stable walking or running over a wide range of speeds on flat ground, then actuation of the ankles must be justified on the basis of improved adherence with the walking surface, better adaptability over non-smooth surfaces, or for ameliorating the shocks associated with the feet impacting the ground. Finally, without feet, the ZMP heuristic is not applicable, and thus underactuation must be explicitly addressed in the feedback control design, leading to the development of new feedback stabilization methods.

For the RABBIT project, we wanted a robot that could run as well as walk. Since we also sought a robot that could perform anthropomorphic gaits, RABBIT had to have at least a hip and two knees,

giving a minimum of four links. For the robot to be able to carry a load, a torso was necessary, making a total of five links. RABBIT is thus a seven degree of freedom mechanism, with four degrees of actuation. In the upright position, with both legs together and straight, the hip is 80 cm above the ground and the tip of the torso is at 1.425 m. Table II specifies the lengths, masses and inertias of each link of the robot.

Why RABBIT walks in a circle and has wheels orthogonal to its sagittal plane. Without active lateral stabilization [9], a biped *walker* can still be designed to maintain its lateral stability by means of "laterally pointing feet", that is, bars or plates attached at the leg ends that extend laterally and prevent the robot from tipping over sideways [10], [7], [8]. But in the case of a *runner*, where a flight phase exists (i.e., ballistic motion—no contact with the walking surface), some means is required to maintain lateral stability. In order for this external stabilization device not to limit the displacement of the robot, the choice of a circular path was made. Hence, our biped robot is guided around a central column by means of a radial bar; see Fig. 2. The same solution for lateral stabilization has already been implemented at the MIT LegLab; see for example [11], [12]. The robot is attached to the radial bar via a revolute joint that is aligned with the axes of the hips, and it is attached to the central column with a universal joint.

With this lateral support device, the robot's sagittal plane is tangent to a sphere centered on the universal joint. As explained in Fig. 3, it follows that the distance between the stance leg end and the central column must be allowed to vary with the position of the hip. To permit frictionless radial displacement of the supporting leg end, wheels directed in the frontal plane (i.e., normal to the sagittal plane) are used. In this way, no mobility of the leg end exists in the sagittal plane of the robot, and thus, with a sufficiently long radial bar (we use one three meters in length), we can accurately model the motion of the robot tangential to the sphere as that of a perfectly planar robot.

Choice of actuation. Specifying the actuation is a key step in the design process of a robot. This includes the choice and sizing of actuation technology. The use of electric motors allows for simpler control, higher bandwidth, and easier construction than hydraulic or pneumatic drives. The choice of the type of electric motor usually comes down to quality measures, such as power to weight ratio [13], [14]. We chose DC motors with Samarium Cobalt magnets, though nearly identical performance in terms of torque density and peak torque could have been had with brushless motors. A gear reducer and belt were used to connect the motors to each of the four actuated joints. The motors for the knees were mounted as close as possible to the hips in order to minimize the inertia of the legs; this decreases the coupling in the dynamic model as well as the required motor torques.

How simulation aided the design process. Once the motor technology was selected, sizing was determined on the basis of dynamic simulations and off-line trajectory optimization [15], [16]. Indeed, in order to check if the proposed structure would be able to walk and run, a simulation study was

conducted. Feasible trajectories were computed, along with the torque needed to achieve them in open loop. One difficulty is that in both flight and single support, our robot is underactuated. During the single support phase, the degree of underactuation is one (five degrees of motion freedom due to the constraint that the stance leg end does not slip, and four actuators), while during the flight phase, the degree of underactuation is three (seven degrees of motion freedom and four actuators). Hence, even though a given motion of the robot may be kinematically realizable, it may not be dynamically feasible [17], so a kinematic analysis combined with an inverse torque model is definitely not sufficient for determining possible walking and running motions.

From a global perspective, we would like the robot to be able to walk and run efficiently, in the sense that the energy cost for a given motion will be as small as possible. Thus, dynamic optimization [18] was used to compute optimal walking and running trajectories, assuming nominal values for the mechanical parameters as well as for the motor characteristics, specifically their torque and speed limits. Reaction forces at the leg ends were calculated to check that all contact conditions were met (the stance leg remains in contact with the walking surface and does not slip). These calculations allowed us to evaluate for each joint the torque-speed curve as a function of walking and running speed, as illustrated in Figs. 4 and 5. By carrying out this analysis for a wide range of walking and running speeds, we were able to determine the total operating range required of each motor, and thereby arrive at its required size. These specifications were then matched to off-the-shelf components, both for the motors and the gear reducers. In the end, Rabbit was designed to be able to walk with an average forward speed of at least 5 km/h and to run at more than 12 km/h.

Impacts, shocks and how to damp them out. An impact or shock occurs in the majority of cases when the swing leg contacts the ground. The only way to avoid a shock is for the velocity of the leg end to be zero at the contact moment. Shocks have obvious deleterious effects on the durability and life time of a mechanical system.

The most affected components are the bearings, gear-reducers, belts and sensors. It is therefore indispensable from the beginning to plan for a source of damping in the system in order to prevent the transmission of large shocks to the most sensitive parts. The magnitude of the shock is determined by the nature of the walking surface (hard, soft, absorbing) and the material used at the end of the leg. The frontal wheels on the leg ends were therefore constructed of a stiff, shock absorbing, polymer. The belts between the motors and the gear boxes were designed to provide additional protection.

Bandwidth and sensing. The speed of response or bandwidth of each axis of the robot is determined by the transfer function of the mechanical powertrain (motors, gears, and belts) and the power amplifiers that drive each motor. In the case of RABBIT, we have approximately a 12 Hz bandwidth in the mechanical portion of the axes and approximately 250 Hz for the amplifiers.

Because RABBIT is an experimental apparatus, we installed a maximal sensor set. The four actuated

joints of the robot are each equipped with two encoders to measure angular position; velocity must be calculated from position. One encoder is attached directly to the motor shaft, while the second is attached to the shaft of the gear-reducer; this configuration allows any compliance between the motor and the joint angle to be detected, though subsequent experimentation has shown that the connection is adequately rigid for control purposes. Identical encoders are used at each joint. The mechanism has three additional encoders. One measures the angle of the torso with respect to a vertical axis established by the central column around which RABBIT walks. The second measures the horizontal (surge) angle of the stabilizing bar with respect to the central column; this allows the distance travelled by the robot to be computed. The final encoder measures the pitch angle of the stabilizing bar, which allows the height of the hips to be measured; in single or double support, this information is redundant, but when both feet are off the ground, as in running, it is not.

The robot was initially equipped with two force sensors, one at the end of each leg, to measure the tangential and normal components of the forces exerted at the contact of the robot and the ground. These turned out to be insufficiently robust, and were replaced with contact switches. The support leg and double support phases are easily distinguished through the positions of the contact switches. Estimating contact moment through swing leg height as determined by the position measurements is not sufficiently accurate.

Mathematical Model

RABBIT is modeled as a planar robot, consisting of five rigid links with mass, connected through rigid, revolute joints. Actuation is provided at the hips and knees, but not at the leg ends. This section describes the mathematical model used for studying walking and running of this mechanism on a flat surface. The key point is that the model is necessarily hybrid, with the dynamics changing as a function of the contact conditions of the legs with the ground.

A walking motion consists of successive phases of *single support* (meaning the stance leg end is touching the walking surface and the swing leg is not) and *double support* (both legs are in contact with the walking surface), while running consists of successive phases of single support and *flight* (there is no contact with the walking surface); see Fig. 6. For the control designs considered here, the double support phase of walking will be modeled as instantaneous, and a rigid impact will be used to model the contact of the swing leg with the ground; these two assumptions simplify the analysis. For a human with feet, the double support phase represents approximately 20% of a walking cycle [19] and the contact is compliant. Control with a non-instantaneous double support phase has been considered in [20], and compliant contact models have been studied in [21]. However, the experimental section will show that our simplifying assumptions of an instantaneous double support phase and a rigid contact model are reasonable for a robot with point feet.

The different phases of walking and running naturally lead to distinct mathematical models, with

the logic for switching among them determined by the number of legs in contact with the ground and the nature of that contact (contact conditions depend on whether or not a foot is slipping, rebounding, etc.). For walking with an instantaneous double support phase, the mathematical model of the biped consists of two parts: the differential equations describing the system during the single support phase, and an algebraic model of the impact.

Seven DOF model. We begin with the most general model of the robot, and then point out special cases as we proceed. A planar mechanism consisting of five rigid links connected via revolute joints in a tree structure with *no contacts or constraints* has seven degrees of freedom (DOF): a degree of freedom associated with the orientation of each link, plus two DOF associated with the horizontal and vertical displacement of the center of mass within the sagittal plane. The state vector of the dynamical model is thus 14-dimensional: there are seven configuration variables required to describe the position of the robot, plus the associated velocities. A convenient choice of configuration variables is depicted in Fig. 6, namely, q_e consists of the two relative angles between the torso and femurs, the two relative angles at the knees, the angle of the torso with respect to the vertical, and the cartesian position of the hips, (x_H, y_H); the velocity is denoted by $\dot{q}_e$. These choices are arbitrary and other choices are common; for example, it is sometimes convenient to attach the cartesian coordinates at the center of mass instead of the hips, or determine all of the angles with respect to a common reference frame.

The model is easily obtained with the method of Lagrange, which consists of first computing the kinetic energy and potential energy of each link, and then summing terms to compute the total kinetic energy, K, and the total potential energy, V [22], [23], [24]. The Lagrangian is defined as $L = K - V$, and a second order dynamical model immediately follows from Lagrange's equation

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{q}_e} - \frac{\partial L}{\partial q_e} = \Gamma,$$

where Γ is the vector of generalized forces and torques applied to the robot. This leads to a model of the form

$$D_e(q_e) \cdot \ddot{q}_e + C_e(q_e, \dot{q}_e) \cdot \dot{q}_e + G_e(q_e) = B_e \cdot u + E_e(q_e) \cdot F^{ext}, \tag{1}$$

where, D_e is the inertia matrix, the matrix C_e contains Coriolis and centrifugal terms, and G_e is the gravity vector; the matrices B_e and E_e are derived from the principle of virtual work and define how the joint torques, u, and external forces, F^{ext}, enter the model. The reader is invited to download the model at [1].

The joint torques are relatively straightforward to model as they consist of the motor torques plus friction in the gear drive. The external forces, on the other hand, represent the forces on the leg ends due to intermittent contact with the walking surface, and how one models them is dependent on the assumed contact conditions. One approach is to model the normal forces on the legs as a compliant contact with the ground via a nonlinear spring-damper [25], [26], and the tangential forces as dynamic

friction [27], [28], [29]. The advantages of this model are that the dynamics are continuous, consisting of ordinary differential equations, and the double support phase is easily addressed. However, there are important drawbacks. For one, a walking surface is usually quite stiff (almost non-compliant), which leads to integration problems for the differential equations. Secondly, no one really knows how to choose reasonable parameters for a compliant contact model, so even though the model seems realistic, it may not be. Finally, the overall model becomes high order and is extremely difficult to analyze.

An alternative model for the external forces due to contact with the ground assumes a rigid contact. This assumption gives a simple algebraic relation for computing the ground reaction forces. In our opinion, the rigid contact model is best for control law design and analysis, while a compliant model is very useful for verifying that the resulting control law is not too dependent on the idealized assumptions used in the rigid contact model; see [21].

Five and three degree of freedom models. Consider the robot in single support on a flat, rigid surface, as in Fig. 6 (a). There are obviously no forces on the swing leg end. If we assume that the stance leg end is stationary (i.e., in contact with the walking surface and not slipping), then the external forces acting on it are easily determined. Indeed, let (x_1, y_1) denote the cartesian coordinates of the stance leg end (see Fig. 6),

$$\left[\begin{array}{c} x_1 \\ y_1 \end{array} \right] = \left[\begin{array}{c} x_H - L_2 \sin(q_1 + q_{32}) - L_1 \sin(q_1 + q_{32} + q_{42}) \\ y_H + L_2 \cos(q_1 + q_{32}) + L_1 \cos(q_1 + q_{32} + q_{42}) \end{array} \right], \tag{2}$$

and assume that single support corresponds to $(x_1, y_1) = (0, 0)$. Computing the accelerations, $(\ddot{x}_1, \ddot{y}_1)$, and equating them to zero allows one to solve for the leg reaction forces as explicit functions of $(q_e, \dot{q}_e)$ and the joint torques:

$$\left[\begin{array}{c} F_1^{ext} \\ F_2^{ext} \end{array} \right] = R(q_e, \dot{q}_e, u); \tag{3}$$

this is the vector of forces necessary for the contact constraint to be maintained [30].

The set of equations (1) and (3) completely specifies the model in single support if the stance leg end is stationary. When solving the equations, one monitors the forces in (3) to ensure that the vertical force is upward, $F_1^{ext} > 0$, and the ratio $|F_2^{ext}/F_1^{ext}|$ is less than the assumed static friction coefficient. Otherwise, the robot is either in flight, so the reaction forces are zero, or slipping, so a dynamic friction model is required.

When designing a controller for walking or running, one normally seeks a solution with no slipping of the stance leg in single support. In this case, it is clear that the hip position is constrained by the angular positions, and thus the robot has only five degrees of freedom. Letting q denote the five angular coordinates depicted in Fig. 6 (a), one can either eliminate the constrained variables from (1), or directly apply the method of Lagrange to the reduced set of coordinates to derive the lower dimensional model

$$D(q)\ddot{q} + C(q, \dot{q})\dot{q} + G(q) = Bu. \tag{4}$$

The reaction forces necessary for this model to be valid are still given by (3), with the hip position (x_H, y_H) obtained from (2). This form of the model is much easier to use and analyze for control design in the single support phase. We note that the degree of underactuation in (4) is equal to one.

If the robot is in double support and appropriate contact hypotheses are made, then a similar analysis can be performed to arrive at a three degree of freedom model (the two additional constraints on the other leg eliminate two more degrees of freedom). In this case, the robot has one degree of *over* actuation, which can be advantageous for control design [20], [31].

