

Stable dynamic walking over uneven terrain

Ian R Manchester¹, Uwe Mettin², Fumiya Iida³ and Russ Tedrake¹

Abstract

We propose a constructive control design for stabilization of non-periodic trajectories of underactuated robots. An important example of such a system is an underactuated “dynamic walking” biped robot traversing rough or uneven terrain. The stabilization problem is inherently challenging due to the nonlinearity, open-loop instability, hybrid (impact) dynamics, and target motions which are not known in advance. The proposed technique is to compute a transverse linearization about the desired motion: a linear impulsive system which locally represents “transversal” dynamics about a target trajectory. This system is then exponentially stabilized using a modified receding-horizon control design, providing exponential orbital stability of the target trajectory of the original nonlinear system. The proposed method is experimentally verified using a compass-gait walker: a two-degree-of-freedom biped with hip actuation but pointed stilt-like feet. The technique is, however, very general and can be applied to a wide variety of hybrid nonlinear systems.

Keywords

Dynamics, legged robots, underactuated robots, transverse linearization, nonlinear hybrid systems, feedback control, rough terrain.

1. Introduction

It has long been a goal of roboticists to build a realistic humanoid robot. Clearly, one of the most fundamental abilities such a robot must have is to walk around its environment in a stable, efficient, and naturalistic manner. When one examines the current state of the art, it seems that one can have either *stability and versatility* or *efficiency and naturalism*, but not all four. This paper reports some recent efforts to bridge this gap.

We propose a general method of exponentially stabilizing a very general class of motions for underactuated robots. In particular, we develop a provably-stable feedback control strategy for efficient “dynamic walking” bipeds over uneven terrain. The major novelty of this approach is that, for the first time in dynamic-walking control, we can provably stabilize arbitrary non-periodic trajectories arriving in real-time from an online motion planner. In this paper we propose the algorithm, prove that it is stabilizing under some reasonable assumptions, and discuss in-depth connections to other methods of control design, such as full-state linear quadratic regulator (LQR) and hybrid zero dynamics (HZD). Furthermore, we verify that the proposed method is feasible and effective in a simple experiment on a compass-gait biped robot.

1.1. Bipedal Walking Robots

The world of bipedal walking robots can be divided into two broad classes. The first, including well-known robots such as the Honda ASIMO and the HRP-2, are controlled using the “zero moment point” (ZMP) principle (see e.g. Vukobratovic and Borovac (2004) and references therein). The main idea is that the center of pressure always remains within the polygon of the stance foot, so the foot always remains firmly planted on the ground. Satisfaction of this principle ensures that all dynamical degrees of freedom remain fully actuated at all times, and control design can be performed systematically using standard tools in robotics. However, the motions which are achievable are highly conservative, inefficient, and unnatural looking.

¹Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, USA

²Department of Engineering Cybernetics, Norwegian University of Science and Technology, Norway

³Institute of Robotics and Intelligent Systems, ETH Zurich, Switzerland

Corresponding author:

Ian R Manchester, Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
Email: irm@mit.edu

The second broad class consists of passive-dynamic walkers and limit-cycle walkers. Inspired by the completely passive walkers of McGeer (1990), these robots forgo full actuation and allow gravity and natural dynamics to play a large part in the generation of motion. They may be completely passive, or partially actuated. Even with partial actuation, the motions generated can be life-like and highly efficient energetically (Collins et al., 2005). However, there is presently a lack of tools for systematic control design and stability analysis, especially for non-periodic motions.

Comparatively little work has been done as yet on walking over uneven terrain, especially for underactuated dynamic walkers. Passivity-based control and controlled-symmetry methods (Asano et al., 2004; Spong and Bullo, 2005) have been analyzed in the context of changing ground slopes (Spong and Bhatia, 2003; Spong et al., 2007), however these results rely on actuators at every link to reshape the potential energy and are thus not applicable to the underactuated walkers we consider in this paper. In Yang et al. (2009), stability of aperiodic walking motions was studied via the virtual-constraints and hybrid zero dynamics approach which achieves stability via high-gain feedback (Grizzle et al., 2001; Westervelt et al., 2003, 2007). The problem of footstep planning has been approached using computational optimal control (Byl and Tedrake, 2009) and experimental studies have shown that a minimalistic open-loop control can achieve stability for the compass-gait walker (Iida and Tedrake, 2009). Recently, more complete planning and control systems have been developed for quadruped walkers – see e.g. Rebula et al. (2007), Byl et al. (2008), and Kolter et al. (2008) – and template-based control has yielded impressive results with a hexapod robot (Saranli et al., 2001). In Shkolnik et al. (2010) a bounding motion over rough terrain was generated for a quadruped, and stabilized using a transverse-linearization technique similar to the method in this paper.

To give the present paper context, in Figure 1 we depict a possible organizational structure for the perception and control of a dynamic walker on uneven terrain. The main components are:

1. **Terrain Perception:** fusion of sensors such as vision, radar, and laser, perhaps combined with pre-defined maps, generating a model of the terrain ahead.
2. **Motion Planning:** uses the terrain map, current robot state, and a model of the robot's dynamics to plan a finite-horizon feasible sequence of footstep locations and joint trajectories. Slow time-scale: motion plan might be updated once per footstep.
3. **Motion Control:** feedback control to stabilize the planned motion in the face of inaccurate modeling, disturbances, time delays, etc. Fast time-scale: typically of the order of milliseconds.
4. **Robot State Sensing and Estimation:** sensors and state estimation algorithms to provide information about the physical state of the robot to all other modules.




Fig. 1. Possible organization of perception and control of a walking robot.

A complete humanoid robot would have all of these components, and many others. In this paper, we focus our attention on the third component: motion control. That is, we assume that the terrain has been sensed and a motion plan generated, and the task remaining is to compute a stabilizing feedback controller which reliably achieves this motion. It is essential for our algorithm that it be paired with a good motion planner. One example would be the algorithm proposed in Shkolnik et al. (2010), which successfully planned bounding motions for a quadruped over very rough terrain. To the authors' knowledge, the proposed method is the first in the literature that can stabilize arbitrary dynamic walking motions coming in real-time from a motion planner.

1.2. Stability of motions and the transverse linearization

Although the focus of this paper is control of non-periodic motions, many of the concepts and techniques used were first introduced for the study of periodic motions, and find their simplest expression in this context. In this section we will briefly introduce these concepts, and later extend them to the non-periodic case.

When studying stability of an autonomous nonlinear system $\dot{x} = f(x)$ about an equilibrium point x^* , the most basic method is to linearize, obtaining a linear comparison system $\dot{z} = Az$ with $A = df(x)/dx|_{x=x^*}$. Asymptotic stability of the linear system is necessary and sufficient for (local) exponential stability of x^* for the original nonlinear system.

When moving from stability of a point to stability of a motion, the situation is slightly more complicated. Supposing an autonomous differential equation $\dot{x} = f(x)$ has a non-trivial periodic solution $x(t) = x^*(t)$ with period T , one could linearize the system about this trajectory and obtain a time-varying linear comparison system $\dot{z}(t) = A(t)z(t)$, which would have periodic coefficients $A(t) = A(t + T)$. However, this system will never be asymptotically stable, since periodic solutions of an autonomous differential equation can never be asymptotically stable. The reasons for this are clear when one considers two solutions starting from different points on the periodic orbit: $x_1(0) = x^*(t_1)$, $x_2(0) = x^*(t_2)$, $t_1 \neq t_2$. After each period, they will each return to their starting points: $x_1(kT) = x^*(t_1) \neq x_2(kT) = x^*(t_2)$ for all k , whereas asymptotic stability of the motion would imply that $\|x_1(t) - x_2(t)\| \rightarrow 0$ as

$t \rightarrow \infty$. Note that for a controlled system $\dot{x} = f(x, u)$ one could apply, for example, a time-varying linear controller to stabilize to a periodic trajectory, however the closed-loop system is then time-varying and not autonomous. A discussion on this approach in the context of underactuated walking is given in Section 7.2.

