Chapter 4.7

Hybrid Zero Dynamics Control of Legged Robots

Aaron D. Ames** and Ioannis Poulakakis**[††]

**Mechanical and Civil Engineering and Control and Dynamical Systems, California Institute of*
Technology (Caltech), United States [††]*Department of Mechanical Engineering, University of*
Delaware, Newark, DE, United States

4.7.1 BIPEDAL ROBOTS WITH HZD CONTROLLERS

The method of hybrid zero dynamics (HZD) offers a paradigm for designing feedback control laws that induce reliable dynamically-stable walking and running motions on bipedal robots, all the while providing analytically tractable guarantees of performance. The method has been introduced by Jessy Grizzle, Eric Westervelt, and their collaborators in Grizzle et al. (2001), Westervelt et al. (2003), Westervelt (2003), Morris and Grizzle (2009), where geometric nonlinear control tools are developed to generate provably stable limit-cycle walking motions in a class of bipedal robots by dealing directly with their underactuated and hybrid nature; the book (Westervelt et al., 2007) provides an integrative perspective. At its core, the method relies on restricting the dynamics of the robot on a lower-dimensional attractive and invariant subset of its state space. This is achieved by defining a set of holonomic output functions with the control objective being to drive these outputs to zero. Through this process, a lower-dimensional dynamical system emerges from the closed-loop dynamics of the robot that governs the existence and stability properties of its behavior.

Beyond its theoretical value, the method has been successful in experimentally generating robust walking motions on the planar bipedal robot Rabbit (shown in Fig. 4.7.1); see Chevallereau et al. (2003), Westervelt et al. (2004) for details regarding these experiments. Rabbit's successful walking experiments prompted the extension of the HZD method to stabilize bipedal running (Chevallereau et al., 2005). However, while initial experiments have been successful in exciting running on Rabbit, the resulting motions could not be sustained due to actuator limitations (Morris et al., 2006). Given that elastic energy storage elements – e.g., in the form of tendons in animals (Alexander, 1988) or springs (Raibert, 1986) in robots – play a significant role in the realization of running motions, subsequent research efforts on the HZD method concentrated on its implementation on compliant robots. Based on theoretical tools developed in Morris (2008), Morris and Grizzle (2009), the notion of compliant hybrid zero dynamics has been introduced in Poulakakis (2008),
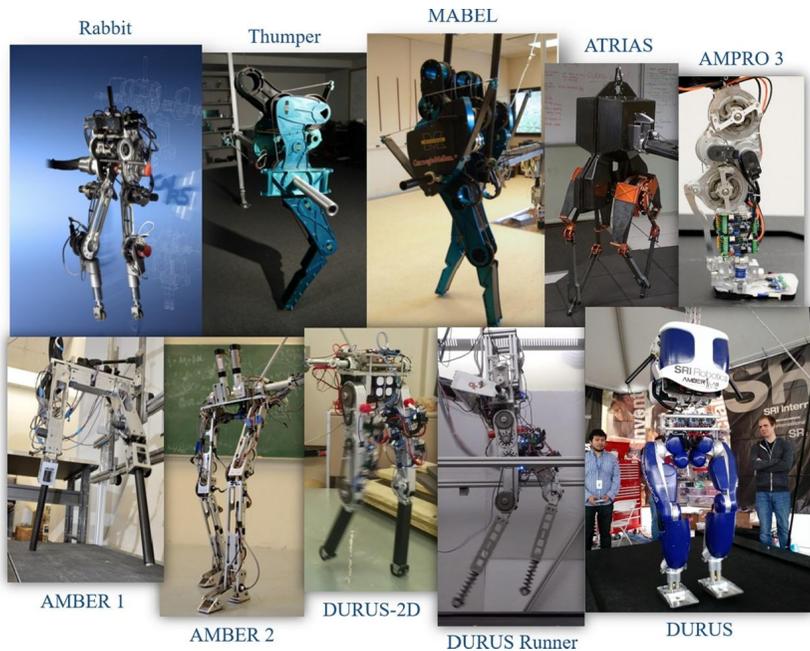
FIGURE 4.7.1 A collection of robots and robotic assistive devices for which HZD-based methods have successfully resulted in stable bipedal locomotion.

Poulakakis and Grizzle (2009b) and further refined in Sreenath (2011) to induce robust walking (Sreenath et al., 2011) and running (Sreenath et al., 2013) motions in experiments with the compliant bipedal robot MABEL[4] shown in Fig. 4.7.1. In addition to these experiments, MABEL provided an excellent platform for validating advanced locomotion controllers for accommodating unexpected large ground-height variations (Park et al., 2013), and for testing alternative Lyapunov-based HZD control schemes as in Ames et al. (2014a) that afford greater flexibility in incorporating constraints such as actuator torque saturation (Galloway et al., 2015). These ideas have been translated to other robots, including the use of Lyapunov-based HZD techniques to realize walking on the planar robot DURUS-2D (Cousineau and Ames, 2015) (shown in Fig. 4.7.1), along with running on the same platform (Ma et al., 2017).

Building upon the successes of HZD, which focused on underactuated walking due to its clear separability from methods that require full actuation (e.g., zero moment point (ZMP) based frameworks (Kajita et al., 2003, 2006;

4. MABEL and its monopedal version Thumper depicted in Fig. 4.7.1 have been designed and constructed by Professor J. Hurst in a collaborative effort between The University of Michigan and Carnegie Mellon University; see Grizzle et al. (2009) for an overview and Hurst et al. (2007), Hurst and Rizzi (2008), Hurst (2008) for details relevant to the underlying design philosophy.

Vukobratović et al., 2006; Goswami, 1999a, 1999b)), the HZD methodology has been proven to be extensible to walking behaviors that more clearly resemble humans. When locomoting, humans naturally display different discrete phases that correspond to their changes in contact with the world (Ackermann, 2007). In the context of robots, this multi-contact behavior can be represented by a hybrid dynamical system model of walking consisting of discrete domains, wherein dynamics of the robot change discretely as a function of its contact with the ground; this results in phases of full, under- and over-actuation (Ames et al., 2011; Vasudevan et al., 2013). The ability of HZD to handle underactuation motivated its extension to this multi-domain locomotion scenario. In particular, using human locomotion data as inspiration, the framework of human-inspired control extended HZD to the full-actuated case through the notion of *partial* hybrid zero dynamics (PHZD) (Ames, 2014). Combining this approach with the HZD methods for underactuated walking, e.g., on AMBER 1, resulted in the ability to consider multi-domain HZD, thereby achieving multi-contact walking behaviors on bipedal robots, including ATRIAS (Hereid et al., 2014), AMBER 2 (Zhao et al., 2014b, 2015a), and the prosthesis AMPRO (Zhao et al., 2016b) (see Fig. 4.7.1). These results, both formally and experimentally, indicated that HZD can provide a mathematical framework for realizing human-like walking behaviors on robotic systems.

The last frontier for HZD-based methods was their extension to three-dimensional (3D) walking and realization on humanoid robots; the challenges and approaches are described in detail in Grizzle et al. (2014, 2010). This will enable us to engage dynamically moving bipeds in motion planning tasks, including navigation in environments cluttered by obstacles as in Gregg et al. (2012), Motahar et al. (2016), Veer et al. (2017). While 3D robot walking utilizing HZD had long proved feasible in simulation, and had even proven realizable on small-scale humanoids like the NAO robot (Ames et al., 2012a; Powell et al., 2013), bridging the gap between this theory and the experimental realization on full-scale humanoid robots is a difficult task. This is, at its core, a function of the fact that HZD uses the entire dynamics of the robot to generate walking gaits in the context of a constrained nonlinear programming problem. When this optimization problem can be solved, it results in dynamic and efficient gaits – yet as the complexity of the robot increases, solving the problem becomes more difficult making the translation to hardware evermore challenging. During the design and development of the humanoid robot DURUS (shown in Fig. 4.7.1), structure in nonlinear optimization problem necessary to generate gaits was discovered and exploited to allow for rapid gait generation; bringing the time needed to obtain a stable walking gait from hours to a few minutes (Hereid et al., 2016); importantly, due to the presence of springs in the ankles of DURUS, this was done in the context of multi-domain walking that exploits

PHZD to achieve stability. Not only was stable walking achieved, but it was done so in a sustained fashion with the end result being the public demonstration of DURUS at the DARPA Robotics Challenge, wherein it walked continuously for over 2 1/2 hours covering over 2 km – all on a single 1.1 kWh battery (Reher et al., 2016a). Additionally, these results were extended to the case of multi-contact walking with natural heel–toe behaviors thereby demonstrating human-like humanoid locomotion (Reher et al., 2016b). Moreover, in both cases, the walking realized on DURUS was the most efficient walking realized on a bipedal humanoid robot (Collins and Ruina, 2005). The final lessons from the realization of HZD on a variety of platforms was that it provides a powerful method for realizing dynamic walking behaviors on bipedal and humanoid robots.

4.7.2 MODELING LEGGED ROBOTS AS HYBRID DYNAMICAL SYSTEMS

Walking and running behaviors can be modeled as distinguished periodic orbits of mechanical systems that are strongly nonlinear and hybrid in nature. For example, a simplified walking cycle consists of successive phases of single support (swing phase) and double support (impact phase). On the other hand, running comprises phases where a leg is in contact with the ground (stance phase) and phases where the system is in the air following a ballistic motion under the influence of gravity (flight phase). This combination of continuous dynamics and discrete transitions among them is characteristic of legged locomotion and it gives rise to hybrid system models, which are the focus of this section.