Rigid impact model and leg swapping. An impact occurs when the swing leg end touches the walking surface. For control design, we have used a rigid contact model. It is assumed that: the impact is instantaneous; the impulsive forces due to the impact may result in an instantaneous change in the velocities, but there is no instantaneous change in the positions; and, the contact of the swing leg end with the ground results in no rebound and no slipping of the swing leg, and the stance leg lifting from the ground without interaction. This results in the double support phase being instantaneous. Under these assumptions, the seven degree of freedom dynamic model (1) can be used to determine an expression for $\dot{q}^+$, the vector of angular velocities just after impact, in terms of the configuration of the robot at impact and $\dot{q}^-$, the vector of angular velocities just before impact [32], [33]. The post-impact velocity is then used to re-initialize the model for the next step. Since the model (4) made a choice of the stance leg, a change of coordinates is necessary since the former swing leg must now become the stance leg and vice versa. It is convenient to include this coordinate swap as part of the impact map. The final result is an expression for $x^+ = (q^+, \dot{q}^+)$ in terms of $x^- = (q^-, \dot{q}^-)$, which is written as

$$x^+ = \Delta(x^-). \tag{5}$$

It is possible to derive a similar model that allows a non-trivial interval of time for the double support phase.

A hybrid model of walking. Walking is modeled as the alternation of phases of single support and double support, where the double support phase is instantaneous, and no slipping occurs during the single support phase. Letting $x = (q, \dot{q})$, the dynamic model (4) and the rigid impact model (5) together yield a nonlinear system with impulse effects [34]. Assume that the system trajectories possess finite left and right limits, and denote them by $x^-(t) := \lim_{\tau \nearrow t} x(\tau)$ and $x^+(t) := \lim_{\tau \searrow t} x(\tau)$, respectively. The model is then:

$$\begin{array}{rclr} \dot{x} &=& f(x) + g(x)u, & x^- \notin S, \\ x^+ &=& \Delta(x^-), & x^- \in S, \end{array} \tag{6}$$

where

$$S := \{(q, \dot{q}) \mid y_2(q) = 0\} \tag{7}$$

is the set of points where the swing leg touches the ground. In simple words, a solution of the model is specified by the single support model until an impact occurs. Impact occurs when the state attains the

set S (swing leg touches the ground), which represents the walking surface. At this point, the impact with the ground results in a very rapid change in the velocity components of the state vector. The impulse model, Δ, compresses the impact event into an instantaneous moment in time, resulting in a discontinuity in the velocities. The ultimate result of the impact model is a new initial condition from which the single support model evolves until the next impact. In order for the state not to be obliged to take on two values at the "impact time", the impact event is described in terms of the values of the state just prior to impact at time t^-, and just after impact at time t^+. These values are represented by x^- and x^+, respectively. A representation of the model as a simple hybrid system is shown in Fig. 7.

An Interesting Property of the Model, or, The Role of Gravity in Walking

RABBIT has no actuation at the leg ends. So, what causes the robot to rotate about the support leg end and thus advance forward in a step? The answer is gravity! Let σ be the angular momentum of the robot about the stance leg end, which is assumed to act as a pivot (i.e., it does not slip and remains in contact with the walking surface). The angular momentum balance theorem says that the time derivative of the angular momentum about a fixed point equals the sum of the moments of the external forces about that point. Since the motor torques act internally to the robot, their contribution to the moment balance is zero, leaving only gravity:

$$\dot{\sigma} \;=\; M \cdot g \cdot x_c, \tag{8}$$

where, x_c is the difference between the x-coordinate of the stance leg end and the x-coordinate of the center of mass of the robot, M is the total mass of the biped, and g is the gravity constant. In this regard, RABBIT functions just like a passive bipedal walker [7], [8].

So what is the role of the actuators at the hips and knees? The actuators directly act on the shape or posture of the robot, thereby changing the position of the center of mass, and, thus, the moment arm through which gravity acts on the robot. In addition, the posture of the robot has a large effect on the energy lost at impact [35] and whether or not the required contact conditions at the leg ends are respected. The challenge is to bring all of this together in a manner that ensures the creation of a desired asymptotically stable, periodic motion.

What would change if the robot had feet and actuation at the ankle? Then, the additional torque could potentially be used to allow for faster locomotion and even larger basins of attraction [36]. On the other hand, the control authority of the ankle is limited by the contact constraints of the foot and the ground [5], and thus a sophisticated control law design is required to take full advantage of an actuated foot. Three of the most technologically advanced biped robots today, namely, Honda's Asimo, Sony's SDR-4X, and the University of Munich's Jogging Johnnie, are walking on the basis

of the ZMP heuristic, which imposes a conservative walking motion where the support foot remains flat on the ground throughout the stance phase. This is done specifically to avoid dealing with the underactuation that occurs when the foot rotates about the toe in an anthropomorphic gait; see Fig. 8 (as motivation for allowing toe roll, it is interesting to note that analytical work in [37] shows that plantarflexion of the ankle, which initiates toe roll, is an efficient method to reduce energy loss at the subsequent impact of the swing leg). Our work on point feet forces us to deal directly with this underactuation; indeed, conceptually, a point foot corresponds to continuous rotation about the toe throughout the entire stance phase (e.g., walking like a ballerina). A more complex hybrid model of walking with feet is also indicated in Fig. 8.

What's Behind Stable Walking: The Hybrid Zero Dynamics

To properly control RABBIT, we need a control theory for a class of systems with both continuous and discrete dynamics, and fewer actuators than degrees of freedom; in addition, we want asymptotically stable orbits rather than equilibrium points, and we have multiple objectives, such as walking, running, and balancing. In this section and the next we attempt to summarize in as non-mathematical terms as possible the key control ideas we have developed for walking and balancing, while still conveying a sense of what it takes to achieve provably stable control solutions; the reader wishing a complete mathematical treatment is referred to more technical publications at appropriate places. All of our stability claims will be local as global stability probably does not make much sense for bipedal mechanisms. The literature on biped robots is quite extensive. The reader seeking a general overview would do well to start with [38], [3], [39], [40], in that order. Some control-oriented works that we have found especially illuminating, because of their emphasis on analytical aspects of walking, running, and balancing are [41], [42], [43], [44], [36], [9], [45], [46]. For an insightful analysis of another system that exhibits limit cycles and impacts, see [47].

Time-invariance, or, self-clocking of periodic motions. The controller designs that we propose for walking do not involve trajectory tracking. Why? We mentioned one reason in the Introduction; here is another. In a controller based upon tracking, if a disturbance affects the robot and causes its motion to be retarded with respect to the planned motion, for example, the feedback system is obliged to play catch up in order to regain synchrony with the reference trajectory. Presumably, what is more important is the orbit of the robot's motion, that is, the path in state space traced out by the robot, and not the slavish notion of time imposed by a reference trajectory (think about how you respond to a heavy gust of wind when walking). A preferable situation, therefore, would be for the robot in response to a disturbance to converge back to the periodic orbit, but not to attempt otherwise re-synchronizing itself with time. One way to achieve this is by parameterizing the orbit (i.e., the walking motion) with respect to (a scalar-valued function of) the state of the robot, instead of time [48], [49], [50]. In this way, when a disturbance perturbs the motion of the robot, the feedback

controller can focus solely on maintaining limb positions and velocities that are appropriate for that point of the orbit, without the additional burden of re-synchronizing with an external clock. As a bonus, the controller is time invariant, which helps analytical tractability.

Feedback as a mechanical design tool: the notion of virtual constraints. Let us start by considering something more familiar and less complicated than a biped. Fig. 9 (a) depicts a planar piston in an open cylinder. The system has one DOF, which means that a model can be given in terms of the angle of the "crank", θ_1, and its derivatives. Fig. 9 (b) represents the planar piston without the constraints imposed by the walls of the cylinder. The system now has three degrees of freedom involving three coupled equations in the angles $\theta_1, \theta_2, \theta_3$, and their derivatives. One degree of motion freedom is imposed on the three degree of freedom model through two constraints: (a) the center of the piston lies always on a vertical line passing through the rotation point of the crank and (b), the angle of the piston head is horizontal throughout the stroke. Mathematically, this is "equivalent" to imposing

$$
\begin{aligned}
0 &= L_1 \cos(\theta_1) + L_2 \cos(\theta_1 + \theta_2), \\
\pi &= \theta_1 + \theta_2 + \theta_3,
\end{aligned}
\tag{9}
$$

where L_1 is the length of the crank, and L_2 is the length of the second link (due to the existence of multiple solutions, one must choose the solution corresponding to the piston being above the crank). These two constraints can be imposed through the physical means of the cylinder walls shown in Fig. 9 (a), or, through the use of additional links as shown in Fig. 9 (c).

If the system is appropriately actuated, the constraints can also be *asymptotically* imposed through feedback control. To see this, assume that the joints θ_2 and θ_3 are actuated. Define two outputs in such a way that zeroing the outputs is equivalent to satisfying the constraints; for example

$$
\begin{aligned}
y_1 &= L_1 \cos(\theta_1) + L_2 \cos(\theta_1 + \theta_2), \\
y_2 &= \theta_1 + \theta_2 + \theta_3 - \pi.
\end{aligned}
\tag{10}
$$

The constraints will then be asymptotically imposed by any feedback controller that asymptotically drives y_1 and y_2 to zero; for the design of the feedback controller, one could use computed torque, PD control, etc.

The outputs (10) are expressed as implicit functions of the actuated joint angles. Sometimes it is more convenient to express them in an explicit form. As long as $L_1 < L_2$, the constraints (9) can also be re-written as explicit functions of the crank angle, θ_1, per

$$
\begin{aligned}
\theta_2 &= \pi - \theta_1 - \arccos(\tfrac{L_1}{L_2}\cos(\theta_1)), \\
\theta_3 &= \arccos(\tfrac{L_1}{L_2}\cos(\theta_1)),
\end{aligned}
\tag{11}
$$

leading to the alternate output functions

$$
\begin{aligned}
y_1 &= \theta_2 - (\pi - \theta_1 - \arccos(\tfrac{L_1}{L_2}\cos(\theta_1))), \\
y_2 &= \theta_3 - \arccos(\tfrac{L_1}{L_2}\cos(\theta_1)).
\end{aligned}
\tag{12}
$$

We have used both explicit and implicit forms of the constraints when controlling the biped.

When constraints are imposed on a system via feedback control, we call them *virtual constraints*. The planar three DOF piston of Fig. 9 (b) can be virtually constrained to achieve asymptotically the same *kinematic behavior* as the one DOF piston in Fig. 9 (a); the resulting dynamic models are different because the constraint forces are applied at different points of the 3 DOF piston. The virtual constraints can be imposed through the virtual constraints given by (9) or those in (11). In the case of a robot, the advantage of imposing the constraints on the mechanism virtually (i.e, via feedback control) rather than physically (i.e, through complicated couplings between the links or the environment), is evident: the robot can then be "electronically reconfigured" to achieve different tasks, such as walking at different speeds, going up stairs, and running.

Constraint forces and torques. There are forces or torques associated with imposing virtual constraints on a system [51], just as there are with physical constraints [30]. To see this, suppose that the outputs (12) have been successfully zeroed by some feedback controller. Then $y(t) \equiv 0$ implies that all of its derivatives are identically zero as well. Since the torques show up in the accelerations, the interesting equation is $\ddot{y}(t) \equiv 0$. Upon substitution of the model equations and the constraints, and after verifying that certain matrix inverses exist, one arrives at an expression for the torque required to maintain the constraints as a function of the unactuated variable, θ_1, given by

$$u = u^*(\theta_1, \dot{\theta}_1). \tag{13}$$

Since u^* is also the asymptotic value of any feedback controller that asymptotically imposes the constraints, making it small over a given motion is clearly related to efficient walking, running or balancing. For more details, see the Appendix.

Swing-phase control of RABBIT through virtual constraints. Consider now our 5-link, 4-actuator, planar, bipedal walker, RABBIT, in single support. Several different constraint choices have been explored, with the common element being that four independent outputs are regulated since the robot has four actuators. In [21], a largely Cartesian view is taken. The choice was made to regulate the angle of the torso, the height of the hips, and the position of the end of the swing leg (both horizontal and vertical components). Over each step of a normal walking gait, the horizontal position of the hips is monotonically increasing. Hence, along a step, the desired torso angle, hip height and swing leg end position were expressed as functions of the horizontal position of the hip. These four functions were chosen so that, as the hips advance, the torso is erect at a nearly vertical angle, the height of the hips rises and falls during the step, the swing leg advances from behind the stance leg to in front of it, tracing a parabolic trajectory. In [52], [53], [54], the desired motion of the robot is described in terms of the evolution of relative joint angles. The four outputs are selected to be colocated with the actuators: the two relative angles of the torso with the femurs and the two relative angles of the knees. Two additional choices have been used for the monotonic parameter. The angle of the virtual support leg, that is, the line connecting the stance leg end to the hips, is clearly

monotonic over a step whenever the horizontal component of the hips is monotonic; see Fig. 10. The angle of the stance tibia with respect to the ground seems to be monotonic as well over most walking gaits. Consequently, the four relative angles can be virtually constrained as explicit functions of the angle of the virtual support leg, or the support tibia.

Mathematically, four outputs of the form

$$y \;\; = \;\; h_0(q) - h_d(\theta(q)) \tag{14}$$

are constructed. The function h_0 specifies four independent quantities that are to be controlled, while $\theta(q)$ is a scalar function of the configuration variables that is independent of h_0 and is monotonically increasing along a non-pathological step. The function $h_d(\theta(q))$ specifies the virtual constraints as a function of $\theta(q)$. The condition $y \equiv 0$ imposes the constraints. One may interpret the quantity $\theta(q)$ as playing the role of time and the function h_d as taking the place of a desired time trajectory; in this way, the evolution of the posture or "shape" of the robot will be "synchronized" to an internal variable.

It is hoped that the general principles of how a control law can be constructed are starting to become clear; all of the details can be found in [53]. What probably remains unclear at this point is: how the impacts come into the picture; how to determine if a given set of constraints leads to a walking motion, and if it does, is the walking motion stable; how to design the constraints so that a walking motion is efficient in a certain sense; and how to ensure that contact conditions at the leg ends are met. These issues are addressed through the hybrid zero dynamics HZD.

The hybrid zero dynamics: a low-dimensional description of the closed-loop, hybrid mechanism. Once again, let us begin the discussion with the planar piston instead of the biped, and for definiteness, let us consider the three DOF piston of Fig. 9 (b), along with the set of outputs (12). The dynamics of the system compatible with the outputs being identically zero, that is, the constraints (11) being perfectly respected, is called the *zero dynamics* [51]. In the case of the planar piston, it is clear that the zero dynamics has one degree of freedom. This example illustrates again the important and well-known principle that an N degree of freedom mechanical system plus M (independent and holonomic) constraints leads to an $N - M$ degree of freedom mechanical system [22]. When the constraints are applied virtually instead of physically, the zero dynamics describe the *exact* behavior of the closed-loop system whenever the system is initialized so that the constraints are exactly satisfied, and the applied feedback controller maintains the outputs exactly zeroed; otherwise, the zero dynamics describe the *asymptotic* behavior of the closed-loop system as long as it is initialized sufficiently well that the feedback controller manages to drive the outputs asymptotically to zero.