A better concept by which to study stability of periodic systems is *orbital stability*. The basic idea is that the motion of the system should get closer to the *set* of states defining the target periodic solution, not a *particular* solution of the system. Let $\text{dist}(y, \mathcal{X})$ be the distance from a point y to its closest neighbor in a set \mathcal{X} , that is: $\text{dist}(y, \mathcal{X}) = \inf_{x \in \mathcal{X}} \|y - x\|$. Then we can make the following definitions:

Definition 1. Let \mathcal{X} be the set of states through which the periodic motion being studied passes, i.e. $\mathcal{X} = \{x : x = x^*(t), t \in [0, T]\}$. The periodic solution $x^*(t)$ is said to be asymptotically orbitally stable if there is a real number $b > 0$ such that for any initial condition x_0 with $\text{dist}(x_0, \mathcal{X}) < b$, the solution of the differential equation $\dot{x}(t) = f(x(t))$ starting from $x(0) = x_0$ has the property

$$\text{dist}(x(t), \mathcal{X}) \rightarrow 0 \text{ as } t \rightarrow \infty.$$

Furthermore, a periodic solution is said to be exponentially orbitally stable if there exists $c > 0$ and $a(x_0) > 0$ such that

$$\text{dist}(x(t), \mathcal{X}) < a(x_0) \exp(-ct), \quad t \geq 0.$$

The most well-known technique for proving orbital stability is the Poincaré first-return map, in which one constructs a “Poincaré section”: a hypersurface in state space (i.e. a surface of dimension one less than the state space) transversal to the orbit under study (e.g. $S(0)$ in Figure 2). By studying the behavior of the system only at times at which it passes through this surface, one obtains a discrete time system of dimension one less than the original system, referred to as the Poincaré map:

$$x_{\perp}(k+1) = \mathcal{P}[x_{\perp}(k)].$$

The Poincaré map has a fixed point at the periodic orbit: $\mathcal{P}[x_{\perp}^*] = x_{\perp}^*$. Stability or instability of this fixed point corresponds to orbital stability and orbital instability, respectively, of the periodic orbit. Exponential orbital stability corresponds to all the eigenvalues of the linearization of \mathcal{P} being inside the unit circle. This test has frequently been used to verify the stability of both passive and actuated walking robots (see e.g. Hurmuzlu and Moskowitz, 1986; Goswami et al., 1997; Grizzle et al., 2001; Westervelt et al., 2003; Wisse et al., 2005; Westervelt et al., 2007; Hobbelen and Wisse, 2007, and many others).

In the context of robot locomotion, one disadvantage of the Poincaré map is that it does not give a continuous representation of the dynamics of the system transverse to the target orbit, but focuses only at *one* point on the orbit. This means it has limited use for constructive control design.

When considering locomotion over uneven terrain, another significant shortcoming is that the method of




Fig. 2. A visualization of Poincaré surfaces and transverse linearization of a periodic orbit (red) and a trajectory converging to it (black).

Poincaré sections is only defined for periodic orbits. It can be used to study biped walking on flat ground or constant slopes, but on uneven ground where we have no reasonable expectation of periodic orbits it is not applicable.

Another tool for analysis of orbital stability is the *transverse linearization* (Urabe, 1967; Hale, 1980; Hauser and Chung, 1994). This tool has a number of advantages over Poincaré map analysis, particularly for control design (Shiriaev et al., 2008a; Freidovich et al., 2008b), but has so far seen relatively little use in robotics. In Song and Zefran (2006) and Shiriaev et al. (2008b) the transverse linearization was used to stabilize periodic walking motions, and in Manchester (2010); Manchester et al. (2010) it was used to estimate regions of attraction to periodic walking cycles.

Motivated by the structure of a mechanical system with n degrees of freedom, let us consider a system with a state space of dimension $2n$. The idea is to introduce a new coordinate system which moves along the target periodic motion. At each point on the target motion, one of the coordinates, ρ , is exactly aligned with the flow of the system, and the other $2n - 1$ coordinates, denoted x_{\perp} , are orthogonal to the flow of the system. This can be visualized via the related concept of a *moving Poincaré section* or *transverse dynamics*. This is a continuous family of $(2n - 1)$ -dimensional surfaces transverse to the desired trajectory, with one member of the family present at *every* point along the cycle ($S(t)$ for all t in Figure 2).

In a neighborhood of these new coordinates, the dynamics of the original system can then be written in the form:

$$\dot{\rho} = 1 + g_1(\rho, x_{\perp}), \quad (1)$$

$$\dot{x}_{\perp} = A(\rho)x_{\perp} + g_2(\rho, x_{\perp}). \quad (2)$$

for $\rho \in [0, T]$ representing the phase around the orbit, and $g_1(\cdot)$ and $g_2(\cdot)$ containing higher order terms (Urabe, 1967; Hale, 1980; Hauser and Chung, 1994). The transverse linearization of a periodic orbit is the linear part of the transversal dynamics:

$$\dot{z}(t) = A(t)z(t), \quad t \in [0, T].$$




Fig. 3. A visualization of Poincaré surfaces and transverse linearization of a non-periodic trajectory with impulses.

Exponential stability of this linear comparison system is equivalent to exponential orbital stability of the target periodic motion of the nonlinear system (Urabe, 1967; Hale, 1980; Hauser and Chung, 1994). This analysis can be extended to a class of hybrid nonlinear systems by linearizing the jumps and considering a jump linear system (Freidovich et al., 2008b). Note that the dynamics of ρ do not influence local orbital stability, and thus do not need to be explicitly computed.

When considering a controlled system, in contrast to the Poincaré map a transverse linearization provides a continuous representation of the relationship between controls and transverse coordinates of the form

$$\dot{z}(t) = A(t)z(t) + B(t)u(t), \quad t \in [0, T].$$

Stabilization of this system is equivalent to exponential orbital stabilization of the target motion of the nonlinear system, and can be approached using standard techniques in linear control theory.

Furthermore, the concept of a transverse linearization can be extended to non-periodic motions with impulses, such as walking over uneven terrain, as visualized in Figure 3. In this case, the transverse linearization will be of the form

$$\begin{aligned} \dot{z}(t) &= A(t)z(t) + B(t)u(t) \quad \text{for } t \neq t_j, \\ z(t_j^+) &= F_j z(t_j^-), \end{aligned}$$

where $t_j, j = 1, 2, 3, \dots$ are the times at which impacts occur in the target trajectory $q^*(t)$. Here, F_j is the linearization of the effect of the impact on the transverse state x_\perp .

In the remainder of this paper we will develop this idea, and use it to design a stabilizing control law. Note that unlike the case of an equilibrium or periodic orbit, stability of a non-stationary linearization about a trajectory of a nonlinear system does not necessarily imply stability of the trajectory (Leonov and Kuznetsov, 2007). However, under some assumptions which are reasonable for walking robot control we can prove stability using Lyapunov's direct method.

2. System model and motion plan

Although the technique we will propose can be applied to a wide range of nonlinear hybrid systems, in this paper we will concentrate on hybrid mechanical (Euler–Lagrange) systems. Let q be a vector of generalized coordinates, and u be a vector of forces and torques which can be assigned, then the dynamics of the system can be written in the following form (Hurmuzlu and Marghitu, 1994; Spong et al., 2006; Westervelt et al., 2007):

$$M(q) \ddot{q} + C(q, \dot{q}) \dot{q} + G(q) = B(q) u \quad \text{for } q \notin \mathcal{Q}$$

$$\left. \begin{aligned} q^+ &= \Delta_{q,i} q^- \\ \dot{q}^+ &= \Delta_{\dot{q},i}(q^-) \dot{q}^- \end{aligned} \right\} \text{ whenever } q^- \in \mathcal{Q}_k, k = 1, 2, 3, \dots \quad (3)$$

where $M(q)$ is the inertia matrix, $C(q, \dot{q})$ is the matrix of Coriolis and centrifugal terms, $G(q)$ is the gradient of the potential energy field, and $B(q)$ describes the effects of actuators on the generalized coordinates. The sets \mathcal{Q}_j represent a family of switching surfaces, e.g. for a walking robot, states at which the foot of the swing-leg hits the ground, and a new step begins. Throughout the paper, for discontinuous signals the superscripts $-$ and $+$ refer to the value of a signal immediately before and after a jump, respectively. We will also use the notation $q(t_j^-)$ and $q(t_j^+)$ for the value of $q(\cdot)$ immediately before or after t_j .

We assume that a motion-planning algorithm has been implemented which provides a trajectory of desired states:

$$q(t) = q^*(t) \in \mathbb{R}^n, \quad t \in [0, \infty).$$

The entire trajectory may be specified in advance, or it may be constructed over time in a receding horizon framework. The trajectory $q^*(t)$ is piecewise smooth, with occasional time instants $t_j, j = 1, 2, \dots$ at which an impact occurs, i.e. $q(t_j^-) \in \mathcal{Q}_k$ for some k . The following two assumptions will be satisfied for any reasonable motion plan for a walking robot.

Assumption 1. *There exists $\tau_1 > \tau_2 > 0$ such that $\tau_1 \geq t_{j+1} - t_j \geq \tau_2$ for all j .*

That is, the footsteps do not get infinitely long or infinitely short.

Assumption 2. *For all t_j , the vector $[\dot{q}^*(t_j^-) \dot{q}^*(t_j^+)]^T$ is linearly independent of the $2n-1$ vectors spanning the tangent plane of the switching surface at $q^*(t_j)$.*

That is, all impacts are “real” impacts, not grazing touches of the switching surface. This is a necessary assumption since with a grazing touch arbitrarily small deviations from the trajectory can result in no impact occurring, invalidating the analysis based on an impulsive linear system.

3. Receding-horizon control design

The main contribution of this paper is a proposal for a control design method that is guaranteed to stabilize the system to its motion plan.