4.7.2.1 Continuous Dynamics

Let Q be the configuration space of a robot with n degrees of freedom, i.e., $n = \dim(Q)$, with coordinates $q \in Q$; examples of coordinate choices for various bipeds are shown in Fig. 4.7.2. For the sake of definiteness, it may be necessary to choose Q to be a subset of the actual configuration space of the robot so that global coordinates can be defined,[5] i.e., such that Q is embeddable in $\mathbb{R}^n$, or more simply $Q \subset \mathbb{R}^n$. Consider the equations of motion for a robot given in the general form by the Euler–Lagrange equations (Murray et al., 1994; Spong et al., 2006):

$$D(q)\ddot{q} + C(q, \dot{q})\dot{q} + G(q) = Bu \qquad (4.7.1)$$

where $D(q)$ is the mass matrix, and $C(q, \dot{q})\dot{q}$, $G(q)$ are vectors containing the centrifugal and Coriolis forces and the gravitational forces, respectively, and

---

5. At various points of this chapter we will assume that certain matrix functions have full rank; it may be necessary to carefully choose Q to satisfy these conditions.
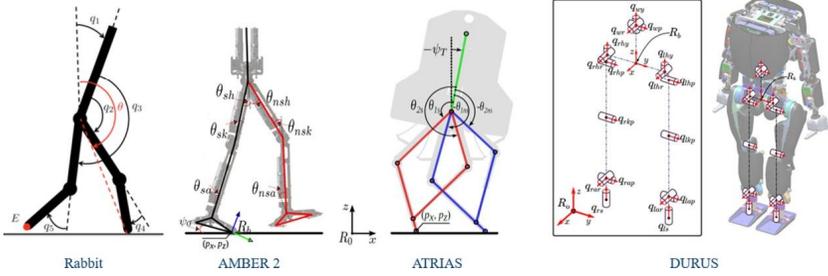
FIGURE 4.7.2 Examples of the configuration space for a collection of bipedal and humanoid robots; in this case, Rabbit (Westervelt et al., 2003), AMBER 2 (Zhao et al., 2014b), ATRIAS (Hereid et al., 2014) and DURUS (Hereid et al., 2016).

$B \in \mathbb{R}^{n \times m}$ is the actuation matrix which determines the way in which the torque inputs, $u \in U \subset \mathbb{R}^m$, actuate the system (where here U is the set of admissible control inputs). Importantly, the actuation matrix changes based upon the actuation type of the robot: in the case of *full actuation*, this matrix is full rank, while in the case of *underactuation* this matrix has rank $m < n$ indicating that it is not possible to actuate all of the degrees of freedom of the system, and if $m > n$ the system is *overactuated*, i.e., there is more control authority than degrees of freedom in the system.

Selecting the state vector x to include the configuration variables and the corresponding rates, that is $x = [q^T \; \dot{q}^T]^T \in \mathbb{R}^{2n}$, and noticing that

$$\frac{d}{dt}\begin{bmatrix} q \\ \dot{q} \end{bmatrix} = \begin{bmatrix} \dot{q} \\ -D(q)^{-1}\left(C(q,\dot{q})\dot{q} + G(q)\right) \end{bmatrix} + \begin{bmatrix} 0 \\ D(q)^{-1}B \end{bmatrix} u \qquad (4.7.2)$$

results in the following state-space form of the continuous dynamics (4.7.1):

$$\dot{x} = f(x) + g(x)u. \qquad (4.7.3)$$

4.7.2.2 Discrete Dynamics

In the basic walking model – more advanced models are discussed in Section 4.7.2.4 below – the continuous dynamics (4.7.3) represents the swing phase, which evolves until the leg hits the ground, thereby resulting in an impact. It is this impact that is the basis for the hybrid dynamical system model that underlies walking and running motions (Westervelt et al., 2007; Grizzle et al., 2014; Ames et al., 2011; Haddad et al., 2006). In particular, we consider the height of the swing foot and the surface defined by this height being zero. If $p_{\text{toe}}^{\text{v}}$ denotes the height of the toe of the swing leg, then the surface

$$S = \{x \in TQ \mid p_{\text{toe}}^{\text{v}}(q) = 0 \text{ and } \dot{p}_{\text{toe}}^{\text{v}}(x) < 0\} \qquad (4.7.4)$$
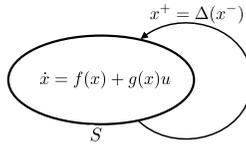
FIGURE 4.7.3 Representation of the hybrid model of the biped as a system with "impulse" effects.

is the *switching surface* (or the *guard* in the terminology of hybrid systems). Upon reaching the switching surface, the system undergoes an impact, which in basic walking models represents an instantaneous double support phase. The end result is an update law

$$x^+ = \Delta(x^-) \qquad (4.7.5)$$

mapping the pre-impact states, x^-, to the post-impact states, x^+. This impact model includes both a "change of coordinates" in the configuration variables corresponding to swapping the swing and stance legs, together with a discrete change in the velocity of the system determined by a plastic impact of the swing foot with the ground (which causes the stance foot to leave the ground and thus become the swing foot). More details on deriving the impact map Δ can be found in the book (Westervelt et al., 2007).

4.7.2.3 Hybrid Control System

The end result of these constructions is a *system with impulsive effects* or a *hybrid control system*:

$$\mathscr{HC} : \begin{cases} \dot{x} = f(x) + g(x)u, & x \in D \backslash S, \\ x^+ = \Delta\left(x^-\right), & x^- \in S, \end{cases} \qquad (4.7.6)$$

where $D = \{x \in TQ \mid p_{\text{toe}}^{\text{v}}(q) > 0\}$ is the *domain* of the system, i.e., we require the swing foot to be above the ground. Note that the domain is often restricted to the admissible domain through the inclusion of friction constraints. The system (4.7.6) is depicted in Fig. 4.7.3. We also note that sometimes systems of this form are written as a tuple

$$\mathscr{HC} := (D, U, S, \Delta, (f, g)), \qquad (4.7.7)$$

as is more common in the hybrid systems literature (Lygeros et al., 2003; Goebel et al., 2009; van der Schaft and Schumacher, 2000; Lamperski and Ames, 2013).

4.7.2.4 Advanced Models of Locomotion

The hybrid system model of a walking robot that we have considered so far includes only a single continuous and discrete domain. For robots with more complex mechanical characteristics, i.e., springs and nontrivial feet, hybrid system models with more complex discrete structure are needed. Interestingly enough, these hybrid models naturally relate to human locomotion models, i.e., humans tend to display a specific discrete domain structure when walking, wherein their contact points with the ground change throughout the gait (Ames et al., 2011; Vasudevan et al., 2011, 2013). Other examples of locomotion models evolving on multiple domains can be found in running motions, due to the alternation between stance and flight phases (Poulakakis and Grizzle, 2009a; Sreenath et al., 2013). These more complex models for locomotion can again be modeled as a hybrid system.

The key element to advanced models of robotic walking and running is an oriented graph, $\Gamma = (V, E)$, that indicates how the contact points change throughout the course of a gait, i.e., the vertices of this graph (V) correspond to different collections of contact points with the ground, and the transitions (described by edges E) occur when these contact points change; see Fig. 4.7.4 for examples in the case of multi-contact locomotion and Fig. 4.7.8 for a running model. The end result is a hybrid control system model of the form

$$\mathscr{HC} := (\Gamma, D, U, S, \Delta, FG) \tag{4.7.8}$$

where, in this case, $D = \{D_v\}_{v \in V}$ is a collection of domains, $U = \{U_v\}_{v \in V}$ is a collection of admissible inputs, $S = \{S_e\}_{e \in E}$ is a set of switching surfaces, $\Delta = \{\Delta_e\}_{e \in E}$ is a set of impact maps with $\Delta_e : S_e \subset D_{\text{source}(e)} \to D_{\text{target}(e)}$ and $FG = \{(f_v, g_v)\}_{v \in V}$ is a collection of control systems of the form (4.7.3). It is important to note that the degree of actuation changes for each domain, i.e., on some domains the system may be underactuated, on some it might be fully actuated, and on others it can be overactuated. The specific methods for constructing hybrid system models as they relate to the changing contact points of the robot can be found in Ames et al. (2011).

To provide a concrete example, consider the multicontact model of the bipedal robot AMBER 2 shown in Fig. 4.7.4. As described in Zhao et al. (2014b, 2015a), this model consists of three domains, D_{v^+}, D_{v^i}, and D_{v^-}, that depend on how the robot's contact points (heel and toe) change throughout the course of a step. The dynamics on each of these domains changes with the change in contact points. Importantly, each of these domains display a different actuation type: D_{v^+} is overactuated, D_{v^i} is fully actuated, and D_{v^-} is underactuated. Other examples of multidomain walking, and the
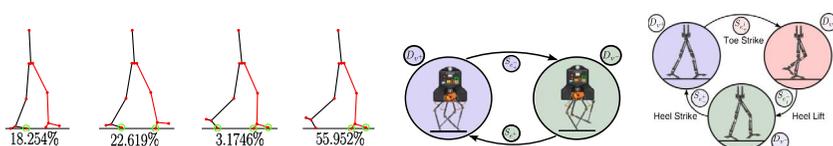
FIGURE 4.7.4 (Left) The discrete domains, and the percentage of the step spent in a domain, for a human walking (as obtained from experimental data), (Middle) the discrete two domain structure associate with the hybrid system model of ATRIAS, and (Right) the discrete three-domain structure for AMBER 2 when walking with articulated feet (Zhao et al., 2014b).

application of hybrid zero dynamics to these systems to achieve robotic running, will be discussed in Section 4.7.6.2. Finally, an interesting class of multidomain hybrid models emerges naturally in the context of planning the motion of dynamically walking bipeds amidst obstacles (Gregg et al., 2012; Motahar et al., 2016). This can be achieved through the sequential composition of primitive limit-cycle walking motions each stabilized through HZD as in Motahar et al. (2016), Veer et al. (2017).