These same general ideas are true of the biped when walking, with one important difference: the biped model is hybrid, since it has a swing phase and an impact phase; consequently, the notion of the zero dynamics must be extended to allow for impacts. This situation has been analyzed carefully

in [53], leading to the notion of the *hybrid zero dynamics*. The basic idea is the following: Let Z be the surface of all points in the state space of the swing phase model of the robot corresponding to the outputs being identically zero. Because of the manner in which we have chosen the virtual constraints, this is a two dimensional surface. The swing phase zero dynamics are the dynamics of the swing phase model restricted to Z; that is, subject to the virtual constraints being exactly imposed. Since the swing phase model has five DOF and there are four constraints, the swing phase zero dynamics has one DOF. It follows that the swing phase zero dynamics may be written as a pair of first order equations, and, with a little help from the angular momentum balance theorem, they can be shown to have the special form

$$
\begin{aligned}
\dot{\theta} &= \tfrac{1}{I(\theta)}\sigma_Z, \\
\dot{\sigma}_Z &= Mgx_c(\theta),
\end{aligned}
\tag{15}
$$

where σ_Z is the angular momentum of the robot about the pivot point of the stance leg, restricted to Z, and $I(\theta)$ plays the role of an inertia.

Since every solution of the zero dynamics corresponds to a solution of the five degree of freedom model, we must consider what happens at the moment when the swing leg touches the ground in the corresponding solution of the five degree of freedom model. At the contact instant, the impact model must be applied, resulting in a new initial condition of the five DOF model. If this new initial condition lies outside of Z, then the swing-phase zero dynamics cannot be used to compute the next phase of the solution to the model. If, on the other hand, the initial condition resulting from the impact model lies in Z, then a solution of the zero dynamics can be continued, and the evolution of the high degree of freedom robot is exactly and completely described by the one degree of freedom model; see Fig. 11. In the latter case, where impacts in Z are mapped back into new initial conditions in Z, the constraints are said to be *invariant* under the impact map. Achieving this requires that extra care be taken when the constraints are designed [53]; from now on, we assume that this has been done. The hybrid zero dynamics consist of the zero dynamics of the swing phase in combination with the impact map, yielding a one degree of freedom hybrid system. Upon defining $z = (\theta, \sigma_Z)$, the hybrid zero dynamics are written as

$$
\begin{aligned}
\dot{z} &= f_{\text{zero}}(z), & z^- &\notin S \cap Z, \\
z^+ &= \begin{bmatrix} \theta^+ \\ \delta_{\text{zero}} \cdot \sigma_Z^- \end{bmatrix}, & z^- &\in S \cap Z,
\end{aligned}
\tag{16}
$$

where f_{zero} is given by (15) and δ_{zero} is a constant that is computed from restricting Δ to Z. This one DOF hybrid model is of the form (6).

Determining the existence and stability of limit cycles. The hybrid zero dynamics are extremely valuable as an analytical tool, and, as will be seen in the next subsection, also as a feedback design tool. Determining the existence and stability properties of limit cycles is typically a very difficult task for a system as complex as RABBIT. The standard application of Poincaré's method requires

knowledge of the solutions of the closed-loop system between two nine-dimensional surfaces; see Fig. 12 (a). Since the system is very nonlinear, the required computations can only be done numerically and do not yield much insight for feedback design. However, thanks to the hybrid zero dynamics, the stability analysis can be rigorously decomposed into two parts [49], [53]: first making sure that the feedback controller asymptotically drives the outputs to zero (see Fig. 12 (b)), which is a relatively easy thing to do, and then determining the existence and stability of limit cycles in the hybrid zero dynamics; see Fig. 12 (c)-(d). (References [49], [53] assume that the errors are driven to zero in finite time [55]; when the virtual constraints are invariant under the impact map, it seems that sufficiently fast exponential convergence is enough, though such a result has not been published.) Since the hybrid zero dynamics are two dimensional, a complete stability analysis is possible; indeed, the Poincaré map, ρ_{zero}, can be computed in closed form [53]. We present the main results of doing this in two steps:

a) there exists a periodic solution to the hybrid zero dynamics if, and only if, $\delta^2_{\text{zero}} \neq 1$ and

$$\frac{\delta^2_{\text{zero}}}{1 - \delta^2_{\text{zero}}} V_{\text{zero}}(\theta^-) + V^{max}_{\text{zero}} < 0, \tag{17}$$

where

$$\begin{aligned} V_{\text{zero}}(\theta) &:= -\int_{\theta^+}^{\theta} I(\xi) M g x_c(\xi) \cdot d\xi \\ V^{max}_{\text{zero}} &:= \max_{\theta^+ \leq \theta \leq \theta^-} V_{\text{zero}}(\theta). \end{aligned}$$

b) there exists an exponentially stable periodic solution if and only if (17) holds and

$$0 < \delta^2_{\text{zero}} < 1. \tag{18}$$

These formulas are mentioned for two reasons: i) they provide simple analytical expressions that can be profitably used in feedback design; and ii), because the stability conditions involve inequalities, they are robust to a certain amount of modelling error. Fig. 13 depicts a typical limit cycle induced by the controller.

A physical interpretation of the necessary and sufficient conditions for the existence of an exponentially stable orbit involves the essential interplay of kinetic and potential energy that is taking place throughout a step. Analyzing this may help to make the inherent robustness more clear [56]. The zero dynamics (15) are Lagrangian, with kinetic energy $K_{\text{zero}} := \frac{1}{2}\sigma^2_Z = \frac{1}{2}\left(I(\theta)\dot{\theta}\right)^2$, and potential energy $V_{\text{zero}}(\theta)$ [53]. (Note that the kinetic and potential energy terms of the zero dynamics are NOT the kinetic and potential energy terms of the robot model restricted to the zero dynamics manifold! As they have been written here, they do not have units of energy. This could be accomplished by scaling both terms by $I(\theta^-)$, for example.) Between impacts, the total energy $K_{\text{zero}} + V_{\text{zero}}$ is conserved along solutions of the zero dynamics [22]; it follows that energy may be gained or lost only at impacts. This property is similar to the energy conservation in the case of an inverted pendulum subject only to

gravity. For the studied gait, the angular momentum, σ_Z, is always positive. In the beginning of the single support phase, the center of mass of the robot is behind the support leg end. Thus, by (8) and (15), gravity initially decreases the angular momentum of the robot, and $V_{\text{zero}}(\theta)$ increases. If the angular momentum is not large enough, then the angular momentum goes to zero while the center of mass is still behind the support leg end, and, due to gravity, the robot falls backward. If the initial angular momentum is sufficiently large to overcome the potential energy barrier corresponding to V_{zero}^{max}, the center of mass will move past the support leg end, inducing the reverse exchange of energy, until the swing leg impacts the ground, see Fig. 14. An impact induces a change in the total energy in two ways. A constant change of V_{zero} occurs at impact, from $V_{\text{zero}}(\theta^-)$ at the end of the step to $V_{\text{zero}}(\theta^+)$ at the beginning of the step; see Fig. 14. The angular momentum changes also, through multiplication by δ_{zero}. From this, one can compute an angular momentum just before impact, σ_Z^*, that results in the conservation of the total energy *during the impact*, so that periodicity is enforced. Condition (17) stipulates that $\delta_{\text{zero}}\sigma_Z^*$ must be large enough to overcome the barrier posed by gravity V_{zero}^{max}. For the periodic orbit, the total energy has a constant value $V_{\text{zero}}(\theta^-) + \frac{1}{2}\left(\sigma_Z^*\right)^2$. Since the angular momentum is scaled by δ_{zero} at impact, the same is true of the difference between the angular momentum and its value on the periodic orbit, given by $\sigma_Z^- - \sigma_Z^*$. Thus, if angular momentum decreases at impacts, then it converges to σ_Z^*. Exponential stability is thus ensured by condition (18).

From the above analysis, it follows that once an exponentially stable orbit exists for the model of the robot, modeling errors will tend to destroy it only if they are sufficiently large to drive the angular momentum of the robot to zero before its center of mass is above the support leg end. Interpreted loosely, deliberate forward gaits, that is, gaits with a periodic motion that has significant angular momentum reserve at the point of maximum potential energy, will be quite robust; modelling error will significantly alter the average walking speed before it destabilizes the robot.

Designing a good walking motion. We have developed two approaches to designing virtual constraints that result in good walking motions, in the sense that they are energetically efficient, require low peak torques, respect all required contact conditions, and achieve a desired average walking speed. One method uses directly the hybrid zero dynamics, while the second only uses them indirectly to verify closed-loop stability. The direct method is presented first.

In [53], Bézier polynomials are used to parameterize the output (14), yielding

$$y = h_0(q) - h_d(\theta(q), a), \tag{19}$$

where a is a vector of real coefficients. The Bézier polynomials make it very easy to satisfy the invariance condition, so that the hybrid zero dynamics are guaranteed to exist. A cost function of the form

$$J(a) := \frac{1}{sl(T)} \int_0^T ||u^*(t, a)||^2 dt, \tag{20}$$

is posed, where T is the time to complete a step, $sl(T)$ corresponds to step length, and $u^*(t, a)$ is the vector of constraint torques that result from zeroing the output (19); see (13). A sequential, quadratic programming package is used to minimize $J(a)$ with respect to a, subject to a number of optimization constraints: the existence of an asymptotically stable orbit, i.e, inequalities (17) and (18); a desired walking rate; adequate contact conditions; maximum deflection of stance leg and swing leg knees; and actuator limitations. Whenever the optimization problem has a feasible solution, the result is an *asymptotically stable, closed-loop system* that meets natural kinematic and dynamic constraints. We have used this to design controllers that achieve asymptotically stable walking for a wide range of speeds [57].

The indirect method is based on first computing a periodic solution to the model equations that is optimal with respect to energy consumption, for example, and respects actuator limitations and contact conditions [58]. In a second step, output functions are constructed, based on the idea of virtual constraints, that are zero along the periodic solution [21]. The periodicity of the solution automatically guarantees that the resulting virtual constraints will be invariant under the impact map *and* that condition (17) is satisfied. Furthermore, for all trajectories analyzed to-date, condition (18) is also satisfied, which guarantees closed-loop stability.

An alternative design method. A quite different way to go from a periodic solution of the model equations to a time-invariant controller has been developed in [56]; see also [59], [60]. Consider a periodic solution of the model as a curve in the configuration space of the robot for a single step. The curve has a beginning and an end determined by the double support condition. Introduce a parameter, s, that is similar to arc-length in that $s = 0$ at the beginning of the curve and $s = 1$ at the end, with intermediate values of s parameterizing the posture of the robot, $q_d(s)$, as it progresses from the beginning of a step to the end. The condition $q(s) - q_d(s)$ defines the virtual constraints to be imposed by the control law. The freedom in how s itself evolves as a function of time, from its initial value of zero to its final value of one, can be used to augment the four joint torques (already available for control) with the acceleration $\ddot{s}$; this makes the system now look like it is fully actuated: five degrees of motion freedom and five controls. Consequently, a dynamic state feedback controller can be found that drives a vector of five outputs, $y = q(s) - q_d(s)$, asymptotically to zero.

An advantage of this approach is that a monotonic parameter that replaces time is automatically produced, so the control designer does not have to find one *a priori*. From a theoretical perspective, this idea may be especially useful for applying the method of virtual constraints to mechanisms with many more degrees of freedom than RABBIT. A potential disadvantage is that, since the evolution of s must be determined from the model, it is unclear how sensitive the closed-loop system may be to model uncertainty. Further work may clarify this issue. We will experiment with this method on the prototype.

Initial Walking Experiments Under HZD Control

A controller based on virtual constraints and the hybrid zero dynamics (HZD) was implemented on RABBIT, and it led to successful walking on the first try! To our knowledge, this is the first implementation of an analytically derived control law on a bipedal robot with actuated knees and torso, as well as the first time such a robot has successfully walked on the very first attempt. The first two sets of walking experiments that we have completed and the experimental set-up are described here.

Experimental environment. To implement the HZD control algorithm, a dSPACE system was selected as the real-time control platform. With dSPACE, run-time software is created by automatic translation and cross-compiling of SIMULINK diagrams, allowing the controller software to be developed in a high-level language. In addition, dSPACE provides low-level computation, digital-to-analog and analog-to-digital conversion, as well as a user interface, all in a single package. This obviates the need for low-level I/O programming and facilitates debugging. The control computations were performed with a sample period of 1.5 ms (667 Hz).

During the first set of experiments, RABBIT's workspace was limited to approximately 120° of rotation about the central tower, and the lateral support bar was set at 1.5 m. This allowed RABBIT to take at most nine or ten steps, which is not enough to demonstrate stability of the walking motion. This will be addressed in the second set of experiments.

The system model versus reality. Modelling error is a concern with any control scheme. As RABBIT is a complex mechanism, there are many sources of error between the idealized Lagrangian dynamics and the actual system. These include:

• **friction** in the motor-belt-gear reducer system that is used at each joint, as well as friction at the universal joint of the tower that supports the power electronics and the dSPACE system;

• **unmodeled flex dynamics** caused by cabling, torsion of the gears in the joints, and flexing of the counter-balance bar;

• **parameter inaccuracies** due to poor estimates of link inertias and the additional inertia of the tower that contains the dSPACE module and power electronics;

• **non-rigid impacts** due to compliance at the end of the leg and in the walking surface; and

• **digital implementation issues** such as sampling effects, quantization, velocity estimation from position measurements.

Our control designs have not explicitly taken into account any of these potential problems. Hence, will the closed-loop system be sufficiently robust to these sorts of uncertainties? To the extent possible, this was verified through a detailed simulation model of the robot where many of these model imperfections were applied to the robot in closed-loop with the HZD controller. For example, it was easy to check that friction in the knee and hip joints could be overcome with a robust design of the

controller used to zero the outputs. Simulation work reported in [21] showed that significant changes in the impact model may result in the robot's average speed changing, but asymptotic stability was apparently maintained. An exhaustive study of the effects of changing all mass and inertia parameters by ±20% showed a similar result, in that, while the average walking speed and peak torques generally changed, stability was always preserved. The same was true for applying a constant force at the hips to simulate static friction in the universal joint of the central tower. These simulation studies supported our qualitative analysis of the closed-loop stability conditions (recall Fig. 14) and gave us the confidence to proceed with the initial control experiments even prior to completion of the parameter identification phase of the project.