The basic strategy is to repeatedly compute the transverse linearization about the target trajectory looking a few steps ahead, then compute a receding-horizon optimal control – also known as model predictive control – which is then transformed and applied to the nonlinear system. Note that exponential stabilization of time-varying systems, even linear systems, is a non-trivial problem. For time-invariant or periodic linear systems one can compute constant or periodic gain matrices, respectively, which exponentially stabilize the system. This is not true in general for time-varying systems, however the receding-horizon method is a computationally feasible solution to this problem (see e.g. Kwon et al. (1983); Mayne et al. (2000) and many others). Arguing stability of a trajectory of a nonlinear system via time-varying linearization is also far from straightforward (Leonov and Kuznetsov, 2007); however, the method by which we construct the controller allows us to construct a Lyapunov function by which this can be proven.

That is, find the feedback controller that minimizes the following cost function:

$$J(x, u) = \int_{t_i}^{t_{i+h}} [z(t)^T Q(t) z(t) + v(t)^T R(t) v(t)] dt + \sum_{j=i}^{i+h} z(t_j)^T Q_j z(t_j)$$

subject to the constraint $z(t_{i+h}) = 0$, where $z(t)$ is the state of the transverse linearization, and $v(t)$ is a virtual control signal.

Assumption 3. *There exists $\alpha_i > 0$, $i = 0, \dots, 5$, such that $\alpha_0 I \leq Q(t) \leq \alpha_1 I$, $\alpha_2 I \leq R(t) \leq \alpha_3 I$, and $\alpha_4 I \leq Q_j \leq \alpha_5 I$ for all t and j .*

We propose the following receding-horizon strategy, computed for time intervals beginning at each footstep and looking h footsteps ahead:

1. Set $i = 0$.
2. Receive from the motion-planner the desired trajectory $q^*(t)$ for the time interval $[t_i, t_{i+h})$. Compute the transverse linearization of the system (3) about the target trajectory:

$$\dot{z}(t) = A(t)z(t) + B(t)v(t) \text{ for } t \in [t_i, t_{i+h}), t \neq t_j, \quad (4)$$

$$z(t_j^+) = F_j z(t_j^-) \text{ for } j = i, i+1, \dots, i+h-1. \quad (5)$$

3. Solve the following jump-Riccati equation backwards in time from t_{i+h} to t_i with a final condition $Z(t_{i+h}) = 0_{(n-1)^2}$:

$$\begin{aligned} \frac{d}{dt} Z(t) &= Z(t)A(t)^T + A(t)Z(t) - B(t)R(t)^{-1}B(t)^T \\ &\quad + Z(t)Q(t)Z(t) \text{ for } t \neq t_j, \\ Z(t_j^-) &= \{F_j^T Z(t_j^+) - F_j + Q_j\}^{-1} \text{ for } j = i, i+1, \dots, i+h-1. \end{aligned} \quad (6)$$

4. Until the next footstep impact is detected, apply the following state-feedback controller:

$$\begin{aligned} u &= N(q, \dot{q})^{-1} [K(s)x_{\perp} - W(q, \dot{q})], \\ K(s) &= -R^{-1}(s)B_f(s)^T Z(s)^{-1}, \\ s &= \arg \min_{t \in [t_i, t_{i+1})} \|x - x^*(t)\| \end{aligned} \quad (7)$$

where $N(\cdot)$ and $W(\cdot)$ are partial feedback-linearizing terms, given in Section 4.1.

5. For the next footstep, set $i = i + 1$ and return to stage 2.

We note that the jump-Riccati equation (6) is different to that usually associated with the LQR design. In fact, $Z(t)$ is the *inverse* of the cost-to-go matrix. Therefore, setting $Z(t_{i+h}) = 0_{(n-1)^2}$ is equivalent to infinite cost on the final condition, i.e. a strictly-constrained final state. However, the time at which this state constraint is applied continues to recede into the future, and therefore it is never actually enforced in closed loop. This simple trick is useful in guaranteeing stability of the closed loop system, and can be relaxed somewhat if needed (Mayne et al., 2000).

In order to prove stability, we make some technical assumptions:

Assumption 4. *The hybrid transverse linearization system is uniformly completely controllable, as defined in Kalman (1960).*

This assumption essentially states that there is sufficient dynamical coupling between the unactuated and actuated links of the system. It is always satisfied with reasonable walking robot designs, and planned trajectories which are “close” to periodic, such as walking over uneven terrain.

Assumption 5. *There exist constants $\epsilon > 0$ and $\delta > 0$ such that, for a given state q, \dot{q} , $\|q - q^*(s)\| < \delta$ implies $\|\dot{x}_{\perp} - A(s)x_{\perp} + B(s)u(t)\| < \epsilon$ for $q \notin Q_j$ for any j , where $s = \arg \min_t \|q - q^*(t)\|$. Also, there exist $\bar{\epsilon} > 0$ and $\bar{\delta} > 0$ such that $\|q - q^*(t_j)\| < \bar{\delta}$ implies that $\|x_{\perp}(t_j^+) - F_j x_{\perp}(t_j)\| < \bar{\epsilon}$ when $q \in Q_j$ for some j .*

This is a uniform smoothness assumption on the dynamics around the trajectory. It essentially states that there are uniform (in time) bounds guaranteeing how well the transverse linearization locally approximates the system’s dynamics. It is a technical condition which will always be satisfied for reasonable walking motions.

We are now ready to state the main theoretical result of the paper:

Theorem 1. *If Assumptions 1, 2, 3, 4, and 5 are satisfied, then the controller (6), (7) locally exponentially orbitally stabilizes the planned motion of the hybrid nonlinear system (3).*

The proof is given in Appendix B.

4. A method for analytical construction of the transverse linearization

In this paper, we use a construction of the transverse linearization specifically designed for underactuated mechanical systems (Perram et al., 2003; Shiriaev et al., 2005, 2008a) which has several components in common with the technique of virtual holonomic constraints and hybrid zero dynamics (Grizzle et al., 2001; Westervelt et al., 2003, 2007). The transversal coordinates can be analytically constructed and have a very natural interpretation in terms of synchronization functions of the various joints, and a scalar kinetic-energy-like coordinate. Recently, an alternative construction of the transverse linearization for more general hybrid nonlinear systems has been introduced (Manchester, 2010), which allows the control method from the previous section to be applied to a very wide class of systems.

To begin our construction we represent the target trajectory in terms of virtual holonomic constraints. For each interval $[t_j, t_{j+1})$, $j = 0, 1, 2, \dots$, choose one generalized coordinate or some scalar function of the generalized coordinates $\theta := \Theta_j(q)$ which evolves monotonically along a desired trajectory.

Remark 1. *In the case of the compass-gait walker, which we will consider in Sections 5 and 6, we will take θ to be the “ankle angle”. It is a reasonable assumption that for any useful walking motion, this angle evolves monotonically over any given step. This representation is common in walking robot control (Westervelt et al., 2007).*

Since it evolves monotonically θ can then be considered as a reparameterization of time, and hence the nominal trajectories of all other coordinates over each interval $[t_j, t_{j+1})$ can be given as well-defined functions of θ :

$$\begin{aligned} q_1^*(t) &= \phi_1^j(\theta(t)), \\ &\vdots \\ q_n^*(t) &= \phi_n^j(\theta(t)) \quad \forall t \in [t_j, t_{j+1}). \end{aligned}$$

Having thus defined the functions $\phi_1^j, \dots, \phi_n^j$, one can define variables representing deviations from the nominal trajectory:

$$\begin{aligned} y_1(t) &:= q_1(t) - \phi_1^j(\theta(t)), \\ &\vdots \\ y_n(t) &:= q_n(t) - \phi_n^j(\theta(t)) \quad \forall t \in [t_j, t_{j+1}), \end{aligned} \quad (8)$$

where $y_m(t) = 0$ for all m implies the system is on the nominal trajectory, i.e. the motion of the joints is synchronized.

Consider now the quantities θ, y_1, \dots, y_n . Since the system is n -dimensional, these $n + 1$ quantities are excessive coordinates for the system, and hence one can be dropped. The coordinate which is dropped must be locally independent of θ . Without loss of generality, let us assume we drop y_n , and our new coordinates are $y = [y_1, \dots, y_{n-1}]^T$ and θ .

Remark 2. *When the conditions $y_m = 0$ for all m are enforced via feedback action, the functions $\phi_1^j, \dots, \phi_n^j$ are often referred to as virtual holonomic constraints (Grizzle et al., 2001; Westervelt et al., 2003, 2007). Our control strategy does not require that these constraints be strictly enforced to guarantee stability, they are simply used as a set of coordinates. However, we retain the terminology “virtual constraints”.*

4.1. Construction of the continuous part of the transverse linearization

The representation of trajectories introduced above allows us to construct, at any θ , a set of transversal coordinates. The first $2n - 2$ coordinates are given by the virtual constraints defined in Equation (8) and their derivatives:

$$y_i = q - \phi_i(\theta), \quad (9)$$

$$\dot{y}_i = \dot{q}_i - \frac{d\phi_i^j(\theta)}{d\theta} \dot{\theta}, \quad i = 1, \dots, n-1. \quad (10)$$

Now, the dynamics of y can be written in the form

$$\ddot{y} = W(y, \theta, \dot{y}, \dot{\theta}) + N(y, \theta)u,$$

where the functions $W(\cdot)$ and $N(\cdot)$ are simply the nonlinear dynamics in the variables y, \dot{y} , and follow from substitution of (9) and (10) into the equations of motion (3) (see, for example, Shiriaev et al. (2008a) for details). For systems of underactuation degree one, the dynamics of y can be feedback linearized via a controller of the form:

$$u = N(y, \theta)^{-1} [v - W(y, \theta, \dot{y}, \dot{\theta})], \quad (11)$$

where v is a virtual control signal, resulting in the dynamics $\ddot{y} = v$. Note that this is just a *partial* feedback linearization: since there are $n - 1$ actuators, we can feedback linearize $2n - 2$ states out of a total of $2n$. The partial feedback linearization is not strictly necessary, but it makes the system behave closer to a linear system and can thus improve the performance of a controller based on a linear approximation.