4.7.3 VIRTUAL CONSTRAINTS FOR LOCOMOTION

Central to the HZD approach is the introduction of *virtual* constraints. These constraints represent relations among the robot's degrees of freedom that correspond to preferred postures during the realization of a walking or running gait. They are formulated as functions of the configuration variables of the form $h(q) = 0$, $q \in Q$, and can thus be interpreted as *holonomic* constraints, the enforcement of which effectively restricts the robot's motion on low-dimensional surfaces embedded in its higher-dimensional state space. It should be emphasized however, that the key difference with the classical notion of holonomic constraints from analytical mechanics (Goldstein et al., 2002) is that virtual holonomic constraints are imposed on the system via its actuators, not via workless constraint forces.

4.7.3.1 Virtual Constraints

In our setting, we would like to "force" a set of coordinates – those over which we have control – to follow desired patterns. Doing so both enforces certain patterns with regard to walking motions and reduces the overall dimensionality of the system to a reduced dimensional space, thus giving rise to a lower-dimensional dynamical system, namely the *zero dynamics*. Mathematically, we consider the difference between an actual output, y_a, and a desired output, y_d, expressed via

$$y(q) := y_a(q) - y_d(\tau(q), \alpha) \in \mathbb{R}^m, \tag{4.7.9}$$

where the desired function begins as a function of time, $y_d(t, \alpha)$, dependent on a parameter set α, and converted to a function of the configuration variables through a *parameterization of time* often chosen to be of the general form,

$$\tau(q) = \frac{\theta(q) - \theta^+}{\theta^- - \theta^+}, \tag{4.7.10}$$

where $\theta : Q \to \mathbb{R}$ is a phase variable, $\theta^+ = \theta(q^+)$ is its value post-impact, $\theta^- = \theta(q^-)$ is its value pre-impact, and therefore $\tau : Q \to [0, 1]$ throughout the course of a step. To provide some intuition, as Fig. 4.7.2 indicates, the phase variable θ can be chosen to correspond to the angle of the line connecting the hip with the toe of the support leg, which is a monotonically increasing quantity that captures "progression" of the support leg into the step; see Westervelt et al. (2007) for details. It is important to emphasize that the outputs (4.7.9) depend only on the configuration variables, hence the term virtual holonomic constraints.

4.7.3.2 Designing Virtual Constraints for Locomotion Tasks

Based upon the framework of virtual constraints, the main idea is to consider a vector of output variables y_a in (4.7.9), with one output for each actuator. These outputs capture quantities that are of interest, e.g., angles in the system or other geometric relationships, such as the position of the center of mass or the height of the swing foot. The goal is to drive these outputs to evolve according to a collection of desired behaviors as represented by y_d, which is a function of the phase variable τ and a set of parameters α that allow "tuning" the constraints according to desired specifications. The art of gait design is to pick y_d so that it displays certain properties so that driving $y \to 0$ in (4.7.9) guarantees stability of the system. A concrete way of selecting y_d in (4.7.9) is through the use of Beziér polynomials of degree M, i.e., for $i = 1, \dots, m$,

$$y_d(\tau(q), \alpha)_i = \sum_{k=0}^{M} \frac{M!}{k!(M-k)!} \alpha_{k,i} \tau(q)^k (1 - \tau(q))^{M-k}. \tag{4.7.11}$$

The use of Beziér polynomials is only one choice of functions for the design of the desired evolution y_d in (4.7.9), which offers some flexibility in imposing desired boundary conditions on the different phases that compose a cyclic locomotion pattern. More details about certain key properties of these polynomials and on how to use them in the context of gait design can be found in Westervelt et al. (2007).

Virtual constraints, designed via Beziér polynomials, provides a computationally efficient way of constructing virtual constraints. Yet, since the desired

behavior is given by polynomials, they do not necessarily capture the virtual constraints present in human walking. That is, we could instead seek *human-inspired* virtual constraints whose design is inspired by human locomotion data. In this light, human data suggests that for certain collections of outputs, y_a, appear to act like the time solution to a mass–spring–damper system, i.e., humans appear to display simple "linear" behavior when the proper collection of virtual constraints are considered (Sinnet et al., 2011a; Ames, 2012; Powell et al., 2012; Huihua et al., 2012; Zhao et al., 2014b). This motivates the following *mass–spring–damper* desired output:

$$y_d(\tau(q), \alpha)_i = y_{\text{MSD}}(\tau(q), \alpha)_i$$
$$:= e^{-\alpha_{4,i}\tau(q)}(\alpha_{k,i}\cos(\alpha_{k,i}\tau(q)) + \alpha_{k,i}\sin(\alpha_{k,i}\tau(q))) + \alpha_{k,i}$$
$$(4.7.12)$$

for $i = 1, \ldots, m$, which is simply the time solution to a linear mass–spring–damper system, i.e., a second order linear system. Human data has been calculated from a variety of actual output combinations, y_a, and it has been shown that y_{MSD} accurately describes (with high correlation) these outputs; examples include the position of the hip, the position of the center of mass, and the knee angles (Ames, 2014; Sinnet et al., 2014).

Another class of virtual constraints developed for *fully actuated* walking robots considers both velocity modulating and position modulating virtual constraints. More specifically, in this case we can modulate both the position of the robot – through the virtual constraints defined in (4.7.9) – and its traveling speed. Moreover, we would like to do this in a general fashion that will allow for different collections of virtual constraints depending on the bipedal robot being considered and the desired behavior to be achieved. To regulate the velocity of the robot in an explicit fashion, we consider the following virtual constraints (Ames, 2014):

$$y_1(q, \dot{q}) = \frac{\partial \theta(q)}{\partial q} \dot{q} - v, \tag{4.7.13}$$

$$y_2(q) = y_{2,a}(q) - y_{2,d}(\tau(q), \alpha) \tag{4.7.14}$$

where y_2 are the position modulating outputs as defined in (4.7.9), $\theta : Q \to \mathbb{R}$ is the phase variable of the virtual constraint (4.7.9), and v is the desired velocity. For example, we may wish to explicitly control the forward velocity of the center of mass to regulate the robot's speed; in this case θ would be the position of the center of mass, and v would be the desired velocity. In doing so, it is often useful to consider the following modified form for the parameterizations of time:

$$\tau(q) = \frac{\theta(q) - \theta(q^+)}{v}, \tag{4.7.15}$$

where q^+ is the post-impact configuration of the robot. Therefore, τ directly couples the phase of the robot to the forward progression of the velocity modulating output. As in the case of purely position modulating virtual constraints (4.7.9), the goal is to construct a controller that drives $y_1 \to 0$ and $y_2 \to 0$ to force the robot to progress forward in a desired fashion while displaying the coupling dictated by y_2.

4.7.4 USING FEEDBACK CONTROL TO IMPOSE VIRTUAL CONSTRAINTS

As discussed in Section 4.7.3, the goal is to drive $y \to 0$ in order to force the actual outputs, y_a, to the desired outputs, y_d, i.e., in order to achieve $y_a \to y_d$. This objective can be achieved through the use of a core tool in nonlinear control: *feedback linearization* (Sastry, 1999). The end result is a controller that drives the system to the *zero dynamics* surface and renders this surface invariant through the continuous dynamics. Therefore, applying this control law implies that the full dynamics of the robot will ultimately evolve on a low dimensional space *for the continuous dynamics*. The next section will discuss how to achieve this through the full hybrid dynamics of the robot and, thereby, realize periodic walking and running motions in the hybrid models of Section 4.7.2.

4.7.4.1 Feedback Linearization

The goal of feedback linearization is to uncover a relationship between the output and the control input. This is achieved by differentiating the output until this relationship is revealed. To be concrete, let us consider differentiating y in (4.7.9) along solutions of the continuous dynamics (4.7.3). We have

$$\dot{y}(q,\dot{q}) = \frac{\partial y(q)}{\partial q}\dot{q}. \tag{4.7.16}$$

Since none of the inputs appear in this equation, we differentiate a second time to obtain

$$\ddot{y}(q,\dot{q}) = \frac{\partial}{\partial q}\left(\frac{\partial h(q)}{\partial q}\dot{q}\right)\dot{q} + \frac{\partial y(q)}{\partial q}\ddot{q} \tag{4.7.17}$$

and substituting in the dynamics (4.7.1) yields

$$\ddot{y}(q,\dot{q}) = \underbrace{\frac{\partial}{\partial q}\left(\frac{\partial h(q)}{\partial q}\dot{q}\right)\dot{q} + \frac{\partial y(q)}{\partial q}\left[-D^{-1}(q)\left(C(q,\dot{q})\dot{q} + G(q)\right)\right]}_{L_f^2 y(q,\dot{q})}$$

$$+ \underbrace{\frac{\partial y(q)}{\partial q} D^{-1}(q) B}_{L_g L_f y(q,\dot{q})} u \qquad (4.7.18)$$

where, since we differentiated y twice to obtain the input, the virtual constraint is a *relative degree two output* in the terminology of nonlinear control (Isidori, 1995; Sastry, 1999). In the context of the mixed position and velocity modulating outputs (4.7.13)–(4.7.14), the result will be a mixed relative degree, as will be discussed in more detail in Section 4.7.4.3 below.