Controller selection and implementation. As described previously, the hybrid zero dynamics are a finitely parameterized, low-dimensional description of the closed-loop dynamics during walking. This enables us to tune the walking motions for efficiency, while respecting kinematic constraints, such as minimum hip height and step length, and dynamic constraints, such as walking rate and stability. We designed an HZD controller with a step length of 0.39 m and an average forward velocity of 0.75 m/s. An animation of the designed walking motion is available at [1]. The controlled variables are the relative knee angles and the relative angles between the torso and femurs. In other words, the controlled variables $h_0(q)$ correspond to the actuated variables.

For all of the reported experiments, the outputs were zeroed by independent, joint-level, PD control; see Fig. 15. This was feasible due to the high gear ratio of 50:1 between the motors and the links [61, p.85]. The use of a decoupled controller for zeroing the outputs had a number of benefits. Between experiments, RABBIT would be secured with the feet off the ground, since, without power, the mechanism would obviously fall. It turned out that a set of gains for the PD controllers could be found that resulted in stable operation for all three possible contact conditions: double support (initial pose on the ground), single support (normal walking) and flight (robot secured in the air). Finally, swing leg touchdown was determined by measuring the forces applied to the leg ends. When touchdown was detected, the robot's coordinates were permuted so that the same controller could be used independently of whether the inside or outside leg was the current stance leg.

The first walking experiment. Prior to placing the robot on the ground, basic implementation issues such as measurement conventions, amplifier gains, and tuning of the PD controller gains were resolved. The experiment presented here was the first experiment where RABBIT was placed on the ground and required to take more than one step forward or backward.

The experiment began with the robot elevated in the air. The links were servoed to a nominal open stance position; the state feedback controller was then activated and the robot was placed on the ground. The closed-loop system is asymptotically stable in double support with zero forward velocity ($\dot{\theta} = 0$). To initiate walking, an external force was applied to RABBIT by means of pulling on a rope

attached to its hip. Application of this force, while crude in the way it was applied, was a means to provide the robot with an initial forward velocity sufficient to enter the basin of attraction of the 0.75 m/s controller. The velocity of the hips at the point of maximum hip height for the first step was approximately 0.38 m/s.

To our delight, once walking was initiated, RABBIT continued to walk until the limited workspace was exhausted; this corresponded to 1.89 m or five steps. To prevent RABBIT from walking into a wall, an external force was applied; one of us simply caught RABBIT's torso in his hand, bringing the robot to a stop. Again, this was crude, but effective. We immediately arranged for room in the lab and repeated the experiment, this time allowing the robot to travel 3.4 m or nine steps. While this limited range of motion does not validate the stability of the closed-loop system, our simulation experience tells us that, if the mechanism is going to fall, it usually does so by the third step or so. Videos of the both experiments are available at [1].

For the initial experiments, the robot was equipped with sensors to measure the tangential and normal components of the forces exerted on each leg end. The support leg and swing leg are therefore easily distinguished, as are the single support and double support phases; this facilitates the logic required for leg commutation. Fig. (16) shows typical normal forces recorded during our walking experiments. The flat sections correspond to the leg being off the ground. From these plots, it can be seen that the double support phase is very short (less than 20 ms), indicating that a rigid impact model seems to be appropriate. After several experiments, the force senors degraded; apparently they were unable to withstand the forces due to impact. This degradation can be seen in the drifting zero values recorded for the forces.

A nice bonus: reverse is built into the controller. When RABBIT is under HZD control, an additional feature that may be a bit unexpected at first glance is the ability to make RABBIT walk backward simply by pushing on the torso; a video of this is available at [1]. This ability arises because the hybrid zero dynamics are parameterized by the robot's unactuated state, θ. By applying a sufficiently large force on the robot's torso opposite to the direction of forward walking, θ decreases; recall that we only have to overcome the moment arm due to gravity in order to stop the robot's center of mass from rotating about the stance leg end, which stops the robot from advancing. The feedback controller still enforces the virtual constraints, so RABBIT nicely raises the swing leg and walks backward! We believe that the ability to stop the robot's forward motion by simply applying a force on the torso also increases the safety of performing experiments on RABBIT.

The second walking experiment. In this experiment, the robot has been moved to a laboratory that permitted us to use a 3 m bar. The force sensors have been removed and contact switches installed at the leg ends for the purpose of detecting impact with the ground. A nominal controller was designed to achieve a walking speed of 0.7 m/s. The robot was first raised off the floor and servoed to a nominal

open stance position; the state feedback controller was then activated and the robot was placed on the floor. To initiate walking, one of us simply gripped the robot's torso and gave it a push forward. This provided adequate forward velocity to attain the basin of attraction of the closed-loop system. The robot completed more than 70 steps before being deliberately stopped; for a video, see [1]. The speed of the robot was deliberately varied by adjusting the angle of the torso (increased forward lean provides increased walking speed). Figs. 17 and 18 report typical time-traces for the joint angles and the commanded torques while the torso angle is held constant. It is seen that the virtual constraints are being accurately enforced. From the video, it is clear that the closed-loop system consisting of the robot and the HZD controller has achieved asymptotically stable walking.

Initial lessons learned. From a control design point of view, it was very important that algorithm development started at least three years before the robot was ready for testing. This gave us time to consider new ideas and lay the foundation for a completely new theory of walking for this particular biped. Had the robot been ready sooner, there would have been too much temptation to just try stuff, leading to a hack-attack approach to control design.

The experiments have confirmed that our basic approach to control design yields a remarkably robust closed-loop system. The domain of attraction of our designed limit-cycles is sufficiently large that a simple push on the torso suffices to initiate walking. The friction present in the joints due to the gear reducers consumes over 15% of the power applied through the motors. However, this has almost no effect on the walking motion as long as the virtual constraints are adequately imposed via feedback. We saw on the mechanism that the virtual constraints could be very accurately imposed through decentralized PD control plus friction compensation (recall Fig. 17). The control law is relatively insensitive to parameter variations in the single support model; in an experiment to be reported elsewhere, we increased the mass of the torso by 10 Kg; stable walking was maintained, though the averaging walking speed was perturbed. For a given set of virtual constraints, Fig. 14 and formula [53, Eq. 55] suggest that average walking speed should vary significantly with the amount of energy lost at impact. Consequently, unless the impact model is accurate, the achieved walking speed will be different than the designed walking speed, and this is what we observed. Furthermore, the laboratory floor is concrete with channels covered by metal grates. We have noted differences in walking speed as the robot passed over the grates. More effort is needed to improve the impact model used in our work.

RABBIT and the overall experimental facility have performed exceptionally well in the initial round of tests. Of course, improvements are always possible. The 3 m lateral stabilization bar needs to be stiffer; we plan to either shorten the bar or build another one of other material. In addition, the inertia of the bars and cables is significant and needs to be taken into account in the model of the robot. RABBIT's shins were designed to be very close together in order to better approximate a planar robot;

in hind sight, they are maybe too close as vibrations from hard impacts can cause them to knock into one another; an easy solution is to widen the hips by 2 or 3 cm. A new set of force sensors will be investigated for the leg ends. Doing leg commutation on the basis of contact switches proved to be quite feasible, but it lacks the flexibility that can be achieved when more than binary information is available; also, it is more difficult to improve the impact model without a measurement of the forces on the leg ends. At one point in our experiments, the encoder that measures the orientation of the torso with respect to the vertical, as established by the lateral stabilization bar, began to rotate with respect to the bar. This drift had a significant destabilizing effect on the closed-loop system. Other means of establishing the orientation of the torso need to be investigated.

Beyond Walking at a Fixed Rate

There is of course more to bipedal robot control than just walking or running at a fixed rate. How to start and stop the robot, how to change walking speeds, and how to transition between walking and running come immediately to mind. In addition, motions which reverse directions present special challenges and motivate the study of periodic balancing. We very briefly discuss how our work extends to some of these tasks.

Walking at multiple speeds. The work presented on the virtual constraints and the hybrid zero dynamics provided analytical control solutions for stable walking at a fixed rate. In [57], we have extended these results to design a controller that takes the robot from a standing position, through a range of walking speeds, and back to a standing position, while providing local stabilization and disturbance rejection. The controller designs are based on two basic ideas: switching control and PI control. The switching controller design is motivated by an idea presented in [62]; the main additional challenge is that the domains of attraction of any two of the individual HZD-controllers may have an empty intersection, and hence a transition controller was required to steer the robot from the region of attraction of one HZD-controller into the region of attraction of a second, nearby, HZD-controller. This provided for walking at multiple discrete speeds. Walking at a continuum of speeds was obtained via an event-based PI controller. The novelty here is that the controller uses integral action to adjust the parameters in an exponentially stabilizing HZD-controller. Parameter adjustment takes place just after impact (swing leg touching the ground). The analysis of the controller is based on the restricted Poincaré map of the hybrid zero dynamics.

On a closely related though slightly different topic, [54] has investigated methods to enhance the rate of convergence to a limit cycle. The analysis, which is also based on the Poincaré map, has shown that both torso angle and step length adjustments are effective for accomplishing this.

Balancing. Balancing consists, to some extent, in viewing the robot in single support as a multi-link inverted pendulum. The goal is to find a feedback control law that induces a non-trivial, limit cycle in the high-dimensional inverted pendulum. We may also wish to move from one limit cycle

to another, possibly by just changing the period or the amplitude of the motion, or both. The key difficulty is that the admissible limit cycles are not completely free: they must respect the constraints imposed by the physics of the workspace and contact forces at the leg end. These restrictions disqualify feedbacks leading to the natural homoclinic orbits (orbits with constant energy levels), that may induce inadmissible motions out of the joint workspace.

What makes this control problem quite different from walking is that ground impacts are not considered in balancing. At first glance, this may seem to simplify the problem, but, upon further reflection, this is not the case. The difficulty lies in the fact that the class of stable, periodic motions that can be achieved by balancing seems to be much smaller that what can be achieved through allowing impacts [35].

The balancing problem has been studied in [63]. Employing the concept of virtual constraints developed earlier, an output of the form $y = h_0(q) - h_d(\theta(q), a)$ is formed and zeroed by a feedback controller. In this case, the resulting zero dynamics is no longer hybrid because there are no impacts. The balancing problem may be restated as finding a set of parameters a that induce a stable orbit. As opposed to walking, where the parameters were held constant during a step, here, a is viewed as an additional control variable (dynamic state feedback) that can be adjusted on line so as to match an exo-system that describes the desired orbit. This idea, which is still under investigation, also allows for the possibility of moving from one orbit to another in a stable manner.

Conclusions

The theoretical and experimental work completed so far has proven that it is possible to achieve asymptotically stable walking on a flat surface without an actuated ankle. This may be a first step toward the realization of bipedal walking robots that are mechanically simpler and less costly than existing robots.

The underactuation of the robot in single support led us to a complete re-thinking of the control problem. Part of the solution was to track "trajectories" that were parameterized by the state of the robot instead of by time; more precisely, these should be called invariant manifolds or surfaces and not trajectories. In this regard, the notions of virtual constraints and hybrid zero dynamics were developed. This led to a control design that simplified the demands placed on the robot, since it did not have to re-synchronize with time after a perturbation, as well those placed on the control designer, by reducing analysis and design to a two-state model. The stability of the closed-loop system was analytically verifiable through easy computations.

The work conducted in the context of this project has shed light on many things that were previously unknown to us. These include the important role of gravity in generating walking motions for a robot such as RABBIT, and how to quantify stability in terms of the tradeoff of potential energy and kinetic energy throughout a periodic orbit. This deeper physical understanding of the behavior of the robot

under closed-loop control led us to experiment with simpler feedback implementations than those used in the analysis, and this led to satisfying experimental results: both because the robot walked well and because we understood the stability mechanism of the controllers. The properties of the invariant surfaces that are created by our controllers were exploited in the development of feedback strategies for changing walking speeds and to accelerate the rate of convergence to the periodic orbit. These strategies will be evaluated soon on the prototype.

We plan to extend our work to more complex problems, such as walking on an uneven surface and avoiding obstacles, both of which require a modification of the periodic walking cycles that we have studied so far. We also plan to study three dimensional walking, with one or two degrees of underactuation. And, finally, we will not limit ourselves to only robots without feet, but will eventually consider integrating foot placement and ankle torque control with various aspects of walking, running and balancing.

Acknowledgments

Appendix

Consider the model (6) and output (14). Differentiating twice yields

$$\ddot{y} = L_f^2 h(x) + L_g L_f h(x) u,$$

where, $L_f h(x) := \frac{\partial h(x)}{\partial x} f(x)$, $L_f^2 h(x) := \frac{\partial L_f h(x)}{\partial x} f(x)$, and $L_g L_f h(x) := \frac{\partial L_f h(x)}{\partial x} g(x)$. It is easy to check that since (6) comes from (4), the decoupling matrix, $L_g L_f h$ is given by $\frac{\partial h}{\partial q}(q) D^{-1}(q) B$, and thus only depends on the configuration variables. If the decoupling matrix is invertible on a given open set $\mathcal{O}$ [51], then several things hold:

• $Z := \{x \in \mathcal{O} \mid h(x) = 0, L_f h(x) = 0\}$ is a smooth submanifold; that is, the constraint surface is nicely regular.

• Setting $\ddot{y} = 0$ yields the torques necessary to identically null the outputs, and hence the torques necessary to satisfy the constraints:

$$u^*(x) = -(L_g L_f h(x))^{-1} L_f^2 h(x).$$

• Let $\Psi(y, \dot{y})$ be such that $\ddot{y} = \Psi(y, \dot{y})$ has a desired evolution. Then this is achieved on the model (6) by

$$u(x) := (L_g L_f h(x))^{-1} \left(\Psi(h(x), L_f h(x)) - L_f^2 h(x) \right).$$

Remark: For many interesting choices of the output, all of the indicated matrix inversions can be avoided in the computation of the control.

REFERENCES

[1] Supplemental material for this article. available at [64]. April 2003.

[2] K. Berns. The walking machine catalogue. *available at http://www.uwe.ac.uk/clawar/*, September 2002.

[3] M. Vukobratovic, B. Borovac, D. Surla, and D. Stokic. *Biped Locomotion.* Springer-Verlag, Berlin, 1990.

[4] K. Hirai, M. Hirose, Y. Haikawa, and T. Takenake. The development of Honda humanoid robot. In *Proc. of the IEEE International Conference on Robotics and Automation, Leuven, Belgium*, pages 1321–1326, May 1998.

[5] A. Goswami. Postural stability of biped robots and the foot-rotation indicator (FRI) point. *International Journal of Robotics Research*, 18(6):523–533, June 1999.

[6] ROBEA: Robotics and Artificial Entities. http://www-lag.ensieg.inpg.fr/prc-bipedes/index.html (in French; for an English version, click on the British flag). January 2003.

[7] T. McGeer. Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82, 1990.