For a full set of transverse coordinates, one more independent coordinate is required. The final coordinate should be a function of the state space which is independent of y and \dot{y} , zero on the target orbit, and it should be possible to compute both its value and its dynamics analytically. It turns out that for mechanical systems, such a function exists in a general form.

For the continuous phase of an underactuated mechanical system, if the relations $y = \dot{y} = 0$ are maintained then the dynamics of the coordinate θ take the following form:

$$\alpha(\theta)\ddot{\theta} + \beta(\theta)\dot{\theta}^2 + \gamma(\theta) = 0, \quad (12)$$

where $\alpha(\cdot), \beta(\cdot), \gamma(\cdot)$ can be computed from the equations of motion of the system – see Appendix C, Equation (20). This special form is due to the fact that the matrix

of Coriolis and centrifugal terms $C(\theta, \dot{\theta})$ is linear in $\dot{\theta}$ for mechanical systems.

An important fact is that a partial closed-form solution of the system (12) can be computed (Perram et al., 2003):

$$\dot{\theta}^2 = \psi(\theta, \theta_0) \dot{\theta}_0^2 + \Gamma(\theta, \theta_0), \quad (13)$$

where $(\theta_0, \dot{\theta}_0)$ is any point on the desired trajectory of the reduced system (12) and the formulas for $\psi(\cdot)$ and $\Gamma(\cdot)$ are given in Appendix C, Equations (21) and (22). By introducing a new variable

$$I = \dot{\theta}^2 - \psi(\theta, \theta_0) \dot{\theta}_0^2 - \Gamma(\theta, \theta_0),$$

we have a clear candidate for the final transverse coordinate: it is clearly independent of y and \dot{y} and is zero when the system is on the target motion. Furthermore, since I has been constructed analytically, the dynamics of I can be explicitly computed (see Appendix C, Equations (23–26)), so we have a complete set of transverse coordinates:

$$x_{\perp} := \begin{bmatrix} I \\ y \\ \dot{y} \end{bmatrix},$$

and explicit formulas for the transverse linearization along the continuous interval $[t_j, t_{j+1})$ can be computed:

$$\dot{z}(t) = A(t)z(t) + B(t)v(t) \quad \text{for } t \in [t_j, t_{j+1}), \quad (14)$$

where $A(t)$ and $B(t)$ are

$$A(t) = \begin{bmatrix} a_{11}(t) & a_{12}(t) & a_{13}(t) \\ 0_{(n-1)} & 0_{(n-1)^2} & I_{(n-1)} \\ 0_{(n-1)} & 0_{(n-1)^2} & 0_{(n-1)^2} \end{bmatrix}, \quad B(t) = \begin{bmatrix} b_1(t) \\ 0_{(n-1)^2} \\ I_{(n-1)} \end{bmatrix}, \quad (15)$$

and with z representing the state of the linear comparison system. The coefficients of this matrix are given in Appendix C, Equation (27).

Note that in this development we have used the fact that the system has a single unactuated degree of freedom to simplify the development by the introduction of a partial feedback linearizing controller. This is not necessary, however, and higher degrees of underactuation can be handled with a small modification (see Freidovich and Shiriaev (2009)).

4.2. Transverse linearization of impacts

Certain care is required in linearizing the impact map. Analysis via the jump-linear system assumes that the jumps occur at a certain time, i.e. t_i for some i . However, small perturbations of the nonlinear system from the target trajectory will not necessarily all hit the switching surface at the same time. Geometrically, one can say that the transversal surfaces are always orthogonal (in the full state space (q, \dot{q})) to the current motion of the system, but the switching surfaces will usually not be.

To solve this problem, we need two projection matrices which go from the transversal surface to the switching surface before an impact, and from the switching surface to the transversal surface after an impact. Suppose $d\Delta_j$ is the linearization of the impact map at time t_j about the nominal trajectory, then

$$z(t^+) = F_j z(t) \quad \text{for } t = t_j, j = 1, 2, \dots, \quad (16)$$

where $F_j = P_j^+ d\Delta_j P_j^-$.

The matrix P_j^+ projects from the transversal surface $TS(t)$ onto the switching surface at the end of the continuous phase, and the matrix P_j^- projects the state back on to the transversal surface after switching. The construction of these projections is given in Appendix C, Equation (28).

This approach was first presented in Shiriaev et al. (2008b). Alternative solutions include searching for a global orthogonalizing transform (Song and Zefran, 2006) and warping the transversal surfaces so as to line up with switching surfaces (Manchester, 2010), however these cannot be applied with the analytical construction of the transverse dynamics which we use in this paper.

5. Experimental setup

We have constructed a two-degree-of-freedom planar biped robot with point feet, a photograph and schematic of which are shown in Figure 4. The robot is mounted on a boom arm with a counterweight, and thus walks in a circular path. The dynamical effect of the boom is approximated by having different values of hip mass for inertial (m_H) and gravitational (m_{Hg}) terms in the model. The robot is fitted with retractable feet to avoid toe-scuffing, however we do not consider the feet as additional degrees of freedom since their masses are negligible.

The robot is modeled in the form of a hybrid mechanical system (3), the parameters of which were estimated via nonlinear least-squares fit on simulation error. The full equations for the model are given in Appendix D. Good fitting required the addition of a friction model for the hip joint consisting of Coulomb and viscous parts:

$$\tau_F = F_C \text{sign}(\dot{q}_1) + F_V \dot{q}_1.$$

The parameters of the model are given in Table 1.

The robot is fitted with optical encoders measuring the angle between the legs and the absolute angle of the inner leg. From these measurements q_1 and q_2 can be calculated. The control law relies on velocities as well, and these are estimated with an observer. The observer structure chosen is one which has previously been used successfully in walking robot control, consisting of a copy of the nonlinear dynamics and a linear correction term (Grizzle et al., 2007). Let \hat{q} and $\hat{\dot{q}}$ be the estimates of the configuration states and velocities. Then the observer is given by:




Fig. 4. Schematic and photograph of compass-gait biped used in the experiments.

Table 1. Parameters of the compass-gait biped.

Parameters	Values
Masses [kg]	$m = 1.3, m_H = 2.2,$ $m_{Hg} = -1.2$
Inertia [kg m ²]	$I_c = 0.0168$
Lengths [m]	$l = 0.32, l_c = 0.0596$
Gravitational constant [m/s ²]	$g = 9.81$
Ratio current/input [A]	$k_I = 1.1$
Motor torque constant [Nm/A]	$k_\tau = 0.0671$
Coulomb friction [Nm]	$F_C = 0.02$
Viscous friction [Nm s]	$F_V = 0.01$

$$\frac{d}{dt} \begin{bmatrix} \hat{q} \\ \hat{\dot{q}} \end{bmatrix} = \begin{bmatrix} M(\hat{q})^{-1} (-C(\hat{q}, \hat{\dot{q}}) \hat{\dot{q}} - G(\hat{q}) + B(\hat{q})u) \\ +L(y - \hat{q}), \end{bmatrix}$$

$$\hat{q}^+ = \Delta_q \hat{q}, \quad \hat{\dot{q}}^+ = \Delta_{\dot{q}}(\hat{q}) \hat{\dot{q}},$$

where y is the measurement of q . A simple choice for the observer gain is $L = [1/\epsilon \quad 2/\epsilon^2]$, which places the eigenvalues of the linearized error system at $-1/\epsilon$. In our experiments we found that $\epsilon = 0.02$ gave a reasonable compromise between speed of convergence and noise rejection.

5.1. Polynomial representation of desired motion

In this paper, we represent planned trajectories in the form of *virtual holonomic constraints*: one coordinate (or function of coordinates) which evolves monotonically is chosen as a “phase” variable, and the desired motions of all other coordinates are represented as functions of this variable.

For the compass biped we take $\theta = q_2$, the “ankle” angle of the stance leg relative to horizontal. Then to specify the path through configuration space for each step j , we need to specify only the inter-leg angle q_1 as a function of the ankle angle: $q_1^* = \phi^j(\theta)$. We chose to construct the ϕ^j functions as Bézier polynomials, which can represent a wide range of useful motions with quite a low number of parameters. For details, see Westervelt et al. (2007, Ch. 6), in which Bézier polynomials were used to design periodic trajectories.