To obtain a controller that drives $y \to 0$ we consider (4.7.18) which can be written in terms of x as

$$\ddot{y}(x) = L_f^2 y(x) + L_g L_f y(x) u, \qquad (4.7.19)$$

where $L_g L_f y(x) \in \mathbb{R}^{m \times m}$ is the *decoupling matrix* that is assumed to be invertible. Therefore, selecting

$$u(x, \mu) = \left(L_g L_f y(x) \right)^{-1} \left[-L_f^2 y(x) + \mu \right], \qquad (4.7.20)$$

where $\mu \in \mathbb{R}^m$ is an auxillary input, results in a linear relationship between the second derivative of y and the new input μ, as in

$$\ddot{y} = \mu. \qquad (4.7.21)$$

That is, the end result is a linear control system of the form

$$\begin{bmatrix} \dot{y} \\ \ddot{y} \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & I \\ 0 & 0 \end{bmatrix}}_{F} \begin{bmatrix} y \\ \dot{y} \end{bmatrix} + \underbrace{\begin{bmatrix} 0 \\ I \end{bmatrix}}_{G} \mu \qquad (4.7.22)$$

where $I \in \mathbb{R}^{m \times m}$ is the identity matrix. Therefore, the control law

$$\mu_\epsilon(y, \dot{y}) = -\frac{K_P}{\epsilon^2} y - \frac{K_D}{\epsilon} \dot{y} \qquad \Rightarrow \qquad \begin{bmatrix} \dot{y} \\ \ddot{y} \end{bmatrix} = \frac{1}{\epsilon} \underbrace{\begin{bmatrix} 0 & \epsilon I \\ -\frac{1}{\epsilon} K_P & -K_D \end{bmatrix}}_{F_{\mathrm{cl}(\epsilon)}} \begin{bmatrix} y \\ \dot{y} \end{bmatrix}$$

$$(4.7.23)$$

where K_P and K_D are chosen so that $F_{\mathrm{cl}}(\epsilon)$ is stable (Hurwitz) for all $0 < \epsilon < 1$. Note that here ϵ forces the system to converge at a user defined rate. Therefore, the control law

$$u^*(x) = \left(L_g L_f y(x) \right)^{-1} \left[-L_f^2 y(x) - \frac{K_P}{\epsilon^2} y - \frac{K_D}{\epsilon} \dot{y} \right] \qquad (4.7.24)$$
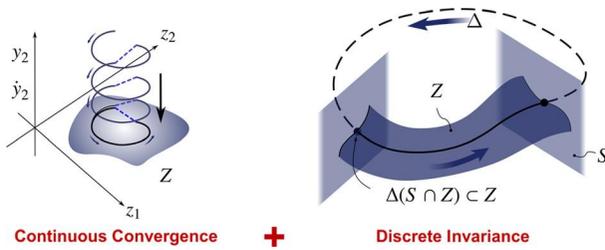
FIGURE 4.7.5   Illustration of the key concepts related to hybrid zero dynamics: continuous convergence to a low dimensional zero dynamics surface Z, coupled with a hybrid invariance condition: $\Delta(S \cap Z) \subset Z$.

drives the output y and its derivative $\dot{y}$ to zero exponentially fast at a rate of $\frac{1}{\epsilon}$.

4.7.4.2   Zero Dynamics

The control law introduced in Section 4.7.4.1 drives $y(x) \to 0$ and $\dot{y}(x) \to 0$. That is, it drives the continuous dynamics to the *zero dynamics surface*, illustrated in Fig. 4.7.5 and defined by

$$Z := \{x \in D \mid y(x) = 0 \text{ and } \dot{y}(x) = 0\}, \qquad (4.7.25)$$

where the dimension of the surface is the degree of underactuation of the system $2(n - m)$; it is reminded that n is the number of the degrees of freedom of the robot and m is the number of actuators. One can find (local) coordinates for the zero dynamics surface, Z, given by $z : D \to \mathbb{R}^{2(n-m)}$ such that $(z, y, \dot{y}) : D \to \mathbb{R}^{2n}$ is a (local) diffeomorphism. To provide a concrete example, suppose that $m = n - 1$ and the angle between the robot and the ground is the first coordinate, q_1. Then, necessarily, $\theta(q)$ is a function of q_1 and $y(q)$ is independent of q_1. We can, therefore, pick coordinates for the zero dynamics as follows (Westervelt et al., 2007):

$$z_1(q) = \theta(q),$$
$$z_2(q) = D(q)_{(1,*)}\dot{q}$$

where $D(q)_{(1,*)}$ is the first row of the inertia matrix in (4.7.1).

Utilizing the coordinates for the zero dynamics and letting $\eta = (y, \dot{y})$ be the coordinates for the dynamics transversal to Z, the system can be represented as

$$\dot{\eta} = \widehat{f}(\eta, z) + \widehat{g}(\eta, z)u, \qquad (4.7.26)$$
$$\dot{z} = w(\eta, z)$$

where

$$\widehat{f}(\eta(x), z(x)) = \begin{bmatrix} \dot{y}(x) \\ L_f^2 y(x) \end{bmatrix}, \qquad \widehat{g}(\eta(x), z(x)) = \begin{bmatrix} 0 \\ L_g L_f y(x) \end{bmatrix},$$

which are written here in the original x coordinates by using the fact that η is defined in terms of y and $\dot{y}$, which, in turn, are functions of x; note that these expressions can be converted to the (η, z) coordinates through the (local) diffeomorphism relating (η, z) with x. Additionally, the control law (4.7.20) with the auxiliary input μ chosen as in (4.7.23) and expressed in terms of the η coordinates results in a linear, time-invariant system $\dot{\eta} = F_{cl}(\epsilon)\eta$ describing the dynamics transversal to Z. Effectively, the feedback control law (4.7.24) of Section 4.7.4.1 ensures that the zero dynamics surface Z is attractive and invariant under the continuous time dynamics of the system – that is, $\eta \to 0$ and $\eta(0) = 0$ implies that $\eta(t) \equiv 0$ for all future times $t \geq 0$. As a result, the *zero dynamics* – that is, the maximal dynamics compatible with the output being identically equal to zero – can be written as

$$\dot{z} = w(0, z). \tag{4.7.27}$$

It is worth mentioning that, to arrive at (4.7.27), the number of outputs equals the number of inputs. In other words, all the inputs available for control are "slaved" to drive the outputs to zero. Depending on the controller's objectives, however, it is possible to define the vector of outputs y so that its dimension is smaller than the dimension of the input vector u, and keep the remaining control inputs for additional control *within* the zero dynamics, which now becomes *controlled*. This may increase the dimension of the zero dynamics, but it provides greater flexibility for developing control action. Examples of this approach include the stabilization of running motions on compliant robots by "shaping" compliance within the zero dynamics (Poulakakis and Grizzle, 2007a, 2007b, 2009b), or by incorporating active force control (Sreenath et al., 2013); see Section 4.7.6.4 below for more details. On a final note, the presence of exogenous inputs in the zero dynamics may result from externally applied forces, giving rise to *forced* zero dynamics, as in Veer et al. (2015, 2016). For example, this is the case when a bipedal robot physically collaborates with a leading external agent – another robot or a human – to transport an object in their workspace (Motahar et al., 2015b). In this case, the objective of the feedback controller is to adapt the robot's locomotion pattern to the externally applied force (Veer et al., 2015, 2016).

4.7.4.3   Partial Zero Dynamics

We can also consider the case in which there is a velocity modulating output (Ames, 2014), i.e., where there are virtual constraints of the form (4.7.13) and (4.7.14). In this case, differentiating y_1 until the control input appears, as was done in (4.7.18) yields

$$\dot{y}_1(q, \dot{q}) = \underbrace{\frac{\partial}{\partial q}\left(\frac{\partial \theta(q)}{\partial q}\dot{q}\right)\dot{q} + \frac{\partial \theta(q)}{\partial q}\left[-D^{-1}(q)\left(C(q, \dot{q})\dot{q} + G(q)\right)\right]}_{L_f y_1(q, \dot{q})}$$

$$+ \underbrace{\frac{\partial \theta(q)}{\partial q}D^{-1}(q)B}_{L_g y_1(q, \dot{q})} u \tag{4.7.28}$$

because y_1 depends on the angular velocity in (4.7.13). Since the control input appears after differentiating once, it implies that y_1 has relative degree one. Therefore, in the case of human-inspired output combinations we mixed relative degree one and relative degree two outputs. That is, we can combine (4.7.28) with (4.7.18) (with y replaced by y_2) to obtain

$$\begin{bmatrix} \dot{y}_1 \\ \ddot{y}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} L_f y_1(q, \dot{q}) \\ L_f^2 y_2(q, \dot{q}) \end{bmatrix}}_{L_f(x)} + \underbrace{\begin{bmatrix} L_g y_1(q, \dot{q}) \\ L_g L_f y_2(q, \dot{q}) \end{bmatrix}}_{A(x)} u \tag{4.7.29}$$

where $A(x)$ is the decoupling matrix that must be full rank. Therefore, analogously to (4.7.20), we have

$$u(x, \mu) = A(x)^{-1}\left(-L_f(x) + \mu\right)$$

$$\Rightarrow \quad \begin{bmatrix} \dot{y}_1 \\ \ddot{y}_2 \end{bmatrix} = \mu \tag{4.7.30}$$

$$\Rightarrow \quad \begin{bmatrix} \dot{y}_1 \\ \dot{y}_2 \\ \ddot{y}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{bmatrix}}_{F} \begin{bmatrix} y_1 \\ y_2 \\ \dot{y}_2 \end{bmatrix} + \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & I \end{bmatrix}}_{G} \mu.$$

Here $\mu = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}$ where μ_1 is the input to $\dot{y}_1$ and μ_2 is the input to $\ddot{y}_2$. As in (4.7.23), we can pick

$$\mu(y_1, y_2, \dot{y}_2) = \begin{bmatrix} -\frac{1}{\epsilon}y_1 \\ -\frac{K_P}{\epsilon^2}y_2 - \frac{K_D}{\epsilon}\dot{y}_2 \end{bmatrix}$$

$$\Rightarrow \quad \begin{bmatrix} \dot{y}_1 \\ \dot{y}_2 \\ \ddot{y}_2 \end{bmatrix} = \frac{1}{\epsilon} \underbrace{\begin{bmatrix} -1 & 0 & 0 \\ 0 & 0 & \epsilon I \\ 0 & -\frac{1}{\epsilon} K_{\mathrm{P}} & -K_{\mathrm{D}} \end{bmatrix}}_{F_{\mathrm{cl}}(\epsilon)} \begin{bmatrix} y_1 \\ y_2 \\ \dot{y}_2 \end{bmatrix} \qquad (4.7.31)$$

such that $F_{\mathrm{cl}}(\epsilon)$ is Hurwitz. Therefore, we have defined a controller that achieves $y_1 \to 0$, $y_2 \to 0$ and $\dot{y}_2 \to 0$ (at a rate of $\frac{1}{\epsilon}$). Moreover, because we considered the velocity modulating virtual constraint, we have that $\dot{\theta}(q) \to v$. Hence, the velocity of the system – as it is captured by the rate $\dot{\theta}(q)$ of the phase variable – converges to the desired value.