[8] S.H. Collins, M. Wisse, and A. Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research*, 20(7):607–615, July 2001.

[9] A. D. Kuo. Stabilization of lateral motion in passive dynamic walking. *International Journal of Robotics Research*, 18(9):917–930, 1999.

[10] T. McGeer. Stability and control of two-dimensional biped walking. Technical Report 1, Center for Systems Science, Simon Fraser University, Burnaby , B.C., Canada V5A 1S6, 1988.

[11] J. Pratt and G. Pratt. Intuitive control of a planar bipedal walking robot. In *Proc. of the IEEE International Conference on Robotics and Automation, Leuven, Belgium*, pages 2014–2021, May 1998.

[12] MIT Leg Lab. http://www.ai.mit.edu/projects/leglab. September 2002.

[13] R.S. Wallace. Miniature direct drive rotary actuators II: Eye, finger and leg. *RAS*, 13:97–105, 1994.

[14] J.M. Hollerbach, I.W. Hunter, and J. Ballantyne. *A comparative analysis of actuator technologies for robotics.* MIT Press. Khatib, O. and Canny, J., and Lozano-Perez, T., 1992.

[15] P.H. Channon, S. Hopkins, and Pham. Optimal walking motions for a biped walking robot. *Robotica*, 10(2):165–172, 1990.

[16] G. Abba and N. Chaillet. Robot dynamic modeling using a power flow approach with application to biped locomotion. *Autonomous Robots*, 6(1):39–52, January 1999.

[17] F. Bullo and K. M. Lynch. Kinematic controllability for decoupled trajectory planning in underactuated mechanical systems. *IEEE Transactions on Robotics and Automation*, 17(4):402–412, August 2001.

[18] C. Chevallereau and P. Sardain. Design and actuation optimization of a 4 axes biped robot for walking and running. In *Proc. of the IEEE International Conference on Robotics and Automation, San Francisco, California*, pages 3365–3370, April 2000.

[19] E. Muybridge. *The Human Figure in Motion.* Dover Publications, New York, 1955. Reprinted in 1989 by Chapman and Hall, Ltd., London, First Edition.

[20] S. Miossec and Y. Aoustin. Walking gait composed of single and double supports for a planar biped without feet. In *Conference on Climbing and Walking Robots*, 2002.

[21] F. Plestan, J.W. Grizzle, E.R. Westervelt, and G. Abba. Stable walking of a 7-DOF biped robot. In *IEEE Transactions on Robotics and Automation*, January 2003 (accepted). to appear; see [64] for a preprint.

[22] H. Goldstein. *Classical Mechanics.* Addison Wesley, second edition, 1980.

[23] M.W. Spong and M. Vidyasagar. *Robot dynamics and control.* John Wiley and Sons, New York, 1991.

[24] E. Dombre and W. Khalil. *Modeling, identification and control of robots.* Hermes Sciences, Europe. Paris, 2002.

[25] D. W. Marhefka and D. Orin. Simulation of contact using a nonlinear damping model. In *Proc. of the IEEE International Conference on Robotics and Automation, Minneapolis, MN.*, pages 1662–1668, April 1996.

[26] C. Canudas-de-Wit, L. Roussel, and A. Goswani. Periodic stabilization of a 1-DOF hopping robot on nonlinear compliant surface. In *Proc. of IFAC Symposium on Robot Control, Nantes, France*, pages 405–410, September 1997.

[27] C. Canudas-de-Wit, H. Olsson, K.J. Astrom, and P. Lischinsky. A new-model for control of systems with friction. *IEEE Transactions on Automatic Control*, 40(3):419–425, 1995.

[28] P.A. Lischinsky. *Compensation de frottement et commande en position d'un robot hydraulique industriel.* PhD thesis, Institut National Polytechnique de Grenoble, 1992. in French.

[29] L. Roussel. *Génération de trajectoires de marche optimales pour un robot bipède.* PhD thesis, Institut National Polytechnique - Grenoble - France, November 1998. in French.

[30] R.M. Murray, Z. Li, and S. Sastry. *A Mathematical Introduction to Robotic Manipulation.* CRC Press, Boca Raton, 1993.

[31] G. Zonfrilli, F. Oriolo and D. Nardi. A biped locomotion for the quadruped robot sony ers-210. In *Proc. of the IEEE International Conference on Robotics and Automation, Washington, DC*, pages 2768–2774, May 2002.

[32] A. Formal'sky. *Locomotion of antropomorphic mechanisms.* Nauka. in Russian, Moscow, 1982.

[33] Y. Hurmuzlu and D.B. Marghitu. Rigid body collisions of planar kinematic chains with multiple contact points. *International Journal of Robotics Research*, 13(1):82–92, 1994.

[34] H. Ye, A.N. Michel, and L. Hou. Stability theory for hybrid dynamical systems. *IEEE Transactions on Automatic Control*, 43(4):461–474, April 1998.

[35] Y. Hurmuzlu and D. Moskowitz. The role of impact in the stability of bipedal locomotion. *Dynamics and Stability of Systems*, 1(3):217–234, 1986.

[36] M.W. Spong. Passivity based control of the compass gait biped. In *Proc. of IFAC World Congress, Beijing, China*, July 1999.

[37] A. D. Kuo. Energetics of actively powered locomotion using the simplest walking model. *Journal of Biomechanical Engineering*, 124:113–120, 2002.

[38] M. Raibert. *Legged robots that balance.* MIT Press, Mass., 1986.

[39] A. Takanishi. Humanoid robots and animal robots - towards entertainment robot market in 21$^{\text{st}}$ century. In *International Symposium on Robotics*, Seoul, Korea, April 2001.

[40] G. Pratt. Legged robots at MIT: what's new since Raibert? *IEEE Robotics and Automation Magazine*, September 2000.

[41] M.H. Raibert. Legged robots. *Communications of the ACM*, 29(6):499–514, 1986.

[42] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. Planning and control of a juggling robot. *International Journal of Robotics Research*, 13(2):101–118, 1994.

[43] C. Francois and C. Samson. A new approach to the control of the planar one-legged hopper. *International Journal of Robotics Research*, 17(11):1150–1166, 1998.

[44] U. Saranli, W.J. Schwind, and D.E. Kodistschek. Toward the control of a multi-jointed, monopod runner. In *Proc. of the IEEE International Conference on Robotics and Automation, Leuven, Belgium*, pages 2676–2682, May 1998.

[45] K. Ono, K. Yamamoto, and A. Imadu. Control of giant swing motion of a two-link horizontal bar gymnast robot. *Advanced Robotics*, 15(4):449–465, 2001.

[46] M.W. Spong and F. Bullo. Controlled symmetries and passive walking. In *Proc. of IFAC World Congress, Barcelona, Spain*, July 2002.

[47] A. V. Roup, D. S. Bernstein, S. G. Nersesov, W. M. Haddad, and V. Chellaboina. Limit cycle analysis of the verge and foliot clock escapement using impulsive differential equations and poincare maps. In *Proc. Amer. Contr. Conf., Arlington, VA*, pages 3245–3250, June 2001.

[48] A. Banaszuk and J. Hauser. On control of planar periodic orbits. In *Proceedings of the 38th IEEE Conference on Decision and Control, Philadelphia, PE*, pages 3830–3836, December 1999.

[49] J.W. Grizzle, G. Abba, and F. Plestan. Asymptotically stable walking for biped robots: Analysis via systems with impulse effects. *IEEE Transactions on Automatic Control*, 46:51–64, January 2001.

[50] E. Westervelt, J.W. Grizzle, and D.E. Koditschek. Zero dynamics of planar biped walkers with one degree of under actuation. In *IFAC 2002, Bareclona, Spain*, July 2002.

[51] A. Isidori. *Nonlinear Control Systems: An Introduction.* Springer-Verlag, Berlin, third edition, 1995.

[52] Y. Aoustin and A. Formal'sky. Stability of a cyclic biped gait and hastening of the convergence to it. In *CLAWAR*, September 2001.

[53] E. Westervelt, J.W. Grizzle, and D.E. Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1):42–56, January 2003.

[54] Aoustin Y. and Formal'sky A. Control design for a biped: Reference trajectory based on driven angles as functions of the undriven angle. *Journal of Computer and Systems Sciences International*, 42(4), 2003.

[55] S.P. Bhat and D.S. Bernstein. Continuous finite-time stabilization of the translational and rotational double integrators. *IEEE Transactions on Automatic Control*, 43(5):678–682, 1998.

[56] C. Chevallereau, A. Formal'sky, and D. Djoudi. Tracking of a joint path for the walking of an underactuated biped. *Robotica*, 2003. accepted for publication; pre-print available at [65].

[57] E.R. Westervelt, J.W. Grizzle, and C. Canudas de Wit. Switching and PI control of walking motions of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(2):308–312, February 2003.

[58] C. Chevallereau and Aoustin. Optimal reference trajectories for walking and running of a biped robot. *Robotica*, 19(5):557–569, September 2001.

[59] C. Chevallereau. Parameterized control for an underactuated biped robot. In *IFAC 2002, Bareclona, Spain*, volume 19 No. 2, pages 362–368, July 2002.

[60] C. Chevallereau. Time scaling control for an underactuated biped robot. *IEEE Transactions on Robotics and Automation*, 18, April 2003.

[61] A.J. Critchlow. *Introduction to Robotics*. Macmillan, New York, 1985.

[62] R. Burridge, Rizzi A., and D.E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555, June 1999.

[63] C. Canudas-de-Wit, B. Espiau, and C. Urrea. Orbital stabilization of underactuated mechanical systems. In *Prof. of IFAC World Congress, Bareclona, Spain*, July 2002.

[64] J.W. Grizzle. Publications on robotics and control. *http://www.eecs.umich.edu/∼grizzle/papers/robotics.html*, April 2003.

[65] C. Chevallereau. Robotics team at IRRcYN. *http://www.ircyn.ec-nantes.fr/irccyn/d/en/equipes/Robotique/Membres*, April 2003.

Component	Model and/or Size	Manufacturer
DC motors	RS420J	Parvex SA, Dijon, France
Motor current drives	RS420 RTS10/20-60	Parvex SA, Dijon, France
Incremental Encoders (motors)	C4, 250 counts/rev	Parvex SA, Dijon, France
Absolute Encoders (joints)	CHM 506 P426R/8192/16 8192 counts/rev	Ideacod, Strasbourg, France
Incremental Encoders (central tower)	GHM5	Ideacod, Strasbourg, France
Gear Reducers	HFUS-2UH, size: 25 ratio: 1/50	Harmonic Drive Technologies, Peabody, MA, US
Real-time Controller	DS1103 400 MHz PowerPC 604e DSP	dSpace, Paderborn, Germany

TABLE I

Components used in RABBIT.

	femur	tibia	torso
length in m	0.4	0.4	0.625
mass in kg	6.8	3.2	12.0
inertia in $kg\ m^2$	0.47	0.20	1.33