A trajectory $q_1^*(\theta)$ can be represented as

$$\phi^j(\sigma) := \sum_{k=0}^M a_k \frac{M!}{k!(M-k)!} \sigma_j^k (1 - \sigma_j)^{M-k},$$

where $\sigma_j \in [0, 1]$ is a mapping of the evolution of θ for the j -th footstep into the interval $[0, 1]$:

$$\sigma_j = \frac{\theta - \theta(t_j^+)}{\theta(t_{j+1}^-) - \theta(t_j^+)}.$$

This method is straightforward to adapt to non-periodic trajectories. For the experiments in this paper, we hand-chose parameters of fourth-order Bézier polynomials which achieved a walking motion over our terrain. Real-time planning in this framework may offer significant efficiency benefits, and will be part of our future work.

6. Experimental results

To test the controller experimentally, a relatively simple task was chosen: the robot should walk flat for two steps, then down two “stairs”, and then continue along the flat.

The steps had a height of 2 cm. A video of a successful experiment can be seen in Extension 1.

The control design was implemented as in Section 3, with constant weighting matrices $Q(t) = Q_j = I_3$ and $R(t) = 1$ for all t and j . The look-ahead horizon was chosen as three footsteps ahead.

For each step, the solution of the jump-Riccati equation took approximately half a second to compute using the `ode45` solver in MATLAB running on a Pentium III desktop computer. This is roughly the time it takes for the robot to complete a step, and it is reasonable to expect that optimized compiled code could perform this task much more quickly. Hence, one can say that the control law could be feasibly computed in real-time, as a part of a dynamic motion-planning and control system.

Figure 5 depicts the results of one experiment. Figure 5(A) is a cartoon of the biped's motion generated from real data, showing the state every 0.3 s, with the current stance leg always indicated in red.

In Figure 5(B) the evolution of the “ankle angle” q_2 is plotted against time for one experiment. During the continuous phases, q_2 serves as our reparameterization of time θ . We note here that, particularly on the second and fourth steps, there is some jitter in the curve. In Figure 5(C) the inter-leg angle q_1 is plotted against time in blue, along with the “nominal” value of q_1 plotted in red. Note that since the nominal value of q_1 is not a function of time but a function of q_2 , defined by the virtual constraint, the jitters in the q_2 measurement lead to jitters in the nominal value of q_1 . Nevertheless, tracking is quite good, and sufficient for the robot to maintain a stable walking trajectory. Figures 5(D) and (E) depict the joint velocities \dot{q}_1 and \dot{q}_2 , obtained from the same observer used in the control system, along with their nominal values as functions of the current value of q_2 . Again, the jitter in q_2 leads to large noise in the velocity estimates. Despite this, good tracking is maintained through all the planned steps. Repeated experiments were performed with similar results each time, indicating good robustness of the control strategy.

7. Discussion

7.1. Further experimental verification

The experiments reported in this paper are a proof of concept, however the terrain is arguably not particularly demanding, with the height variations only around six percent of the leg-length. The reason for this was mainly practical: the robot used has a very small toe extension and larger height variations are physically impossible for the toe clear. Since these experiments were completed the same receding-horizon transverse LQR technique has been applied for stabilization of a quadruped robot bounding over rough terrain (Shkolnik et al., 2010). The trajectories included bounding up stairs and over logs, and the height variations were of the same order as the quadruped's leg-length. Owing to

the complexity of the quadruped model, a different construction of the transverse linearization was used which is valid for any hybrid nonlinear system, not just those with a Lagrangian structure (Manchester, 2010).

Other researchers have successfully performed experiments stabilizing periodic trajectories of the pendubot and Furuta pendulum using the transverse linearization (Shiri-aev et al., 2007; Freidovich et al., 2008a). In these works, the same analytical construction of the transverse linearization was used and a periodic LQR problem was solved to compute a feedback controller.

7.2. Comparison to control via full-state LQR

When it comes to implementation on a real system, an advantage of the method of transverse linearization, however it is constructed, is that it naturally leads to a *static state-feedback* control law of the form $u(t) = k(x(t))$ where $x(t) = [q(t)^T \dot{q}(t)^T]^T$. That is, for each point in the state space, there is a unique control action to be taken. A time-varying linear system is used as an intermediate step in the construction, but by composing this with a projection operator, the complete control law is time-invariant.

In contrast, if one were to design a controller for the full-order dynamics using, for example, time-varying LQR, one would need to stabilize the system to a *particular time-solution* of the nominal trajectory, resulting in a time-varying control law of the form $u(t) = k(x(t), x^*(t), t)$, where $x^*(t)$ is the target solution. This imposes problems of synchronization: suppose the natural underactuated dynamics of the real robot are somewhat slower or faster than the model assumed for motion planning. This will always be the case to a greater or lesser extent since accurately modeling friction effects and energy losses due to impacts remains a difficult problem. If the motion of the real robot is slower than the model, then its trajectory will lag behind the planned trajectory. In the authors' experience, for underactuated systems large deviations in control and the actuated links may be commanded in an attempt to “catch up”, which can force the system to leave the region in which the linear controller is stabilizing, and result in system failure. The transverse-linearization-based controller generally leads to much larger regions of attraction, although much depends on particular properties of the system and the choice of weighting matrices and prediction horizon, so it is difficult to make general statements to this effect. However, in the authors' experience, solving the Riccati equations or controllability grammians for the full-order linearized dynamics of a compass-gait walker, including the impacts, is poorly conditioned numerically, which can indicate very weak controllability of the real system.

One could attempt to solve the synchronization problem by using a projection operator: finding the closest point on the planned trajectory to the current state, and applying the full-order controller computed for that point, thus creating

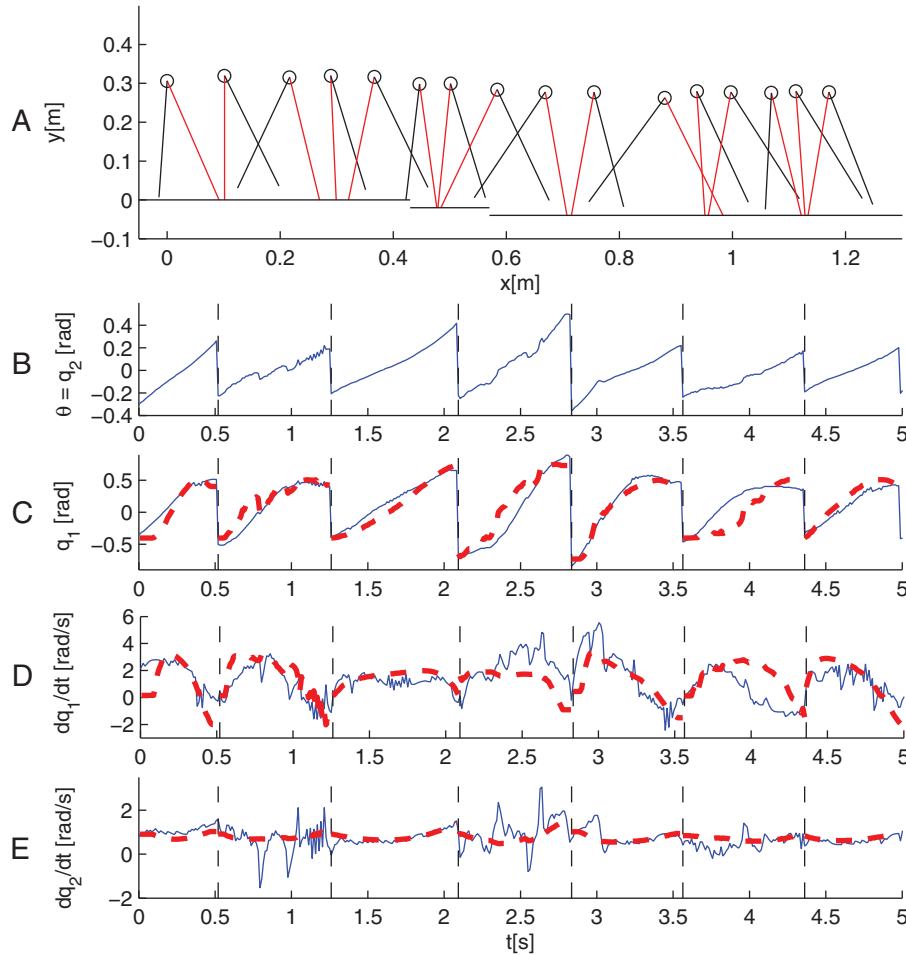


Fig. 5. Results from a successful experiment walking on uneven terrain. See Section 6 for discussion.

a state-feedback control law. However, for a full-order LQR controller such a projection operator is not guaranteed to preserve stability of the closed-loop system. In contrast, the combination of a transverse-linearization controller and a projection operator maintains its stabilizing property.