As was discussed in Section 4.7.4.2, we can consider the surface that the system converges to under the feedback control law introduced in (4.7.29). In this case, while our virtual constraints consist of both relative degree one and relative degree two outputs, we will only consider the surface that the relative degree two outputs converge to, and then study the behavior of the relative degree one outputs on this surface. In particular, consider the *partial hybrid zero dynamics* surface given by Ames (2014)

$$PZ := \{x \in D \mid y_2(x) = 0 \text{ and } \dot{y}_2(x) = 0\}, \qquad (4.7.32)$$

which is rendered attractive by (4.7.31). Writing $\eta = (y_2, \dot{y}_2)$, then we can again write the system in the form given by (4.7.26). The advantage of the partial zero dynamics can be seen easiest in the case of full actuation, i.e., when $n = m$, as is the case for many humanoid robots. In this fully-actuated case, the z dynamics in (4.7.26) become *controlled* with the ankle torque of the robot available to propel the robot forward. This is evidenced by the fact that, in this case, the coordinates for z can be chosen as $z_1 = \theta(q)$ and $z_2 = \dot{\theta}(q, \dot{q})$ wherein, by (4.7.30) and (4.7.31), the q dynamics become linear:

$$\dot{z}_1 = z_2. \qquad (4.7.33)$$

$$\dot{z}_2 = -\frac{1}{\epsilon}(z_2 - v).$$

That is, on the surface PZ the system evolves according to linear dynamics that drive $\dot{\theta} \to v$. It will be seen later that this ensures a stable walking gait in the case of full actuation.

4.7.5 GENERATING PERIODIC MOTIONS

The goal of gait synthesis is to generate periodic walking gaits for a bipedal robot, along with the feedback controller that enforces these periodic motions. This is where *hybrid zero dynamics (HZD)* provides a powerful framework

(Grizzle et al., 2001; Westervelt et al., 2003; Westervelt, 2003; Morris and Griz- zle, 2009). In particular, the feedback controller introduced in Section 4.7.4 rendered the zero dynamics surface, Z, both attractive and invariant *for the con- tinuous dynamics*. Yet, when the system reaches the switching surface S it will be "thrown away" from the zero dynamics surface. This has the potential to destabilize the system even if the dynamics in Z are stable. That is, the dis- crete dynamics in the hybrid system can destroy continuous-time invariance and destabilize the system, even if the continuous dynamics is well behaved. This is the core idea behind hybrid zero dynamics. By ensuring *hybrid invariance* of the zero dynamics (see Fig. 4.7.5),

$$\Delta(S \cap Z) \subset Z, \tag{4.7.34}$$

it prevents the system from being destabilized through impact – in fact, the main result of hybrid zero dynamics is that the condition (4.7.34) implies stability of the overall dynamics provided that the zero dynamics are stable. This section will establish the fundamental results related to hybrid zero dynamics.

4.7.5.1 Hybrid Zero Dynamics

Under the influence of the controllers discussed in Section 4.7.4, the "open- loop" hybrid control system (4.7.6) takes the form of the "closed-loop" hybrid dynamical system

$$\mathscr{H} : \begin{cases} \dot{x} = f_{\text{cl}}(x), & x \in D \backslash S, \\ x^+ = \Delta\left(x^-\right), & x^- \in S, \end{cases} \tag{4.7.35}$$

where

$$f_{\text{cl}}(x) = f(x) + g(x)u(x)$$

and $u(x)$ is the feedback controller given in (4.7.24).

Recall that the feedback controller in (4.7.24) rendered the zero dynamics surface, Z, given in (4.7.25) exponentially stable. Yet, it may be the case that for $x^- \in Z$, the post-impact state of the system $x^+ = \Delta(x^-) \notin Z$. This implies that the pre-impact state is "thrown-away" from the zero dynamics surface. There- fore, if the impacts occur at a rate faster than the controller can stabilize the system, the end result is that the impacts will destabilize the system. Therefore, the core condition that we enforce is *hybrid invariance*, i.e.,

$$\Delta(S \cap Z) \subset Z \quad \text{or, equivalently,} \quad x^- \in S \cap Z \quad \Rightarrow \quad x^+ = \Delta(x^-) \in Z.$$
$$\tag{4.7.36}$$

Enforcing (4.7.36) requires the proper choice of virtual constraints. In particular, recall that the controller in (4.7.24) was synthesized from the virtual constraints given in (4.7.9) which, in turn, depended on the parameter set $\alpha \in \mathbb{R}^k$ with k the total number of parameters of the desired outputs, $y_d(\tau(q), \alpha)$. Therefore, the zero dynamics surface Z depends on the parameters α so that changing the values of these parameters changes the shape of the surface. This allows us to explicitly shape the zero dynamics surface to enforce the hybrid invariance condition (4.7.36), and it can be done systematically in the context of a nonlinear constrained optimization problem of the form

$$\alpha^* = \underset{\alpha \in \mathbb{R}^k}{\mathrm{argmin}}\ \mathrm{Cost}(\alpha) \qquad \text{(HZD Optimization)}$$

$$\text{s.t.} \quad \Delta(S \cap Z) \subset Z \qquad \text{(HZD)}$$

where Cost is a user defined cost function that can be chosen to produce walking gaits with desirable properties. The details of the optimization problem are beyond the scope of this chapter, but can be found in Westervelt et al. (2007), Ames (2014). We only mention here that any physically relevant constraints – such as constraints on torque, angular velocity, ground reaction forces and friction limitations – can be added to the optimization. Adding such constraints ensures the physical realizability of the resulting walking gait (Zhao et al., 2014b; Hereid et al., 2016).

The choice of cost function in (HZD Optimization) can determine the "shape" of the resulting gait, i.e., the overall behavior of the gait. For example, a common choice is a cost function that minimizes the overall torque while maximizing the distance traveled, i.e.,

$$\mathrm{Cost}(\alpha) = \frac{1}{\text{step length}} \int_0^T \|u(\alpha)\|^2 dt \qquad (4.7.37)$$

where $u(\alpha)$ is the feedback controller calculated for a given parameter set. Another common choice of cost function, especially in the context of *human-inspired control* (Ames, 2014), is to use the difference between the outputs as calculated from human data and the desired functions with parameters, α, seeded from human data (Ames, 2012); the end result is typically "human-like" walking gaits. It is interesting to note that, with the proper choice of constraints, one often sees similar gaits independent of specific cost functions as the constraints tend to be a large factor in the resulting look of the gait. Finally, if efficiency is the goal, a cost function that minimizes the cost of transport can be selected (Reher et al., 2016a).

Given constraint parameters that yield well-defined hybrid zero dynamics, the end result is that the system evolves on the zero dynamics surface during the

continuous dynamics and this surface is invariant through impact. Therefore, when supplied with initial conditions in Z, the dynamics of the system evolves according to the *restricted hybrid system*:

$$\mathcal{H}|_Z : \begin{cases} \dot{z} = w(0, z), & z \in Z \backslash S \cap Z, \\ z^+ = \Delta_Z(z^-), & z^- \in S \cap Z, \end{cases} \tag{4.7.38}$$

with w the zero dynamics given in (4.7.27), and $\Delta|_z : S \cap Z \to Z$ the restriction of the impact map to Z. This system is low-dimensional, e.g., for one degree of underactuation it is a two-dimensional system, and the behavior of this system dictates the behavior of the full order dynamics, $\mathcal{H}$, regardless of the dimension of the full order dynamics. For example, for a 23 degree of freedom humanoid robot (as is the case for DURUS), $\mathcal{H}$ will be a 46 dimensional hybrid system, but its behavior will still be completely determined by the behavior of the restricted hybrid system, which may have dimension as low as two (Reher et al., 2016b).