TABLE II

Experimentally measured parameters for RABBIT.
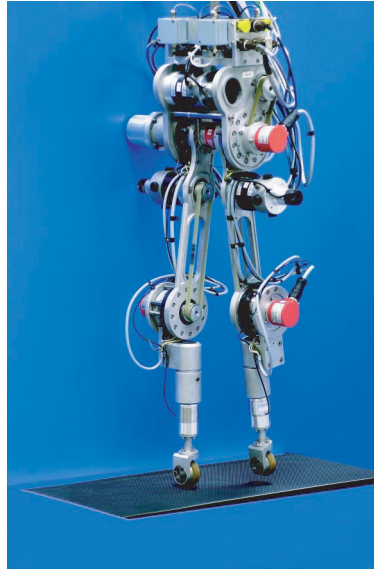
Fig. 1. Photo of RABBIT. The robot was designed to facilitate the development of theoretically-sound control algorithms for walking, running and balancing.
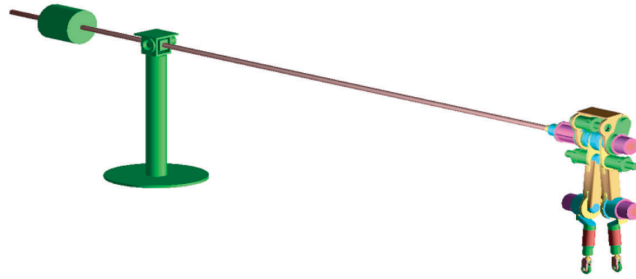
Fig. 2. RABBIT Walking in a Circle, While Looking Like a Planar Biped. The radial bar is attached at the robot's hip, via a revolute joint. The bar only provides lateral stabilization; it does not prevent the robot from falling forward, backward or down. The counter balance can be used to offset the weight of the lateral stabilization bar or to modify the effective gravitational field. Not shown are the dSPACE module and power electronics that are mounted on top of the central tower.
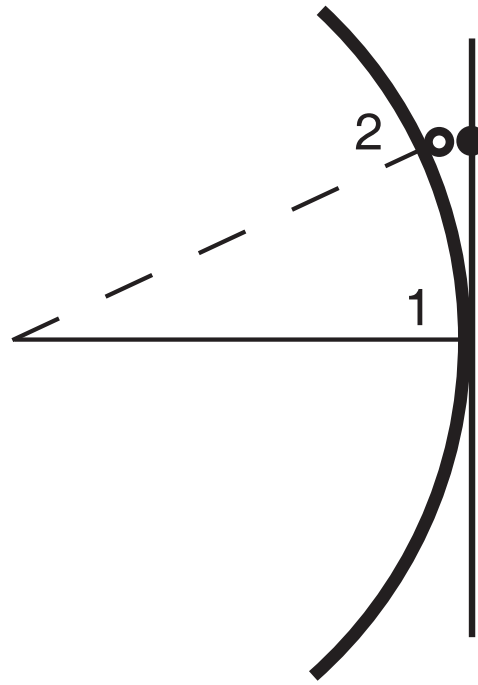
Fig. 3. Top View of RABBIT Walking in a Circle. To see why wheels in the frontal plane are necessary on the leg ends, consider the robot when the hip is in position 1, and the stance leg is in front of the robot, as marked by the solid dot. The leg end must lie on the robot's sagittal plane, which is tanget to the circle, and thus the leg end is not on the circle. However, as the robot advances to point 2 where the hips are now over the stance leg end, the leg end now must touch the circle as shown by the un-filled dot. This mobility is supplied by a wheel that is directed normal to the sagittal plane of the robot. A related but less significant effect is associated with changes in the height of the hip; this bit of geometry is left to the reader.
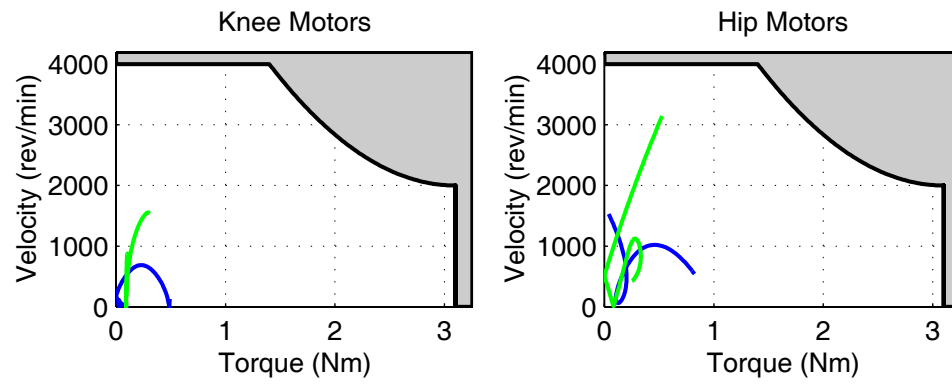
Fig. 4. Plot of Speed Versus Torque for an Optimal Walking Motion. The average speed is 0.75 m/s; the gear ratio is 50:1.
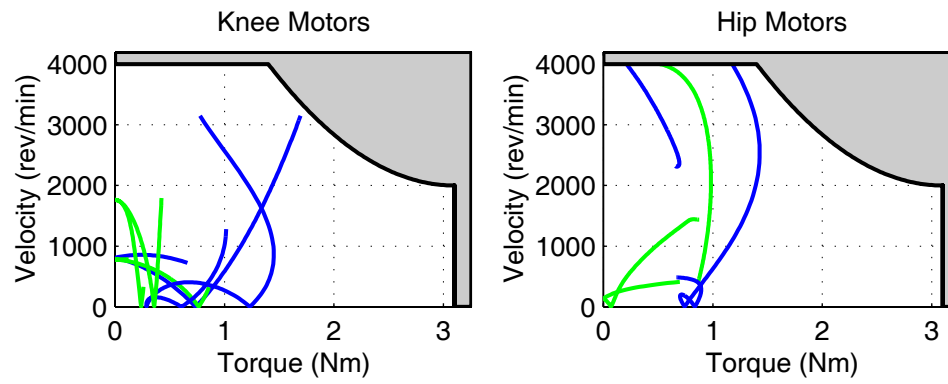
Fig. 5.   Plot of Speed Versus Torque for an Optimal Running Motion. The average speed is 1.2 m/s; the gear ratio is 50:1.
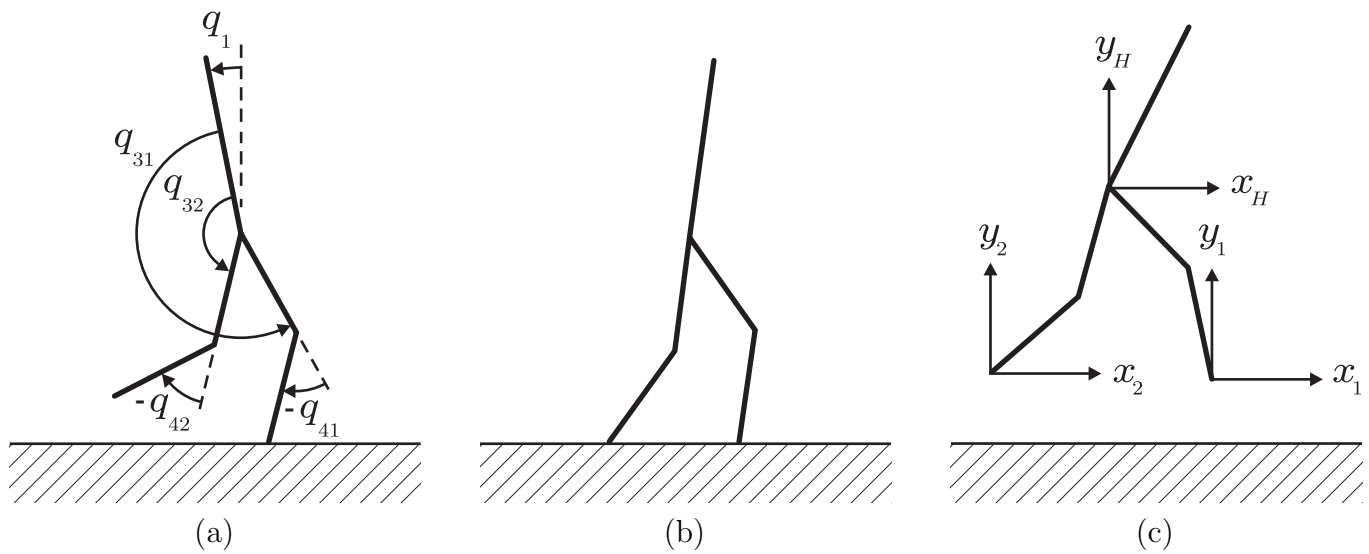
Fig. 6. Different Contact Conditions with the Walking Surface. The robot is shown in (a) single support, (b) double support, and (c) flight phases. The torso is leaned backward in single support to show a positive q_1. To avoid clutter, the coordinate labels have been spread out over the single support and flight phases, even though they apply to all three phases. The length of the tibias will be denoted by L_1 and the length of the femurs by L_2.

$$\varphi(x) = 0$$

$$\dot{x} = f(x) + g(x)u$$

$$x^+ = \Delta(x^-)$$
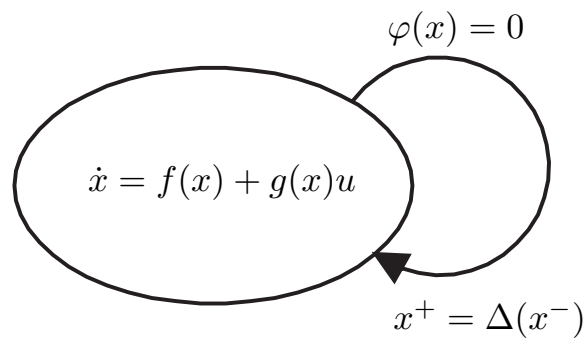
Fig. 7. Single-mode Hybrid Model Corresponding to RABBIT. Key elements are the continuous dynamics of the single support phase, the switching or impact condition, $\varphi = 0$, which detects when the height of the swing leg above the walking surface is zero, and the re-initialization rule coming from the impact map, Δ.

(a)   (b)
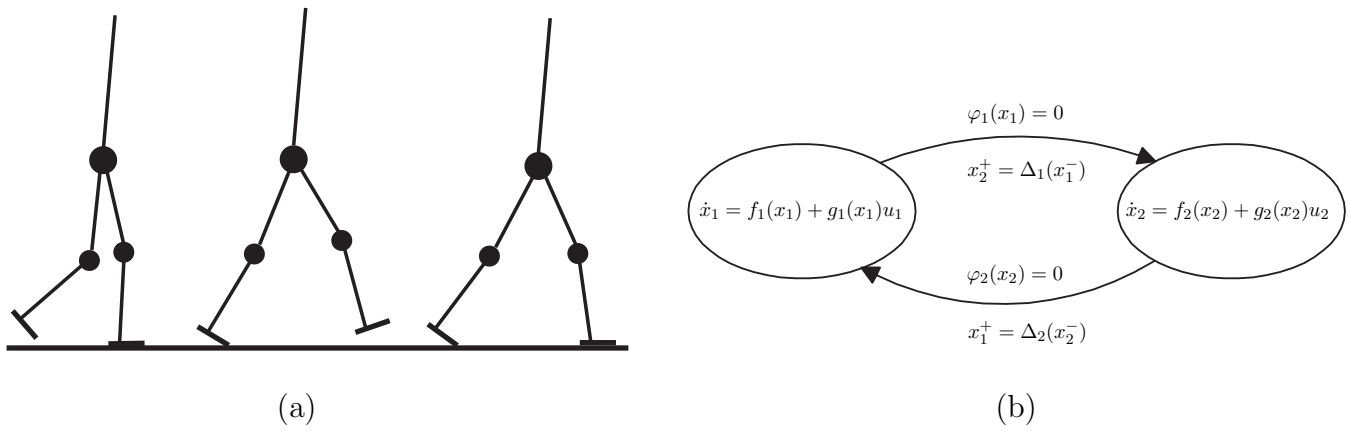
Fig. 8.   Multiple-mode Hybrid Model. (a) A robot with feet in single support with the foot flat on the ground, single
support with rotation of the foot about the toe, and double support. (b) If rotation of the foot is allowed, a hybrid
model of the robot has more than one mode to account for the different contact conditions of the foot with the
ground. Shown here is a two mode model corresponding to the gait depicted in part (a).
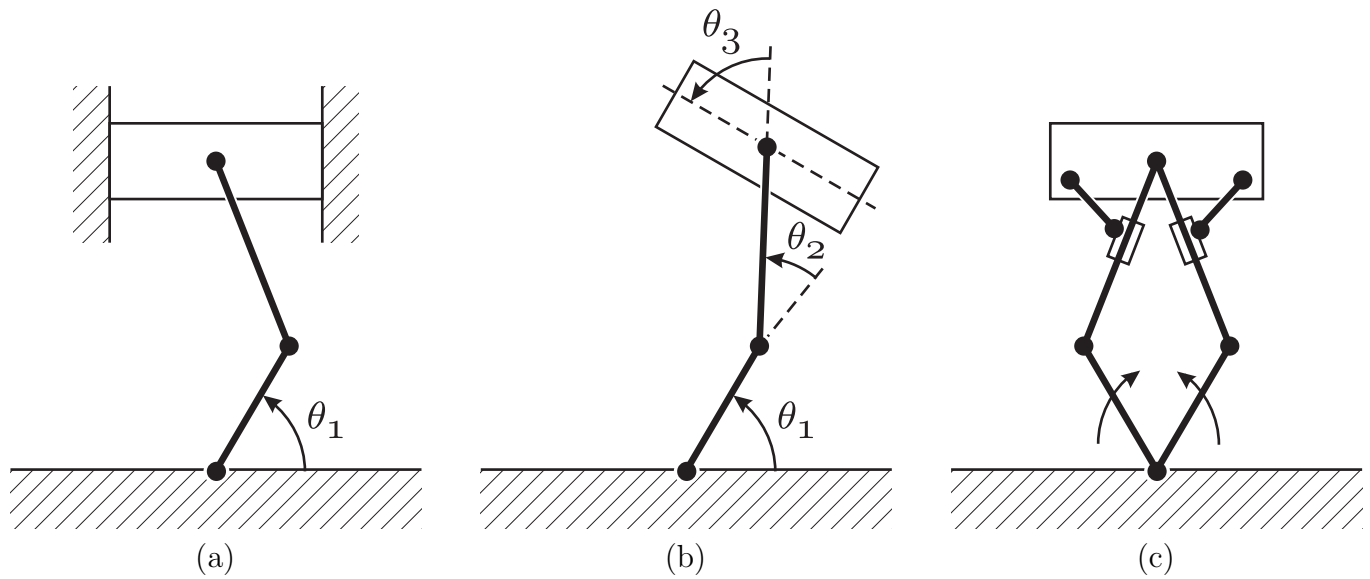
Fig. 9.   Virtual Constraints in a Simpler Context. The dots represent revolute joints. (a) Piston constrained to move in a cylinder; this is a one degree of freedom mechanical system. (b) Piston without the constraints; this is a three degree of freedom mechanical system. (c) Hypothetical, Rube Goldberg realization of the piston constrained via additional links to evolve kinematically as the piston in a cylinder; the curved arrows represent the two cranks rotating synchronously in opposite directions. By realizing link coordination on the robot via virtual constraints, different gaits can be more easily programmed than if the links were coordinated by hardware constraints.
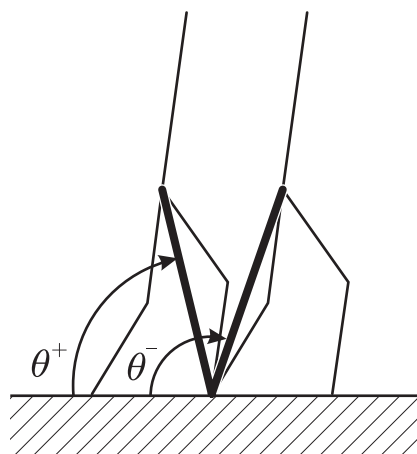
Fig. 10. Virtual Support Leg. The virtual support leg is the fictive link (bold line) between the end of the support leg and the hip. The angle of the virtual leg at the beginning of a step is denoted θ^+ and its value at the end is θ^-, commensurate with the notation used in the hybrid model. Note that as the robot advances forward, θ increases.
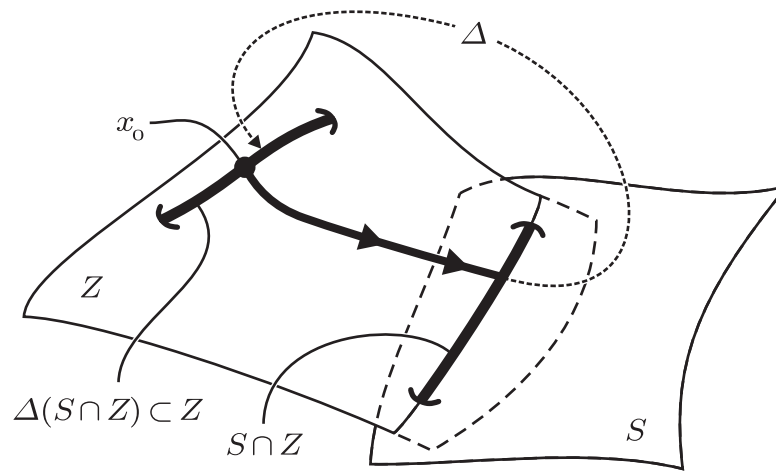
Fig. 11. Invariance of the Constraint Surface Under the Action of the Impact Map. In this case, the hybrid zero dynamics can be defined. They provide a low-dimensional subsystem for studying stability of the closed-loop system. Recall that Z is the constraint surface (i.e., zero dynamics manifold), S is the impact surface, and Δ is the impact map.