7.3. Comparison to control via hybrid-zero-dynamics

The construction of the transverse linearization given in this paper and in Shiriaev et al. (2008a,b) has several technical elements in common with the hybrid-zero-dynamics (HZD) framework for biped control (Grizzle et al., 2001; Westervelt et al., 2003, 2007), and it is worth discussing the similarities and differences between these approaches. The basic difference can be expressed like so: in the HZD framework, a high-gain controller is constructed to drive the $2n-2$ coordinates y and \dot{y} to zero in finite time. The satisfaction of the constraint $y = \dot{y} = 0$ results in a two-dimensional reduced system – satisfying the differential equation (12) – the stability of which can be checked via Poincaré map analysis. Note that it was subsequently proven that orbital stability is

still implied if y and \dot{y} are driven *sufficiently close* to zero (Morris and Grizzle, 2009). In contrast, in our approach, $2n-1$ transverse coordinates – including y and \dot{y} but also the extra coordinate I – are driven towards zero exponentially via a receding horizon control design.

Figure 6 depicts the phase portrait of the reduced dynamics (12) for a particular periodic trajectory of the compass-gait walker. That is, it depicts the dynamics of θ and $\dot{\theta}$ on the manifold defined by the constraint $y = \dot{y} = 0$. We notice that in the center of the phase portrait is a hyperbolic fixed point. Initial conditions with sufficient velocity will pass above this fixed point and complete the walking step, but those with lower velocity fall backwards again, which we consider a failure. Planned motions of walking robots which have been optimized for minimum energy expenditure tend to come quite close to this fixed point, such as the trajectory indicated by the black line.

Since all points on this phase portrait satisfy the virtual constraints, a controller based on the HZD framework will not attempt to stabilize these trajectories to the target orbit. For most reasonable walking robots it can be shown that the dynamics of this reduced system are locally stable, although




Fig. 6. Phase portrait of the reduced dynamics of the compass-gait model. The black line is the target trajectory, the thin vertical red lines indicate the switching surfaces. The dashed blue lines are trajectories using an HZD controller, the dashed red lines are trajectories using a transverse linearization controller.

it is clear that the region of attraction can be very small. In particular, the above-mentioned difficulty in accurately modeling impacts plays an essential role here. If the energy lost during an impact is greater than expected, resulting in θ falling below its value on the target trajectory, the robot will fall backwards, as indicated by the lower dashed blue line in Figure 6. However, a robot controlled using a transverse linearization starting from the same point will be stabilized to the target cycle (the dashed red line). Equally, if the initial velocity after an impact is much greater than expected – perhaps due to underestimation of the ground slope – then the HZD-based walker will complete the walking cycle, as indicated by the upper dashed blue line, but meet the right-hand switching surface with much higher velocity than planned, which may cause the foot to bounce or slip on impact. In contrast, the transverse linearization controlled robot starting from the same velocity will stabilize the system back to the target trajectory during the swing phase, again indicated by the dashed red line.

It is interesting to note that when starting from a lower-than-planned velocity, the robot initially slows down even further before speeding up (the red dashed line dips below the blue dashed line, before crossing it). This is because in order to increase the speed of rotation of the stance leg, which is unactuated, it must push the swing-leg further forward than planned (see Figure 7) to shift the center of mass of the robot forward, thus increasing the forward acceleration. But doing so has a short-term effect of slowing down motion of the stance leg due to coupling between angular momenta. The opposite occurs with the trajectory which starts from a higher-than-planned stance-leg velocity: the swing leg is initially swung right back to shift the center of mass further back, thus slowing rotation of the stance leg. Note that although this strategy of shifting the swing




Fig. 7. Swing leg trajectory as a function of stance leg for the compass-gait walker controlled by transverse linearization. Depicted is the planned motion, i.e. the virtual constraint (black), and the motions to stabilize a reduced initial $d\theta/dt$ (red solid) and an increased initial $d\theta/dt$ (red dashed). In all cases a HZD-controlled walker would follow the black trajectory.

leg forward or backward to affect momentum has an intuitive basis, it is a natural outcome of the general stabilizing controller and is not due to programming of a heuristic strategy.

The main advantage of the HZD framework over the proposed method is that it is much simpler to implement, which can be a great benefit in practice, and typically works very well for planar bipeds. Indeed, once a set of virtual constraints is chosen a Proportional-Derivative (PD) controller is often sufficient for each actuator. However, this advantage is diminished somewhat if the robot has higher degrees of underactuation (e.g. a three-dimensional compass-gait) for which the hybrid zero dynamics may not be stable (Chevallereau et al., 2009), and also for systems with compliance, for which the hybrid zero dynamics manifold may not be invariant through impacts (Morris and Grizzle, 2009). In these cases the problem of finding virtual constraints that result in a stable walking motion is non-trivial. Furthermore, if the planned trajectory is non-periodic, then proof of stability via Poincaré map analysis is not possible.

The main benefits of the proposed method are that it is guaranteed to stabilize any walking trajectory which is stabilizable, under very weak conditions. The extra complexity required for this guarantee is a receding horizon control loop that computes online the transverse linearization – which is given in analytic form – and solutions of a Riccati differential equation. However, in the authors experience this is well within the computational capabilities of most modern robots.

8. Conclusions

In this paper, we have proposed a constructive control technique to stabilize non-periodic motions of underactuated

robots, with a focus on the application of walking over uneven terrain. The controller is computed by constructing a lower-dimensional system of coordinates *transverse* to the target cycle and then computing a receding-horizon feedback controller to exponentially stabilize the linearized dynamics of the transverse states. The proposed approach is supported by theoretical results proving orbital exponential stability of the target trajectory of the original nonlinear system.

The method used in this paper to construct the transverse linearization has several elements in common with the HZD framework. For walking robots with underactuation degree one this is a very natural construction, since the states are synchronization errors of the joints and a scalar kinetic-energy-like state. Other constructions may be more natural for other classes of system (see, for example, Urabe, 1967; Hale, 1980; Shkolnik et al., 2010; Manchester, 2010)).

To the best of the authors' knowledge, this is the first systematic control method which can provably stabilize general non-periodic motions of an underactuated walker. The proposed technique was experimentally verified using a compass-gait biped walker. It was seen that, despite measurement errors and inevitable uncertainties in modeling, the controller reliably stabilized the target motions.

In this paper we solve the receding-horizon control problem via a jump-Riccati equation. If there are hard constraints on states or actuation, a simple extension is to discretize the transverse linearization and solve a quadratic program at each time step if computational resources allow (see Mayne et al., 2000, and references therein).

In closing, we note that closely related methods can be used to compute rigorously verified regions of stability for walking robot motions (Manchester, 2010; Manchester et al., 2010) which can then be used to construct high performance nonlinear control systems (Tedrake et al., 2010). A focus of future work will be experimental evaluation of the control approach in these papers combined with rigorous verification for dynamic walking over very rough terrain.

Acknowledgements

The authors would like to thank Anton Shiriaev, Leonid Freidovich, and Jessy Grizzle for helpful discussions and comments.

Funding

This work was supported by the National Science Foundation (grant number 0746194).

Conflict of interest

The authors declare that they have no conflicts of interest.

References

- Asano F, Yamakita M, Kamamichi N and Luo ZW (2004) A novel gait generation for biped walking robots based on mechanical energy constraint. *IEEE Transactions on Robotics and Automation* 20: 565–573.
- Bainov DD and Simeonov PS (1989) *Systems with Impulse Effects: Stability, Theory and Applications*. Chichester: Ellis Horwood.
- Byl K, Shkolnik A, Prentice S, Roy N and Tedrake R (2008) Reliable dynamic motions for a stiff quadruped. In *Proceedings of the 11th International Symposium on Experimental Robotics (ISER)*, pp 319–328, Athens, Greece.
- Byl K and Tedrake R (2009) Metastable Walking Machines. *The International Journal of Robotics Research*, 28: 1040–1064.
- Chevallereau C, Grizzle JW and Shih C-L (2009) Asymptotically stable walking of a five-link underactuated 3-d bipedal robot. *IEEE Transactions on Robotics* 25: 37–50.
- Collins S, Ruina A, Tedrake R and Wisse M (2005) Efficient bipedal robots based on passive-dynamic walkers. *Science* 307: 1082–1085.
- Freidovich L, Robertsson A, Shiriaev A and Johansson R (2008) Periodic motions of the Pendubot via virtual holonomic constraints: theory and experiment. *Automatica* 44: 785–791.
- Freidovich LB and Shiriaev AS (2009) Transverse linearization for mechanical systems with passive links, impulse effects, and friction forces. In *Proceedings of the Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference*, pp 6490–6495, Shanghai, China.
- Freidovich LB, Shiriaev AS and Manchester IR (2008) Stability analysis and control design for an underactuated walking robot via computation of a transverse linearization. In: *Proceedings of the 17th IFAC World Congress*, pp 10166–10171, Seoul, Korea.
- Goswami A, Espiau B and Keramane A (1997) Limit cycles in a passive compass gait biped and passivity-mimicking control laws. *Autonomous Robots* 4: 273–286.
- Grizzle JW, Abba G and Plestan F (2001) Asymptotically stable walking for biped robots: Analysis via systems with impulse effects. *IEEE Transactions on Automatic Control* 46: 51–64.
- Grizzle JW, Choi JH, Hammouri H and Morris B (2007) On observer-based feedback stabilization of periodic orbits in bipedal locomotion. In *Proceedings Methods and Models in Automation and Robotics*.
- Hale JK (1980) *Ordinary Differential Equations*. Malabar, FL: Krieger.
- Hauser J and Chung CC (1994) Converse Lyapunov function for exponential stable periodic orbits. *Systems and Control Letters* 23: 27–34.
- Hobbelen DGE and Wisse M (2007) A disturbance rejection measure for limit cycle walkers: The gait sensitivity norm. *IEEE Transactions on Robotics* 23: 1213–1224.
- Hurmuzlu Y and Marghitu DB (1994) Rigid body collisions of planar kinematic chains with multiple contact points. *The International Journal of Robotics Research* 13: 82–92.
- Hurmuzlu Y and Moskowitz G (1986) The role of impact in the stability of bipedal locomotion. *Dynamics and Stability of Systems* 1: 217–234.
- Iida F and Tedrake R (2009) Minimalistic control of a compass gait robot in rough terrain. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp 6490–6495, Kobe, Japan.