To examine how the behavior of the restricted hybrid system $\mathcal{H}|_Z$ affects the behavior of the full-order hybrid model $\mathcal{H}$, we consider periodic orbits corresponding to walking gaits of interest. In particular, for the full-order dynamics (4.7.6), let $\phi_t^{f_{cl}}(x_0)$ be the (unique) solution to the continuous dynamics $\dot{x} = f_{cl}(x)$ at time $t \geq 0$ with initial condition x_0 (where we assume local Lipschitz continuity of $f_{cl}(x)$). For $x^* \in S$ we say that $\phi_t^{f_{cl}}$ is hybrid periodic if there exists a $T > 0$ such that $\phi_T^{f_{cl}}(\Delta(x^*)) = x^*$. Given a hybrid periodic solution, we are interested in considering the stability of the corresponding hybrid periodic orbit,

$$\mathcal{O} = \{\phi_t^{f_{cl}}(\Delta(x^*)) : 0 \leq t \leq T\}.$$

To study the stability of this orbit, we first consider the *time-to-impact function*

$$T_I(x) = \inf\{t > 0 : \phi_t^{f_{cl}}(\Delta(x)) \in S, \text{ with } x \in S\},$$

which is well defined by the implicit function applied to the function $H(t, x) = p_{toe}^v(\phi_t^{f_{cl}}(\Delta(x)))$ (where p_{toe}^v is the vertical position of the toe used to define S in (4.7.4)) since it satisfies $H(T, x^*) = 0$. The end result is the *Poincaré map* which is a $P : S \to S$ which is well defined in a neighborhood of x^*, and is given by

$$P(x) = \phi_{T_I(x)}^{f_{cl}}(x).$$

Importantly, the stability of the periodic orbit $\mathcal{O}$ is equivalent to the stability of the Poincaré map viewed as a discrete time dynamical system $x_{n+1} = P(x_n)$ with fixed point $x^* = P(x^*)$, i.e.,

$\mathcal{O}$ is exponentially stable

$\Leftrightarrow \quad x^*$ is an exponentially stable equilibrium point of P,

see Morris and Grizzle (2009) for a proof, and Sastry (1999), Wendel and Ames (2010, 2012), Burden et al. (2011) for more details on Poincaré maps. These constructions can also be applied to the restricted hybrid system given in (4.7.38). That is, given a hybrid periodic orbit $\mathcal{O}_Z \subset Z$, there is the associated *restricted Poincaré map* $P_Z : Z \cap S \rightarrow Z \cap S$ which determines the stability of $\mathcal{O}_Z$. Moreover, given a hybrid periodic orbit $\mathcal{O}_Z$, since Z is an invariant subspace of D it follows that $\mathcal{O} = \iota(\mathcal{O}_Z)$ is a hybrid periodic orbit; here $\iota : Z \hookrightarrow D$ is the canonical embedding.

We now have the technical machinery to state the main result for hybrid zero dynamics. Intuitively, this results states that

$\mathcal{O}_Z$ is exponentially stable $\quad \Rightarrow \quad \mathcal{O} = \iota(\mathcal{O}_Z)$ is exponentially stable.

More formally, we have the following fundamental theorem of hybrid zero dynamics (Westervelt et al., 2007, 2003; Grizzle et al., 2001; Morris and Grizzle, 2009):

Theorem 1 (Hybrid Zero Dynamics). *Consider the hybrid control system $\mathscr{HC}$ given in (4.7.6) with the control law in (4.7.24) applied to obtain the hybrid system $\mathscr{H}$ given in (4.7.35), and assume hybrid zero dynamics (4.7.34), $\Delta(S \cap Z) \subset Z$. If there exists a locally exponentially stable hybrid periodic orbit $\mathcal{O}_Z$ of the restricted hybrid system $\mathscr{H}|_Z$, then there exists an $\overline{\epsilon} > 0$ such that for all $\overline{\epsilon} > \epsilon > 0$ the hybrid periodic orbit $\mathcal{O} = \iota(\mathcal{O}_Z)$ is locally exponentially stable for the full-order hybrid system $\mathscr{H}$.*

The importance of this result is that the zero dynamics provides a *substantially* lower dimensional surface in which to search for stable periodic orbits. In fact, in the case when the robot has one degree of underactuation ($m = n - 1$) closed form expressions can be obtained that guarantee the existence and stability of a hybrid periodic orbit $\mathcal{O}_Z$; this can be added directly to the optimization problem in (HZD Optimization) as a constraint to guarantee that any parameter set produced by the optimization is exponentially stable for the full-order dynamics of the system. In the case of full actuation, even stronger conclusions can be reached.

4.7.5.2 Partial Hybrid Zero Dynamics

In the context of gait generation for humanoid robots, one often has the luxury of dealing with a fully actuated system; in this case, partial hybrid zero dynamics (PHZD) provides a useful tool for gait generation. In particular, for PHZD

we consider virtual constraints of the form (4.7.13) and (4.7.14) that allow for virtual configuration constraints on the robot via y_2, along with the ability to regulate the forward progression of the robot via y_1. Applying this controller developed in Section 4.7.4.3 to the hybrid control system $\mathscr{HC}$ given in (4.7.6) results in a closed-loop hybrid dynamical system $\mathscr{H}$ as in (4.7.35) except, in this case, $u(x)$ is given in (4.7.30). Additionally, we established that this controller resulted in the corresponding partial zero dynamics surface PZ being both attractive and invariant. Therefore, if this surface is invariant through impact:

$$\Delta(S \cap PZ) \subset PZ \qquad \text{(PHZD)}$$

the end result is *partial hybrid zero dynamics*. As in the case of hybrid zero dynamics, we can consider the optimization problem (HZD Optimization) with the constraint (HZD) replaced with (PHZD). Given a parameter set that solves this optimization problem, we have the corresponding restricted hybrid system $\mathscr{H}|_{P}Z$. The advantage, in this case, is that the dynamics $\dot{z} = w(0, z)$ take the simple linear form given in (4.7.33). Since the y_2 dynamics are also linear by choice of controller, the entire hybrid system becomes a *linear* hybrid dynamical system. The structural properties associated with the PHZD motivates the following key partial hybrid zero dynamics result (Ames, 2014).

Theorem 2 (Partial Hybrid Zero Dynamics). *Let $\mathscr{HC}$ given in (4.7.6) be fully actuated, with the control law in (4.7.30) applied to obtain a hybrid system $\mathscr{H}$, and assume partial hybrid zero dynamics (PHZD): $\Delta(S \cap PZ) \subset PZ$. Then, there exists a locally exponentially stable hybrid periodic orbit $\mathcal{O}_{PZ}$ of the restricted hybrid system $\mathscr{H}|_{PZ}$, and an $\bar{\epsilon} > 0$ such that for all $\bar{\epsilon} > \epsilon > 0$ the hybrid periodic orbit $\mathcal{O} = \iota(\mathcal{O}_{PZ})$ is locally exponentially stable for the full-order hybrid system $\mathscr{H}$.*

That is, in the case of fully actuated robots (such as traditional humanoids), we have the following intuitive representation of Theorem 2:

$$\text{Fully Actuated} + \Delta(S \cap PZ) \subset PZ$$
$$\Rightarrow \quad \mathcal{O} = \iota(\mathcal{O}_{PZ}) \text{ is exponentially stable}$$

or, equivalently, the existence of parameters α in (4.7.14) that yield partial hybrid zero dynamics implies a stable walking gait for fully actuated robots. It is important to note that PHZD can also be applied to robots with compliance (or underactuation) – e.g., the humanoid robot DURUS (Hereid et al., 2016; Reher et al., 2016a) – provided that these compliant elements are "normal" to the actuators that allow for forward progression of the robot. In this case, there will exist nontrivial passive dynamics in the partial zero dynamics surface and, therefore, a periodic orbit must be found in this surface to guarantee the existence of a stable periodic orbit in the full-order dynamics.

4.7.5.3 Control Lyapunov Functions

The methods presented thus far involved using feedback to linearize the dynamics of the robotic systems, wherein a linear system was defined to stabilize the virtual constraints. Yet, the control law presented intrinsically ignores the natural dynamics of the system – the main component of the feedback linearization process is to cancel out the dynamics using (4.7.20). Instead of canceling out the nonlinear dynamics, we can leverage them in the context of *control Lyapunov functions*. This has the benefit of yielding an entire class of controllers that stabilize the system. Additionally, these controllers will stabilize periodic orbits in the system in a pointwise optimal fashion.

Let us return to the canonical form of the dynamics obtained before the system was feedback linearized, i.e., the nonlinear system given in (4.7.26). Recall that, for the dynamics in this form, $\eta = (y, \dot{y})$ characterizes the controlled variables of the system while z describes the dynamics encompassed the passive component of the robot. Let us denote by Y the space with η as coordinates so that $D = Y \times Z$. A continuously differentiable function $V_\epsilon : Y \to \mathbb{R}_{\geq 0}$ is a *rapidly exponentially stabilizing control Lyapunov function (RES-CLF)* (Ames et al., 2012b, 2014a) if there exist positive constants $c_1, c_2, c_3 > 0$ such that for all $1 > \epsilon > 0$,

$$c_1 \|\eta\|^2 \leq V_\epsilon(\eta) \leq \frac{c_2}{\epsilon^2} \|\eta\|^2, \tag{4.7.39}$$

$$\inf_{u \in U} \left[L_{\widehat{f}} V_\epsilon(\eta, z) + L_{\widehat{g}} V_\epsilon(\eta, z) u + \frac{c_3}{\epsilon} V_\epsilon(\eta) \right] \leq 0 \tag{4.7.40}$$

for all $(\eta, z) \in Y \times Z$.