(a)

(b)

(c)

(d)
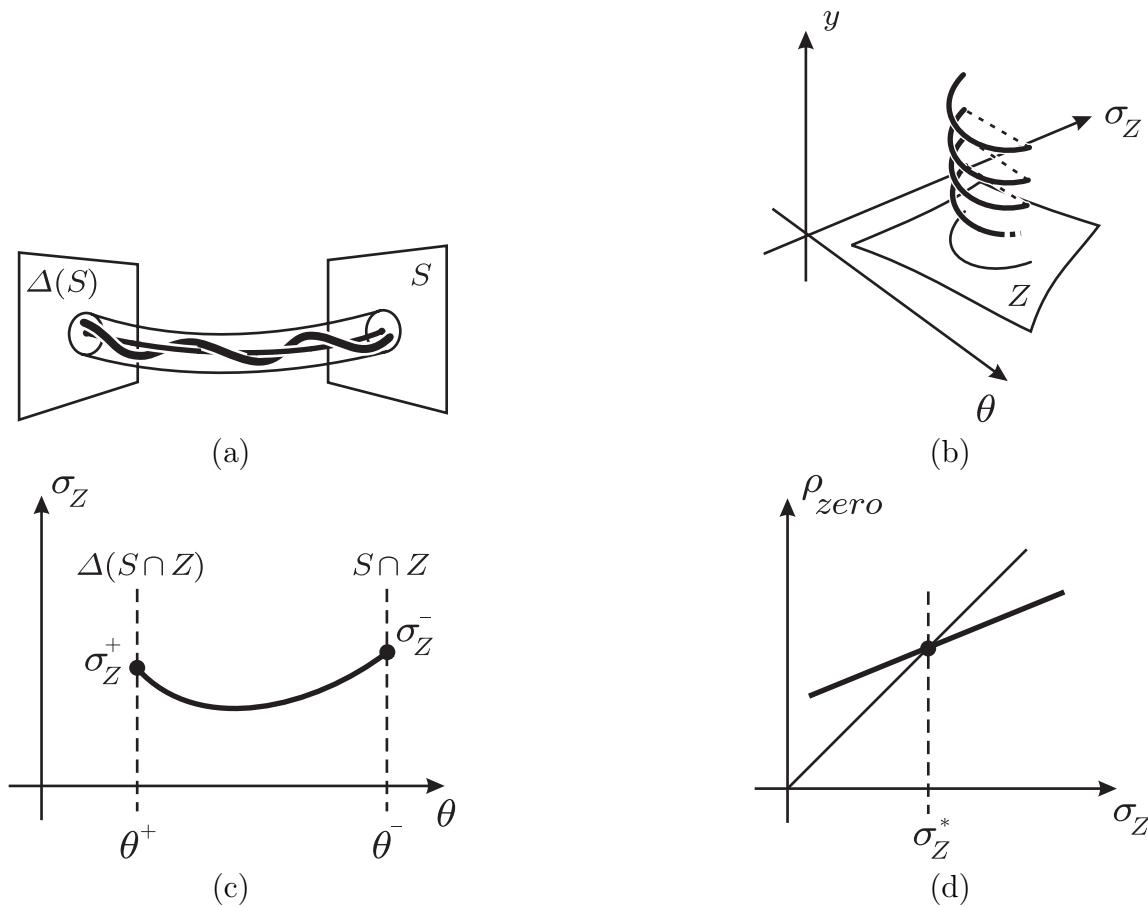
Fig. 12.   The Method of Poincaré Sections. The stability analysis of periodic motions is normally performed by computing the Poincaré return map on the full-order model, which is very difficult to do. Thanks to the hybrid zero dynamics, an analysis of the Poincaré return map becomes tractable and insightful. (a) Poincaré surfaces for the full order hybrid model of the biped; the surfaces S and $\Delta(S)$ are each nine dimensional. Computing the Poincaré return map requires the flow from S to $\Delta(S)$, which is shown in the tube. (b) Conceptualization of the outputs being driven to zero so that the system asymptotically converges to the two-dimensional surface, Z, where the closed-loop system evolves according to the hybrid zero dynamics. (c) A phase plot of a trajectory of the hybrid zero dynamics; a fixed-point (i.e., a limit cycle) occurs when $\delta_{\mathrm{zero}}\sigma_Z^- = \sigma_Z^+$. (d) Typical Poincaré return map for the one degree of freedom hybrid zero dynamics of RABBIT; the fixed-point, σ_Z^*, is the angular momentum just before impact of a stable periodic motion; the initial angular momentum would be $\delta_{\mathrm{zero}}\sigma_Z^*$.
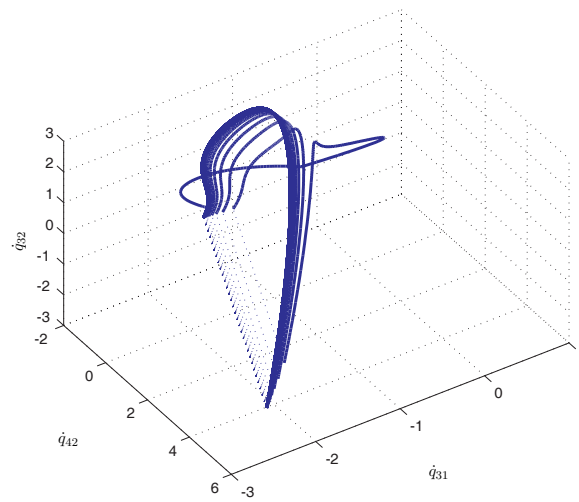
Fig. 13. Projection of a Limit Cycle into a 3-D Slice of the State Space. The simulation shows the controller responding to a perturbation by driving the robot's trajectory back to the limit cycle. The dotted portion is the jump in the trajectory resulting from the impact map.
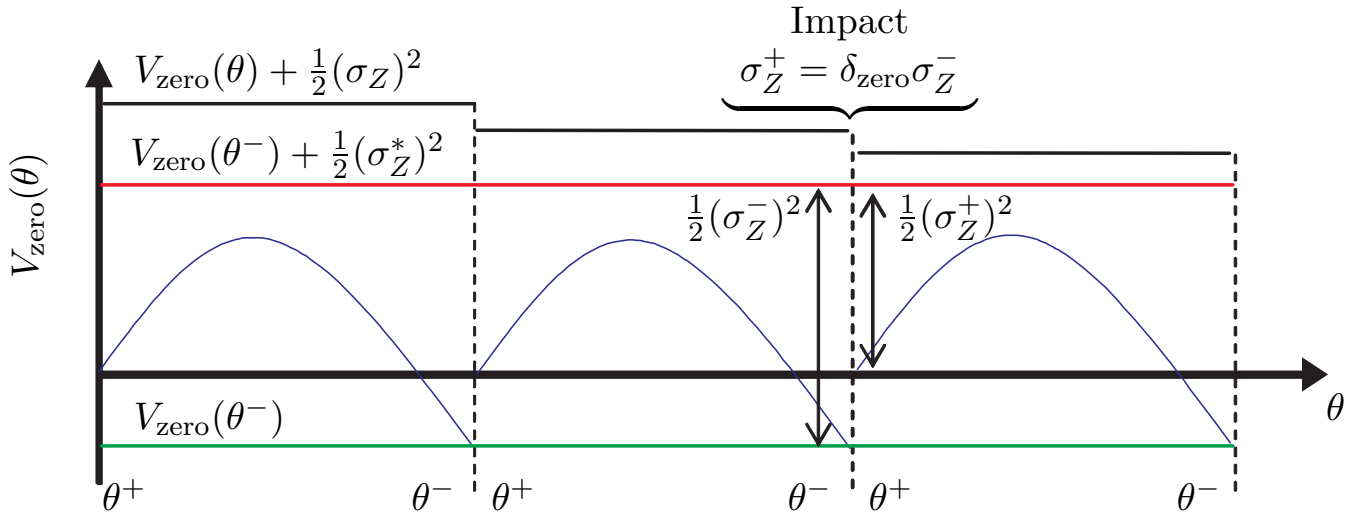
Fig. 14.   A Qualitative Look at Stability Through Energy. The zero dynamics is Lagrangian, and thus throughout the single support phase, the corresponding total energy $V_{\text{zero}}(\theta) + \frac{1}{2}\sigma_Z^2$ is constant. At impact, the change in total energy depends on the angular momentum through $\delta_{\text{zero}}\sigma_Z^-$ and the potential energy through $V_{\text{zero}}(\theta^-)$. The total energy corresponding to the periodic orbit is $V_{\text{zero}}(\theta^-) + \frac{1}{2}(\sigma_Z^*)^2$. Convergence to this total energy level occurs if the angular momentum decreases during impact, namely, $\delta_{\text{zero}} < 1$. From the expression for the existence of a periodic orbit, $\delta_{\text{zero}} < 1$ is equivalent to $V_{\text{zero}}(\theta^-) < 0$.
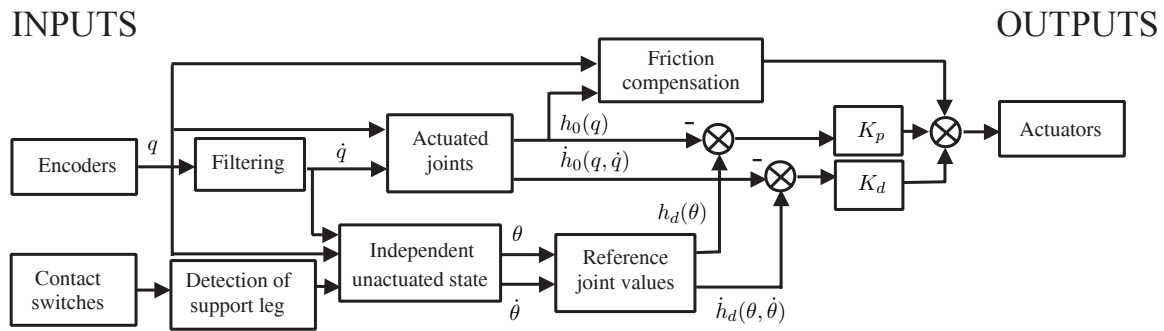
Fig. 15. Block Diagram of the Control System. Contact switches are shown for detecting impacts and determining the support leg. The same structure holds when force sensors are used in place of contact switches.
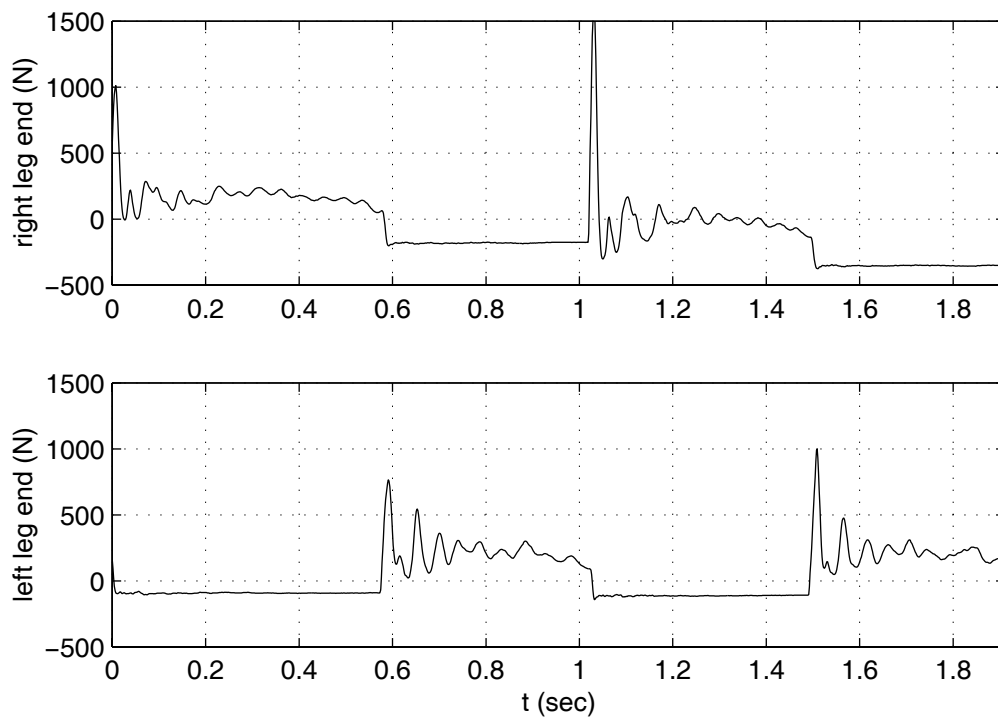
Fig. 16.   Normal Ground Reaction Forces as Measured by the Force Sensors. The flat sections correspond to the leg end being off the ground. From these plots, it can be seen that the double support phase is very short (less than 20 ms), indicating that a rigid impact model is appropriate.
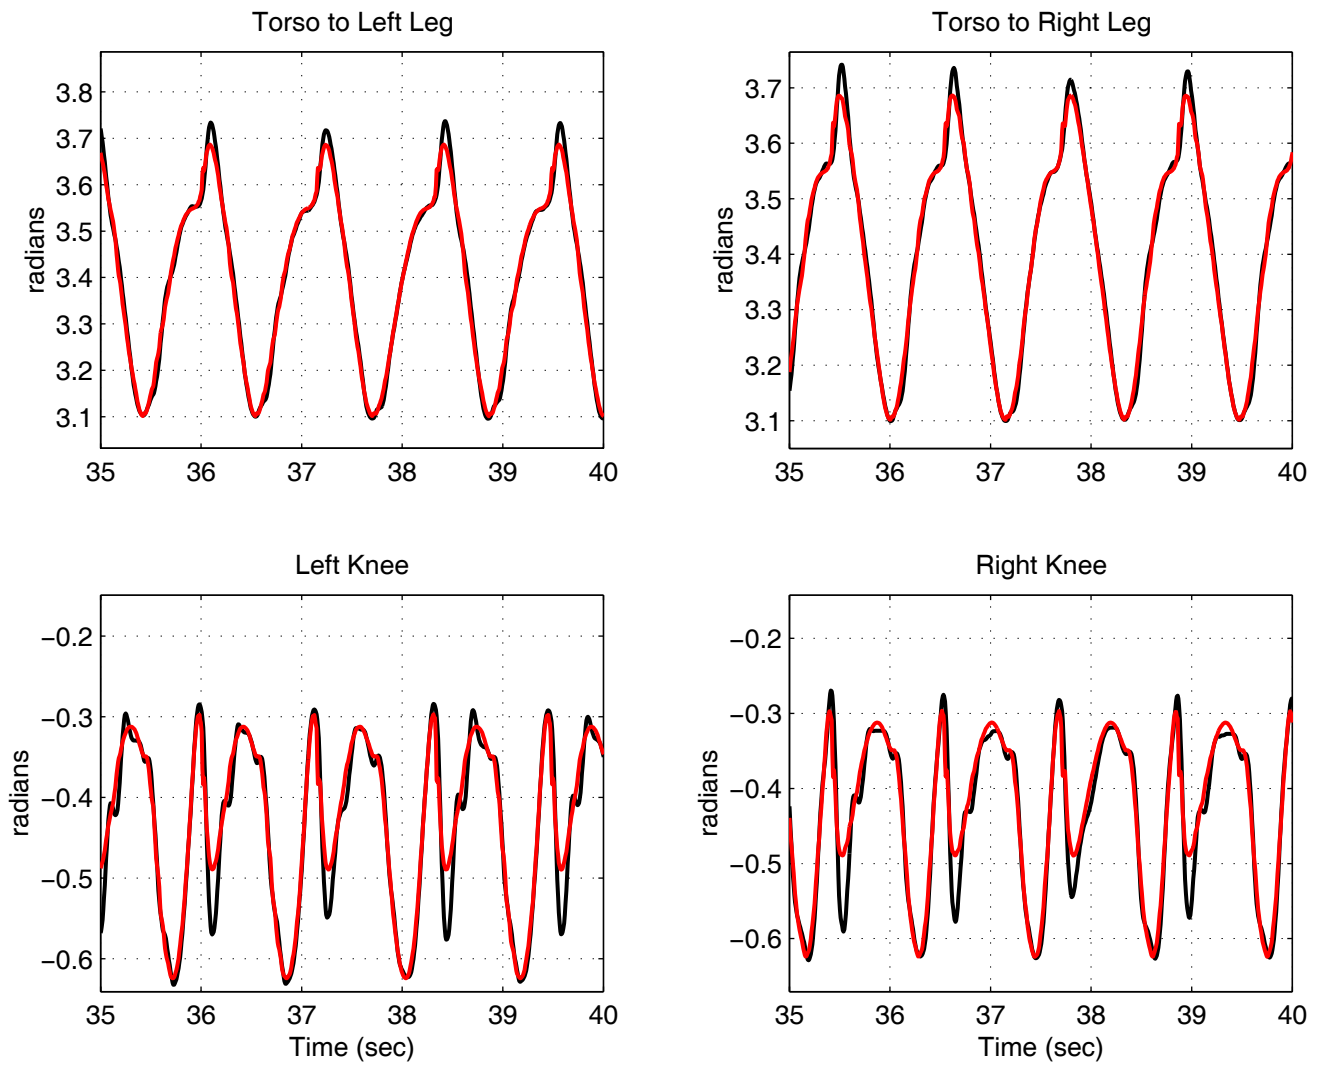
Fig. 17. Comparison of the Measured Joint Trajectories (black) with Their Target Values (red) Arising from $h_d(\theta)$.
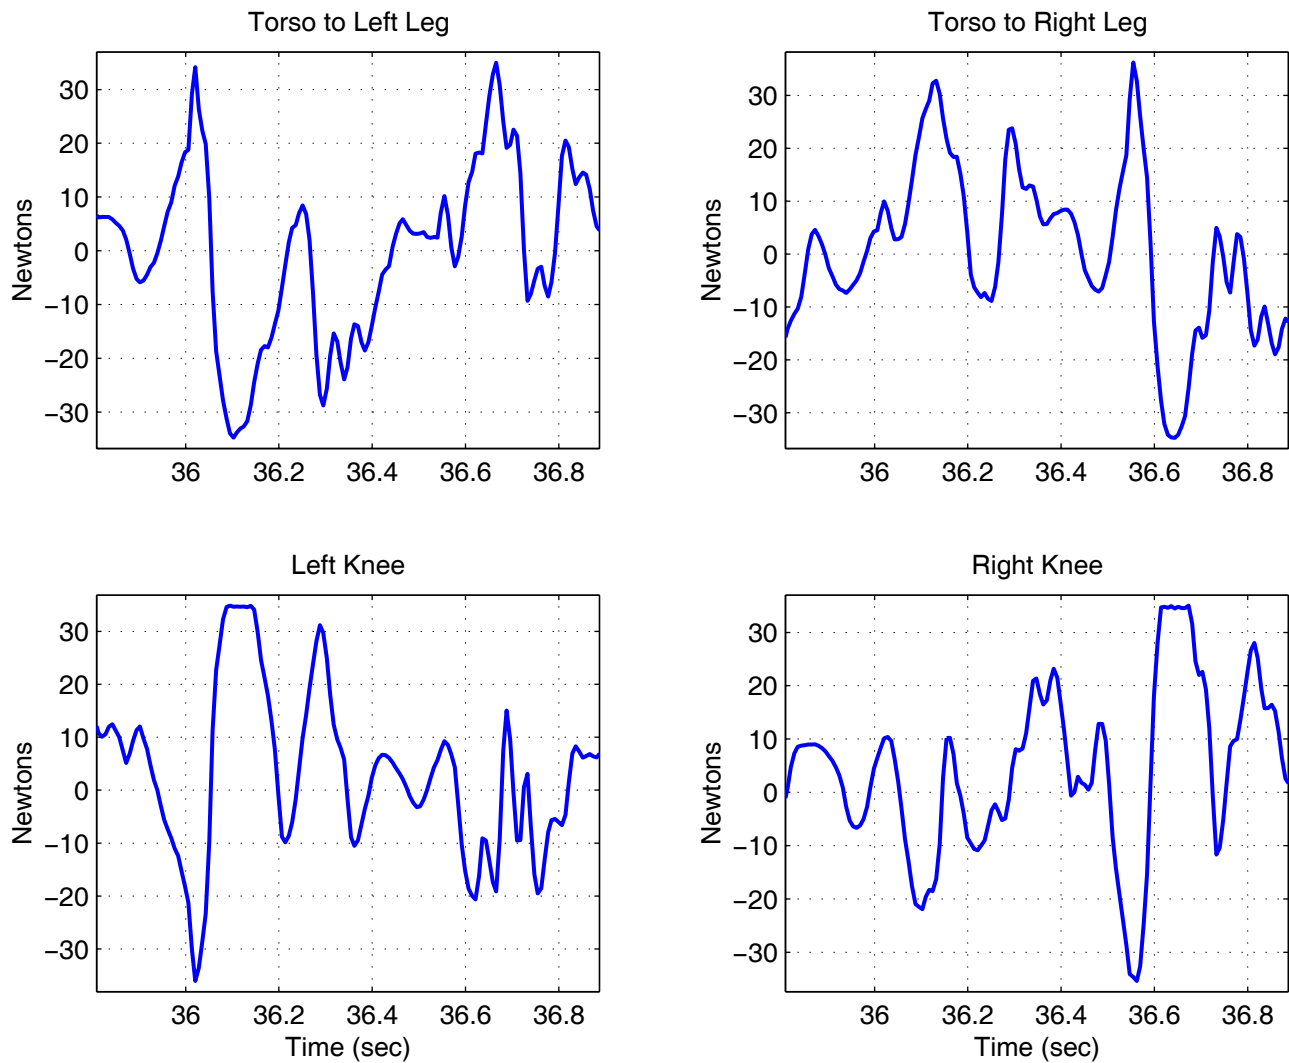
Fig. 18.   Plot of the Commanded Control Values for Each Motor Over Two Steps. The peak torque for each knee occurs during the stance phase on the corresponding leg. The relatively large torque during the swing phase is due to friction in the gear-reducers. For these experiments, the controls were purposefully limited to 35 Nm in order to enhance safety, both of the robot and the experimenters.