- Kalman RE (1960) Contributions to the theory of optimal control. *Bol. Soc. Mat. Mexicana* 5: 102–119.
- Kolter JZ, Rodgers MP and Ng AY (2008) A control architecture for quadruped locomotion over rough terrain. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 811–818.
- Kwon WH, Bruckstein AM, and Kailath T (1983) Stabilizing state-feedback design via the moving horizon method. *International Journal of Control* 37: 631–643.
- Leonov GA and Kuznetsov NV (2007) Time-varying linearization and the Perron effects. *Regular and Chaotic Dynamics* 17: 1079–1107.
- Manchester IR (2010) Transverse dynamics and regions of stability for nonlinear hybrid limit cycles. *Preprint* arXiv:1010.2241.
- Manchester IR, Tobenkin MM, Levashov M and Tedrake R (2010) Regions of attraction for hybrid limit cycles of walking robots. *Preprint* arXiv:1010.2247.
- McGeer T (1990) Passive dynamic walking. *International Journal of Robotics Research* 9: 62–82.
- Mayne DQ, Rawlings JB, Rao CV and Scokaert POM (2000) Constrained model predictive control: Stability and optimality. *Automatica* 36: 789–814.
- Morris B and Grizzle JW (2009) Hybrid invariant manifolds in systems with impulse effects with application to periodic locomotion in bipedal robots. *IEEE Transactions on Automatic Control* 54: 1751–1764.
- Perram J, Shiriaev A, Canudas-de-Wit C and Grognaud F (2003) Explicit formula for a general integral of motion for a class of mechanical systems subject to holonomic constraint. In *Proceedings of the 2nd IFAC Workshop Control Methods for Lagrangian and Hamiltonian Systems*, pp 87–92, Seville, Spain.
- Rebula J, Neuhaus P, Bonnlander B, Johnson M and Pratt J (2007) A controller for the littledog quadruped walking on rough terrain. *Proceedings of the 2007 IEEE International Conference on Robotics and Automation (ICRA)*, pp 1467–1473, Rome, Italy.
- Saranli U, Buehler M and Koditschek DE (2001) RHex: a simple and highly mobile hexapod robot. *The International Journal of Robotics Research* 20: 616–631.
- Shiriaev A, Freidovich L, Robertsson A, Johansson R, and Sandberg A (2007) Virtual-constraints-based design of stable oscillations of Furuta pendulum: theory and experiments. *IEEE Transactions on Robotics* 23: 827–832.
- Shiriaev AS, Freidovich LB and Manchester IR (2008) Can we make a robot ballerina perform a pirouette? Orbital stabilization of periodic motions of underactuated mechanical systems. *Annual Reviews in Control* 32: 200–211.
- Shiriaev AS, Freidovich LB and Manchester IR (2008) Periodic motion planning and analytical computation of transverse linearizations for hybrid mechanical systems. In *Proceedings of the 47th IEEE Conference on Decision and Control*, pp 4326–4331, Cancun, Mexico.
- Shiriaev AS, Perram JW and Canudas-de-Wit C (2005) Constructive tool for orbital stabilization of underactuated nonlinear systems: virtual constraints approach. *IEEE Transactions on Automatic Control* 50: 1164–1176.
- Shkolnik A, Levashov M, Manchester IR and Tedrake R (2010) Bounding on rough terrain with the LittleDog robot. *International Journal of Robotics Research*, to appear.
- Song G and Zefran M (2006) Stabilization of hybrid periodic orbits with application to bipedal walking. In *Proceedings of the American Control Conference*, pp 2504–2509.
- Spong MW and Bhatia G (2003) Further results on control of the compass gait biped. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003 (IROS 2003)*, Vol. 2, pp. 1933–1938.
- Spong MW and Bullo F (2005) Controlled symmetries and passive walking. *IEEE Transactions on Automatic Control* 50: 1025–1031.
- Spong MW, Holm JK and Lee D (2007) Passivity-based control of bipedal locomotion. *IEEE Robotics and Automation Magazine* 14: 30–40.
- Spong MW, Hutchinson S and Vidyasagar M (2006) *Robot Modeling and Control*. Hoboken, NJ: John Wiley and Sons.
- Tedrake R, Manchester IR, Tobenkin MM and Roberts JW (2010) LQR-Trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research* 29: 1038–1052.
- Urabe M (1967) *Nonlinear autonomous oscillations*. New York: Academic Press.
- Vukobratovic M and Borovac B (2004) Zero-moment point – thirty five years of its life. *International Journal of Humanoid Robotics* 1: 157–173.
- Westervelt ER, Grizzle JW, Chevallereau C, Choi JH and Morris B (2007) *Feedback Control of Dynamic Bipedal Robot Locomotion*. Boca Raton, FL: CRC Press.
- Westervelt ER, Grizzle JW and Koditschek D (2003) Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control* 48: 42–56.
- Wisse M, Schwab AL, van der Linde RQ and van der Helm FCT (2005) How to keep from falling forward: elementary swing leg action for passive dynamic walkers. *IEEE Transactions on Robotics* 21: 393–401.
- Yang T, Westervelt ER, Serrani A and Schmiedeler JP (2009) A framework for the control of stable aperiodic walking in underactuated planar bipeds. *Autonomous Robots* 27: 277–290.

Appendix A: Index to Multimedia Extensions

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

Table of Multimedia Extensions

Extension	Type	Description
1	Video	An experimental compass-gait biped walking down steps.

Appendix B: Proof of local stability

Here we provide a sketch of the proof of Theorem 1, since after some preparation the details follow from standard results in linear systems theory and orbital stabilization. We consider the Lyapunov function candidate for the linear comparison system

$$\dot{\mathcal{V}}(z(t), t, \mathcal{K}_i) = z(t)^T P(t) z(t),$$

where $P(t) = Z(t)^{-1}$ and $Z(t)$ is the solution of a finite-time jump-Riccati equation (6) computed at footstep i . That is, $\mathcal{V}(z(t), t, \mathcal{K}_i)$ would be the total “cost-to-go” from a state $z(t)$ assuming that the finite-horizon feedback strategy computed at footstep i were continued until the end of its horizon, rather than just for one footstep.

We state the following three facts about this Lyapunov function candidate:

1. It follows from Assumptions 1, 4, and 3, and standard arguments from optimal control (Kalman, 1960), that there exist $\beta_1 > \beta_0 > 0$ such that

$$\beta_0 I \leq P(t) \leq \beta_1 I. \quad (17)$$

2. Throughout the continuous phase from t_i to t_{i+1} ,

$$\frac{d}{dt} \mathcal{V}(z(t), t, \mathcal{K}_i) = -z(t)^T Q(t) z(t) - u(t)^T R(t) u(t).$$

Therefore, it follows from the bounds on $Q(t)$ in Assumption 3 that

$$\frac{d}{dt} \mathcal{V}(z(t), t, \mathcal{K}_i) \leq -\alpha_0 \|z(t)\|^2 \quad (18)$$

for all t .

3. Let \mathcal{K}_i refer to the strategy of using the finite-time controller calculated at the beginning of step i . Under this strategy, $x(t_{i+h}) = 0$ and remains zero for all $t > t_{i+h}$. After step i , the state is $z(t_{i+1})$. A feasible strategy from here would be to continue with control strategy \mathcal{K}_i . However, a new optimization is performed at step $i+1$ over a new horizon $i+1+h$. Since continuing with \mathcal{K}_i is a feasible strategy, it follows from Assumptions 4 and 3 that there exists a constant $0 < \beta_3 < 1$ such that the new optimal strategy \mathcal{K}_{i+1} must have a cost-to-go which satisfies

$$\mathcal{V}(z(t_{i+1}), t_{i+1}, \mathcal{K}_{i+1}) \leq \beta_3 \mathcal{V}(z(t_{i+1}), t_{i+1}, \mathcal{K}_i); \quad (19)$$

i.e. the Lyapunov function decreases by at least a factor of β_3 at every impulse.