The existence of a RES-CLF yields a family of controllers that rapidly exponentially stabilize the system to the zero dynamics. In particular, we can consider the control values

$$K_\epsilon(\eta, z) = \{ u \in U : L_{\widehat{f}} V_\epsilon(\eta, z) + L_{\widehat{g}} V_\epsilon(\eta, z) u + \frac{c_3}{\epsilon} V_\epsilon(\eta) \leq 0 \}, \tag{4.7.41}$$

wherein it follows that

$$u_\epsilon(\eta, z) \in K_\epsilon(\eta, z) \quad \Rightarrow \quad \|\eta(t)\| \leq \frac{1}{\epsilon} \sqrt{\frac{c_2}{c_1}} e^{-\frac{c_3}{2\epsilon} t} \|\eta(0)\|. \tag{4.7.42}$$

Therefore, picking $\epsilon > 0$ to be a small value increases the rate of convergence of η, i.e., increases the rate of convergence to the zero dynamics surface Z. In addition, this yields specific feedback controllers, e.g., the min-norm controller (Freeman and Kokotović, 1996),

$$m_\epsilon(\eta, z) = \text{argmin}\{\|u\| : u \in K_\epsilon(\eta, z)\}. \tag{4.7.43}$$

The importance of RES-CLFs is made apparent by the following theorem, which says that *any* controller $u_\epsilon \in K_\epsilon$ results in a stable orbit for the full-order dynamics if one exists in the reduced order dynamics:

$$u_\epsilon(\eta, z) \in K_\epsilon(\eta, z) + \mathcal{O}_Z \text{ exponentially stable}$$
$$\Rightarrow \quad \mathcal{O} = \iota(\mathcal{O}_Z) \text{ exponentially stable.}$$

Or, more formally, we have the following result on RES-CLF + HZD (Ames et al., 2014a):

Theorem 3 (Control Lyapunov Functions + HZD). *Consider the hybrid control system $\mathscr{HC}$ given in (4.7.6) with any Lipschitz continuous $u_\epsilon(\eta, z) \in K_\epsilon(\eta, z)$ applied to obtain a hybrid system $\mathscr{H}$, and assume hybrid zero dynamics (4.7.34), $\Delta(S \cap Z) \subset Z$. If there exists a locally exponentially stable hybrid periodic orbit $\mathcal{O}_Z$ of the restricted hybrid system $\mathscr{H}|_Z$, then there exists an $\overline{\epsilon} > 0$ such that for all $\overline{\epsilon} > \epsilon > 0$ the hybrid periodic orbit $\mathcal{O} = \iota(\mathcal{O}_Z)$ is locally exponentially stable for the full-order hybrid system $\mathscr{H}$.*

To provide a specific example of an RES-CLF, we can utilize the constructions in Section 4.7.4.1 to obtain a specific RES-CLF. In particular, recall that the feedback linearizing controller resulted in η dynamics of the form (4.7.22), or in η notation,

$$\dot{\eta} = F\eta + G\mu. \tag{4.7.44}$$

For this linear control system, we can consider the continuous-time algebraic Riccati equations (CARE),

$$F^T P + PF - PGG^T P + Q = 0, \tag{4.7.45}$$

with solution $P = P^T > 0$. One can use P to construct a RES-CLF that can be used to exponentially stabilize the output dynamics (4.7.44) at a user defined rate of $\frac{1}{\epsilon}$. In particular, define

$$V_\epsilon(\eta) = \eta^T \underbrace{M_\epsilon P M_\epsilon}_{P_\epsilon} \eta, \qquad \text{with} \quad M_\epsilon = \mathrm{diag}(\epsilon I, I), \tag{4.7.46}$$

wherein it follows that

$$\dot{V}_\epsilon(\eta, \mu) = L_F V_\epsilon(\eta) + L_G V_\epsilon(\eta)\mu$$

with

$$L_F V_\epsilon(\eta) = \eta^T (F^T P_\epsilon + P_\epsilon F)\eta,$$
$$L_G V_\epsilon(\eta) = 2\eta^T P_\epsilon G.$$

Note that it is easy to verify that $V(\eta) = \eta^T P_\epsilon \eta$ is an RES-CLF with $c_1 = \lambda_{\min}(P)$, $c_2 = \lambda_{\max}(P)$, and $c_3 = \gamma = \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$. This can be seen by noting that, from (4.7.45) and the form of F and G, P_ϵ solves the CARE (Ames et al., 2014a),

$$F^T P_\epsilon + P_\epsilon F - \frac{1}{\epsilon} P_\epsilon G G^T P_\epsilon + \frac{1}{\epsilon} M_\epsilon Q M_\epsilon = 0, \qquad (4.7.47)$$

and noting that $\gamma P_\epsilon \le M_\epsilon Q M_\epsilon$ so

$$\inf_\mu \left[L_F V_\epsilon(\eta) + L_G V_\epsilon(\eta)\mu + \frac{\gamma}{\epsilon} V_\epsilon(\eta) \right] \le \eta^T P_\epsilon G(\frac{1}{\epsilon} G^T P_\epsilon + 2\mu) \le 0,$$

which is satisfied, for example, by $\mu(\eta) = -\frac{1}{\epsilon} G^T P_\epsilon \eta$. And, therefore, V_ϵ is an RES-CLF. We can convert this back to a control law u_ϵ via (4.7.20):

$$u_\epsilon(x) = \left(L_g L_f y(x) \right)^{-1} \left[-L_f^2 y(x) + \mu_\epsilon(\eta(x)) \right] \in K_\epsilon(\eta(x), z(x)) \quad (4.7.48)$$

for $\mu_\epsilon(\eta(x))$ satisfying

$$\dot{V}_\epsilon(\eta(x)) = L_F V_\epsilon(\eta(x)) + L_G V_\epsilon(\eta)\mu_\epsilon(\eta(x)) \le \frac{\gamma}{\epsilon} V_\epsilon(\eta(x)) \qquad (4.7.49)$$

where we converted back to the x coordinates, i.e., $\eta(x) = (y(x), \dot{y}(x))$. This gives concrete conditions that can be checked to stabilize walking gaits in the hybrid system $\mathcal{H}$ according to Theorem 3.

It is important to note that RES-CLFs can also be constructed in the case of partial hybrid zero dynamics. In this case, the linear control system for the output dynamics (4.7.30) is described by $\eta = (y_1, y_2, \dot{y}_2)$. From these dynamics, (4.7.45) can be used to construct an RES-CLF as in (4.7.46). The end result is a reformulation of Theorem 2 so that a stable periodic orbit is guaranteed for the full order dynamics for any $u_\epsilon(\eta) \in K_\epsilon(\eta)$. That is, we obtain stable walking through the entire class of controllers that satisfy the inequality constraint obtained via the CLF condition.

4.7.6 EXTENSIONS OF HYBRID ZERO DYNAMICS

The goal of this section is to consider extensions of the HZD framework with a view toward more rich application domains. In particular, we consider the HZD framework developed in this chapter in the context of optimization-based controllers via CLFs, multidomain hybrid system models, their application to powered prostheses and compliant hybrid zero dynamics.

4.7.6.1 CLF-Based QPs

The advantage of the control Lyapunov functions introduced in Section 4.7.5.3 is that they give a family of controllers which stabilize the system. That is, for any $u_\epsilon(\eta, z) \in K_\epsilon(\eta, z)$ the system has a stable periodic gait (given a stable periodic orbit in the zero dynamics); a specific example of this is given by the traditional feedback linearizing controller. The importance of CLFs goes beyond simply producing a class of controllers – it suggests an optimization-based control framework for bipedal robots and, in fact, nonlinear systems in general. This allows for these control methods to be extended to a variety of application domains from robotic walking, to prostheses to manipulation to safety-critical control methods.

To see how control Lyapunov functions yield optimization-based controllers, we can consider the set (4.7.41) giving the family of stabilizing controllers. Note that this set is affine in the control input u and, therefore, the min-norm controller (4.7.43) can be equivalently stated as a quadratic program (QP) of the form:

$$m(\eta, z) = \underset{u \in U = \mathbb{R}^m}{\mathrm{argmin}} \quad u^T u \qquad (4.7.50)$$

$$\text{s.t.} \quad L_{\widehat{f}} V_\epsilon(\eta, z) + L_{\widehat{g}} V_\epsilon(\eta, z) u + \frac{c_3}{\epsilon} V_\epsilon(\eta) \le 0 \qquad \text{(CLF)}$$

where we assume that $U = \mathbb{R}^m$ to ensure solvability of the QP. Not only is this QP guaranteed to have a solution, but the solution can be written in closed form (see Ames et al., 2014a) and is Lipschitz continuous. Moreover, one can utilize the RES-CLF given in (4.7.46) to explicitly construct the inequality constraint in this QP. Finally, because it is a QP it can be solved in real-time; in fact, the CLF-based QP has been implemented in real-time (e.g., at a 1 kHz loop rate) on MABEL (Galloway et al., 2015) and DURUS-2D (Cousineau and Ames, 2015) to achieve dynamic walking. Additionally, it was implemented at over 5 kHz as an embedded level controller on series elastic actuators in Ames and Holley (2014).

The advantage of the QP formulation of CLFs, as opposed to simply utilizing the closed form min-norm solution, is that it allows for additional constraints and objectives to be unified with the CLF. To provide a concrete example, suppose that we have torque bounds on the actuators given by a scalar $u_{\max}$ (similar ideas extend to actuators with different max torques). While one might typically simply saturate the control input, doing so prevents the controller from taking these torque saturations into account. Therefore, through the CLF-based QP framework, we can incorporate the torque bounds directly into the controller

via the quadratic program:

$$u^*(\eta, z) = \underset{(\delta,\mu)\in\mathbb{R}^{n+1}}{\operatorname{argmin}} \quad u^T H(\eta, z)u + p\delta^2 \tag{4.7.51}$$

$$\text{s.t.} \quad L_{\widehat{f}}V_\epsilon(\eta, z) + L_{\widehat{g}}V_\epsilon(\eta, z)u + \frac{c_3}{\epsilon}V_\epsilon(\eta) \leq \delta, \quad \text{(CLF)}$$

$$u \leq u_{\max}\mathbf{1}, \quad \text{(Max Torque)}$$

$$-u \leq u_{\max}\mathbf{1} \quad \text{(Min Torque)}$$

where $H(\eta, z)$ is positive-definite and $p > 0$ is a large value that penalizes for violations of the CLF constraint. That is, we relax the CLF condition to ensure satisfaction of the physical constraints of the system. While this takes away guarantees on achieving the control objective, it will try to achieve convergence of the CLF in a pointwise optimal fashion when at all possible – the result is that the robot is able to accommodate tighter torque bounds than if one was to simply saturate the control input (see Galloway et al., 2015 for a detailed discussion, and experimental implementation). Note that as one expands the number of constraints in the QP, it is important to be aware of the impact on the resulting solvability and, as a byproduct, the continuity of the solutions to the QP; a discussion can be found in Morris et al. (2013), and conditions on continuity in Morris et al. (2015a).