Author Addresses and Biographies

C. Chevallereau

IRCCyN, Ecole Centrale de Nantes, UMR CNRS 6597, BP 92101, 1 rue de la Noë, 44321 Nantes cedex 03, France, `Christine.Chevallereau@irccyn.ec-nantes.fr`

G. Abba

LGIPM, Université de Metz, IUT de Thionville, Espace Cormontaigne, 57970 Yutz, France, `abba@iut.univ-metz.fr`

Y. Aoustin

IRCCyN, Ecole Centrale de Nantes, UMR CNRS 6597, BP 92101, 1 rue de la Noë, 44321 Nantes cedex 03, France, `Yannick.Aoustin@irccyn.ec-nantes.fr`

F. Plestan

IRCCyN, Ecole Centrale de Nantes, UMR CNRS 6597, BP 92101, 1 rue de la Noë, 44321 Nantes cedex 03, France, `Franck.Plestan@irccyn.ec-nantes.fr`

E.R. Westervelt

Mechanical Engineering Department, Robinson Laboratory, 206 W. 18th Ave., Columbus, OH 43210-1154, `westerve@mecheng.osu.edu`

C. Canudas de Wit

Laboratoire d'Automatique de Grenoble, INPG-ENSIEG, UMR CNRS 5528, BP. 46, 3402, Saint Martin d'Heres, France, `canudas@lag.ensieg.inpg.fr`

J.W. Grizzle

Control Systems Laboratory, Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109-2122, USA, `grizzle@umich.edu`
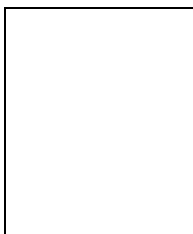
Christine Chevallereau graduated from Ecole Nationale Supérieure de Mécanique, Nantes in 1985, and received the Ph.D. degree in Control and Robotics from Ecole Nationale Supérieure de Mécanique, Nantes, France in 1988. Since 1989, she has been with the CNRS in the Institut de Recherche en Communications et Cybernetique de Nantes. Her research interests include modeling and control of robots, especially control of robot manipulators and legged robots.
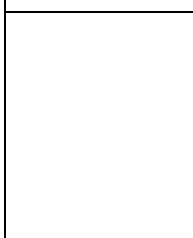
Gabriel Abba joined the Ecole Normale Supérieure de Cachan, France, in 1979, and received the *agrégation* of the Ministry of Education in Electrical Engineering in 1982. He received the Doctorate degree in Electronics and Robotics from the University of Paris XI-Orsay, Paris, France, in 1986. Since 1999, he has been with the University of Metz, where he is professor and head of the Industrial Systems and Maintenance Department and director of the Industrial Engineering and Mechanical Production Laboratory. His research interests include development, modeling and control of robots, especially control of legged robots, visual servoing, and control of high speed drives and spindles for industrial applications. He received the 2002 George S. Axelby award for the best paper published in the Transactions on Automatic Control during the years 2000 and 2001.
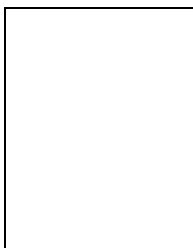
Yannick Aoustin received the Ph.D degree from the University of Nantes, France in 1989. Currently, he is an Associate Professor at the Institute of Research in Communication and Cybernetics of Nantes. His research interests include flexible robots, and, walking and running robots.
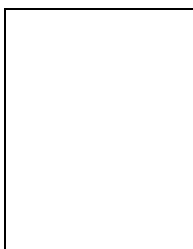
Franck Plestan received the Ph.D. in Automatic Control from the Ecole Centrale de Nantes, France, in 1995. From September 1996 to August 2000, he was with the Université Louis Pasteur, Strasbourg, France. In September 2000, he joined the Ecole Centrale de Nantes, Nantes, France where he is currently an Assistant Professor. His research interests include nonlinear control, theoretical aspects of nonlinear observer design and control of electromechanical and mechanical systems. He received the 2002 George S. Axelby award for the best paper published in the Transactions on Automatic Control during the years 2000 and 2001.
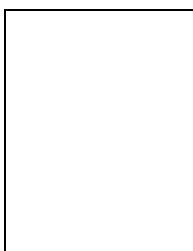
Eric R. Westervelt received the B.S. degree in computer and systems engineering from Rensselaer Polytechnic Institute, Troy, NY, and the Ph.D. degree in electrical engineering from the University of Michigan, Ann Arbor, in 1993 and 2003, respectively. He is currently an Assistant Professor with the Department of Mechanical Engineering at The Ohio State University. His research interests include the theory and practice of control as applied to mechanical systems. Further information is available at http://www.mecheng.ohio-state.edu/~westerve/
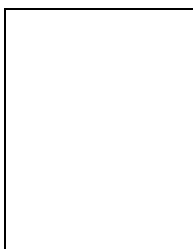
Carlos Canudas-de-Wit was born in Villahermosa, Tabasco, Mexico in 1958. He received his B.Sc. degree in electronics and communications from the Technological Institute of Monterrey, Mexico in 1980. In 1984 he received his M.Sc. in the Department of Automatic Control,Grenoble, France. He was a visiting researcher in 1985 at Lund Institute of Technology, Sweden. In 1987 he received his Ph.D. in automatic control from the Polytechnic of Grenoble (Department of Automatic Control), France. Since then he has been working at the same department as"Directeur de recherche at the CNRS", where He teaches and conducts research in the area of adaptive robot and automotive control. He is also director of a group on Control of Electro-mechanical Systems and Robotics. His research topics includes: adaptive control, identification, robot control, nonlinear observers, control of systems with friction, AC and DC drives, automotive control, and nonholonomic systems. More recent topics include networked control systems. A description of these topics can be found at www-lag.ensieg.inpg.fr/canudas/. He has written a book in 1988 on Adaptive Control of Partially Known Systems:

Theory and Applications (Elservier Publisher). He has also edited a book in 1991 on Advanced Robot Control (Springer-Verlag), a joint 12 author book "Theory of robot control", on Springer Control and Communication Series, in (1997). He has been Associate Editor of the IEEE-Transaction on Automatic Control, from Jan 1992, to Dec. 1997, and of AUTOMATICA, from 1999, to 2002. He teaches undergraduate and graduate courses in robot control and stability of nonlinear systems.
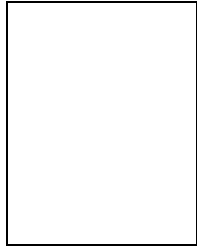
Jessy W. Grizzle received the Ph.D. in electrical engineering from The University of Texas at Austin in 1983. Since September 1987, he has been with The University of Michigan, Ann Arbor, where he is a Professor of Electrical Engineering and Computer Science. His research interests lie in the theory and practice of nonlinear control. He has been a consultant in the automotive industry since 1986, where he jointly holds fourteen patents dealing with emissions reduction through improved control system design. Professor Grizzle is a past Associate Editor of the Transactions on Automatic Control and Systems & Control Letters, and is currently an Associate Editor for Automatica. He served as Publications Chairman for the 1989 CDC, from 1997 through 1999 served on the Control Systems Society's Board of Governors, and was Chair of the IEEE Control Systems Society Fellows Solicitation Committee from 2000 through 2003. He was a NATO Postdoctoral Fellow from January to December 1984; he received a Presidential Young Investigator Award in 1987, the Paper of the Year Award from the IEEE Vehicular Technology Society in 1993, the University of Michigan's Henry Russell Award for outstanding research in 1993, a College of Engineering Teaching Award, also in 1993, elected to Fellow of the IEEE in 1997, and received the 2002 George S. Axelby Award for the best paper published in the IEEE Transactions on Automatic Control during the years 2000 and 2001.