To prove that this implies orbital stability of the target trajectory of the original nonlinear system, we consider the Lyapunov function $\mathcal{V}(x_\perp, \mathcal{K}_i)$, and observe that Assumption 5 together with (18) and (19) imply that for sufficiently small x_\perp there will exist constants $\alpha_6 > 0$ and $0 < \beta_4 < 1$ such that

$$\begin{aligned} \frac{d}{dt} \mathcal{V}(x_\perp(t), t, \mathcal{K}_j) &\leq -\alpha_6 \|x_\perp(t)\|^2, t \neq t_j, j = 1, 2, 3, \dots \\ \mathcal{V}(x_\perp(t_{j+1}), t_{j+1}, \mathcal{K}_{j+1}) &\leq \beta_4 \mathcal{V}(x_\perp(t_{j+1}), t_{j+1}, \mathcal{K}_j), \\ t = t_j, j = 1, 2, 3, \dots \end{aligned}$$

Now, from this and (17) it follows that x_\perp exponentially converges to zero for the original nonlinear system, i.e. the target trajectory is exponentially orbitally stable, via a generalization of Lyapunov's second method to systems with impulse effects (Bainov and Simeonov, 1989, Ch. 13).

Appendix C: Miscellaneous formulas

Functions for the reduced dynamics of a virtually constrained system (12):

$$\begin{aligned} \alpha(\theta) &= B^\perp(q) M(\Phi(\theta)) \Phi'(\theta), \\ \beta(\theta) &= B^\perp(q) [C(\Phi(\theta), \Phi'(\theta)) \Phi'(\theta) \\ &\quad + M(\Phi(\theta)) \Phi''(\theta)], \\ \gamma(\theta) &= B^\perp(q) G(\Phi(\theta)), \\ \Phi(\theta) &= [\phi_1(\theta) \dots \phi_n(\theta)]^T, \\ \Phi'(\theta) &= \frac{d}{d\theta} \Phi(\theta), \quad \Phi''(\theta) = \frac{d^2}{d\theta^2} \Phi(\theta), \end{aligned} \quad (20)$$

and $B^\perp(q)$ is a full rank matrix such that $B^\perp(q) B(q) = 0$.

Functions for the solution of the dynamics θ^2 in Equation (13):

$$\psi(\theta, \theta_0) = \exp \left\{ -2 \int_{\theta_0}^{\theta} \frac{\beta(\tau)}{\alpha(\tau)} d\tau \right\}, \quad (21)$$

and

$$\Gamma(\theta, \theta_0) = 2 \int_{\theta_0}^{\theta} \psi(\theta, s) \frac{\gamma(s)}{\alpha(s)} ds. \quad (22)$$

Dynamics of I away from the target orbit defined by $y = \dot{y} = 0$:

$$\frac{d}{dt} I = \frac{2\dot{\theta}}{\alpha(\theta)} (w - \beta(\theta) I), \quad (23)$$

where

$$w := \alpha(\theta) \ddot{\theta} + \beta(\theta) \dot{\theta}^2 + \gamma(\theta).$$

Note that when $y = \dot{y} = 0$, the dynamics of $\theta, \dot{\theta}$ are given by (12) and $w = 0$.

The dynamics of $\theta, \dot{\theta}$ for the partial feedback-linearized system can be locally approximated as:

$$\begin{aligned} \alpha(\theta) \ddot{\theta} + \beta(\theta) \dot{\theta}^2 + \gamma(\theta) &= g_y(\cdot) y \\ &\quad + g_{\dot{y}}(\cdot) \dot{y} + g_v(\cdot) v, \quad (24) \\ \ddot{y} &= v, \quad (25) \end{aligned}$$

where

$$\begin{aligned} g_y &= g_y(\theta, \dot{\theta}, \ddot{\theta}, y, \dot{y}), \quad g_{\dot{y}} = g_{\dot{y}}(\theta, \dot{\theta}, \ddot{\theta}, y, \dot{y}), \\ g_v &= g_v(\theta, \dot{\theta}, y, \dot{y}) \end{aligned}$$

are smooth functions which are straightforward to compute with computer algebra packages from the difference between (12) and the dynamics for $\theta, \dot{\theta}$ with (11) substituted for u (Shiriaev et al., 2005, 2008a). The dynamics of I are then approximated by substituting (24) into (23):

$$\frac{d}{dt} I = \frac{2\dot{\theta}}{\alpha(\theta)} (g_y(\cdot) y + g_{\dot{y}}(\cdot) \dot{y} + g_v(\cdot) v - \beta(\theta) I). \quad (26)$$

The matrix coefficients for the continuous phase of the transverse linearization:

$$\begin{aligned} b_1(t) &= \dot{\theta}_*(t) \frac{2g_v(\theta_*(t), \dot{\theta}_*(t), 0, 0)}{\alpha(\theta_*(t))}, \\ a_{12}(t) &= \dot{\theta}_*(t) \frac{2g_y(\theta_*(t), \dot{\theta}_*(t), \ddot{\theta}_*(t), 0, 0)}{\alpha(\theta_*(t))}, \\ a_{11}(t) &= -\dot{\theta}_*(t) \frac{2\beta(\theta_*(t))}{\alpha(\theta_*(t))}, \\ a_{13}(t) &= \dot{\theta}_*(t) \frac{2g_y(\theta_*(t), \dot{\theta}_*(t), \ddot{\theta}_*(t), 0, 0)}{\alpha(\theta_*(t))}. \end{aligned} \quad (27)$$

Projection operators for the impact map:

$$\begin{aligned} P_j^- &:= \left(I_4 - \frac{n(t_j^-) m_-^T}{n^T(t_j^-) m_-} \right) \left[L_c(t_j^-) \right]^{-1} \begin{bmatrix} I_3 \\ 0_{1 \times 3} \end{bmatrix}, \\ P_j^+ &:= L_c(t_j^+) \left(I_4 - \frac{n(t_j^+) n^T(t_j^+)}{n^T(t_j^+) n(t_j^+)} \right) \end{aligned} \quad (28)$$

where $n(t) = [\dot{q}_*(t)^T \ddot{q}_*(t)^T]^T$, m_- is a vector normal to the linearization of the switching surface, and L_c is the Jacobian matrix defining the mapping

$$\begin{bmatrix} \Delta I & \Delta y & \Delta \dot{y} \end{bmatrix}^T = L_c(t) \begin{bmatrix} \Delta q(t) & \Delta \dot{q}(t) \end{bmatrix}^T,$$

where Δ represents a small deviation.

Appendix D: Compass biped model

The model of the experimental setup is given by (3), where

$$\begin{aligned} M(q) &= \begin{bmatrix} p_1 & -p_1 + \cos(q_1)p_2 \\ -p_1 + \cos(q_1)p_2 & p_3 + 2p_1 - 2\cos(q_1)p_2 \end{bmatrix}, \\ C(q, \dot{q}) &= \begin{bmatrix} 0 & -\dot{q}_2 \sin(q_1)p_2 \\ -\sin(q_1)(\dot{q}_1 - \dot{q}_2)p_2 & \sin(q_1)\dot{q}_1 p_2 \end{bmatrix}, \\ B &= \begin{bmatrix} k_I k_\tau \\ 0 \end{bmatrix}, \end{aligned}$$

$$G(q) = [\sin(-q_2 + q_1)p_4, -\sin(-q_2 + q_1)p_4, -\sin(q_2)p_5 - \sin(q_2)p_4]^T.$$

The coefficients are defined by the physical parameters of the robot:

$$\begin{aligned} p_1 &= l_c^2 m + (l - l_c), \quad p_2 = m l l_c, \quad p_3 = m_H l^2 + 2m l (l - l_c) \\ p_4 &= m g l_c, \quad p_5 = g(m_H l + 2m(l - l_c)). \end{aligned}$$

The impact model in (3) is derived under the assumption of having an instantaneous and inelastic collision of the swing leg with the ground and no occurrence of slip or rebound (Hurmuzlu and Marghitu, 1994):

$$\Delta_q = \begin{bmatrix} -1 & 0 \\ -1 & 1 \end{bmatrix}, \quad \Delta_{\dot{q}}(q^-) = \Delta_q [H^+(q^-)]^{-1} H^-(q^-),$$

where

$$\begin{aligned} H_{1,1}^+(q^-) &= p_1 + p_3, \quad H_{1,2}^+(q^-) = H_{2,1}^+(q^-) \\ &= p_2 \cos(q_1^-) - p_1 - p_3, \\ H_{2,2}^+(q^-) &= -2p_2 \cos(q_1^-) + p_3 + 2p_1, \\ H_{1,1}^-(q^-) &= p_1 - p_2, \quad H_{1,2}^-(q^-) = p_3 \cos(q_1) - p_1 + p_2, \\ H_{2,1}^-(q^-) &= -p_1 + p_2, \quad H_{2,2}^-(q^-) = p_3 \cos(q_1). \end{aligned}$$