Utilizing the observation that CLFs (and hence control objectives) can be represented as affine constraints in a QP results in a new paradigm for the control of walking robots. In particular, going beyond simply adding torque bounds, one can consider multiobjective controllers consisting of multiple CLF wherein each control objective results in an additional constraint in the QP (Ames and Powell, 2013); for example, in the context of unifying locomotion and manipulation objectives. Additionally, ground reaction forces on the robot also appear in an affine fashion in the dynamics; thus one can use the CLF-based QP framework in the context of force control. Finally, a recent line of work aimed at safety-critical control makes the observation that safety conditions, i.e., set invariance, can be stated in the context of *control barrier functions* which again are affine in the control input (originally formulated in Ames et al. (2014b) and studied in detail in Ames et al. (2016)); this framework has been applied in the context of robotic walking (Nguyen and Sreenath, 2016; Nguyen and Sreenath, 2015, 2016), automotive safety systems (Xu et al., 2015; Ames et al., 2016; Mehra et al., 2015) and swarm robotics (Borrmann et al., 2015; Wang et al., 2016a, 2016b). Therefore, safety conditions can be unified with control objectives, physical constraints, force objectives and safety constraints all in the context of a single optimization-based controller that can be realized in realtime on robotic systems.

4.7.6.2 Multidomain Hybrid Zero Dynamics

The analysis thus far has focused on hybrid system models of walking robots with a single continuous and discrete domain, i.e., single domain hybrid systems. Yet, in the context of advanced walking and running behaviors, it is necessary to consider multidomain hybrid systems models as introduced in Section 4.7.2.4. As indicated in this section, and as motivated by human walking, throughout the course of a step humans naturally display heel–toe behavior in their feet while locomoting (see Fig. 4.7.4). In the context of these types of walking behaviors, the end result is the multi-domain hybrid system model given in (4.7.8); specific examples of this model are shown in Fig. 4.7.4 for the bipedal robots AMBER 2 and ATRIAS (see Zhao et al., 2015a for the specific hybrid system constructions for these robots).

In the context of multidomain hybrid systems, we can extend the concept of hybrid zero dynamics. In particular, we now have a collection of continuous domains $D = \{D_v\}_{v \in V}$ on which we have associated control systems: $\dot{x} = f_v(x) + g_v(x)u_v$. Note that the domains may be of different actuation types, e.g., some may be underactuated while others may be fully actuated or over actuated. For each of these domains, we can define virtual constraints of the form (4.7.9), denoted by $y_{2,v}$; in the case of full (and over) actuation, $y_v = (y_{1,v}, y_{2,v})$ as in (4.7.13) and (4.7.14)), and in the case of over actuation, care must be taken to define constraints that result in a nonsingular decoupling matrix $A(x)$ in (4.7.29). Therefore, we can construct controllers, $u_v(x)$ for each $v \in V$ as in (4.7.20) for the underactuated domains and as in (4.7.30) for the full (and over) actuated domains.

We can consider the zero dynamic surfaces (and partial zero dynamics surfaces) denoted, for notational simplicity, uniformly by

$$Z_v = \left\{ x \in D_v \mid y_{2,v}(x) = 0 \text{ and } \dot{y}_{2,v}(x) = 0 \right\}.$$

Correspondingly, the control laws $u_v(x)$ drive the system to the surface Z_v for each $v \in V$ and, in addition, renders each of these surfaces attractive. To ensure stability of the overall dynamics, we must ensure hybrid zero dynamics for all of the discrete transitions, i.e., *multidomain hybrid zero dynamics (MDHZD)*:

$$\Delta_e(S_e \cap Z_{\text{source}(e)}) \subset Z_{\text{target}(e)}, \qquad \forall e \in E \qquad \text{(Multidomain HZD)}$$

where source(e) and target(e) are the source and target of the edge $e \in E$ of the oriented graph Γ in (4.7.8), respectively. As in Section 4.7.5, if the multidomain hybrid system (4.7.8) has MDHZD, then if there is an exponentially stable periodic orbit contained in Z_v for $v \in V$, then there exists an exponentially stable hybrid periodic in the full order dynamics when the control laws $u_v(x)$ are applied in each domain.
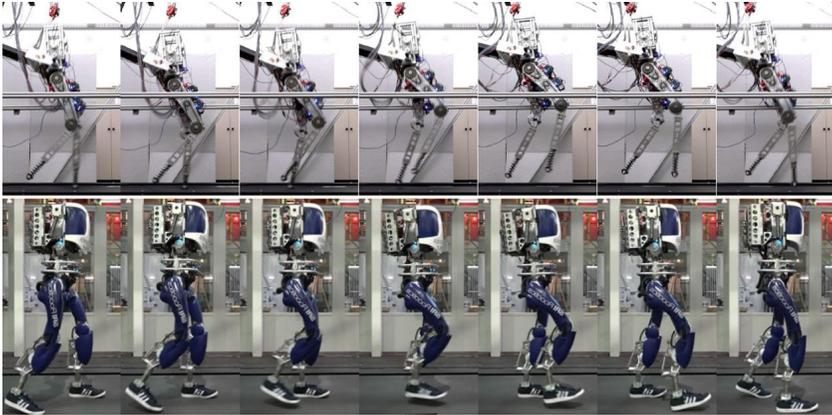
FIGURE 4.7.6 Multicontact walking and running utilizing multidomain hybrid system models realized on the DURUS-2D running robot and the humanoid robot DURUS.

This framework has been applied to numerous bipedal robots to obtain stable walking and running. Specific examples of walking robots include MABEL (Sreenath et al., 2011), AMBER 2 (Zhao et al., 2014b), ATRIAS (Hereid et al., 2014), and DURUS (Hereid et al., 2016). In the case of DURUS, due to the passive springs in the ankles, a two domain hybrid system model was considered. The end result was stable 3D robotic walking, demonstrated publicly during the DARPA Robotics Challenge, where the motion was sustained for over 5 hours with the robot traversing almost 4 km on a treadmill. Importantly, this was the most efficient walking ever realized on a bipedal humanoid robot (Reher et al., 2016a). This can be attributed to the fact that the MDHZD allows for the full dynamics of the robot to be utilized in the generation of walking gaits (through the shaping of the surfaces Z_v) and, importantly for the compliant elements in the system, e.g., springs, to be fully utilized during the walking gait. Recently, these methods were extended to yield a four-domain model of DURUS capturing the natural heel-toe behavior of the foot that humans display when locomoting; the end result was dynamic walking that is efficient and human-like (Reher et al., 2016b) (tiles of this walking gait are shown in Fig. 4.7.6). Finally, note that running motions provide natural examples of two-domain models (Chevallereau et al., 2005; Morris et al., 2006; Poulakakis and Grizzle, 2009b) – see also Fig. 4.7.8 below – and that additional domains can be introduced depending on the control action to enhance control authority over the system, as in the control of running on Thumper (Poulakakis and Grizzle, 2009a), MABEL (Sreenath et al., 2013) and the DURUS-2D runner (Ma et al., 2017) (shown in Fig. 4.7.6).

4.7.6.3 Application to Prostheses

The concepts presented throughout this chapter have natural application to powered prostheses. In particular, a prosthetic device can be simply viewed as a component of a bipedal robot and, with proper representation of the human and the interaction of the human with the device, one can generate controllers for the device via HZD-based methods (Zhao et al., 2011). The core idea to synthesizing prosthetic controllers is to first model the human and the prosthesis as two robotic systems that are coupled at the prosthetic attachment (see Fig. 4.7.7); the parameters of the "human" are taken from measurements of the human and used to generate a corresponding model, and the model of the prosthesis is then added on the affected leg to yield an overall model of the combined human–robot system. This robotic model can then be approached in the same way one would approach generating walking gaits for bipedal robots: the hybrid system model is constructed based upon the desired foot behavior, and gaits are generated through an optimization problem that enforces the HZD conditions together with physical constraints. This idea was first explored in the context of human-inspired control (Sinnet et al., 2011b), and was experimentally validated through the application on both bipedal robots (wherein one leg of the robot plays the role of the prosthesis) (Zhao et al., 2014a), followed by the evaluation with an amputee subject (Zhao et al., 2011). The advantage of the HZD-based approach to designing controllers for powered prostheses is that all of the advanced control and locomotion related concepts of this chapter can be translated to this domain. In particular, multi-domain hybrid system models of locomotion can be utilized to achieve advanced foot behaviors on the device (Zhao et al., 2016b). Additionally, CLF-based QP controllers (as in Section 4.7.6.1) can be realized on prosthetic devices in realtime through a novel model-independent variant (Zhao et al., 2015b); this allows for efficient locomotion that leverages the use of compliant elements as in AMPRO 3 shown in Fig. 4.7.6.

4.7.6.4 Compliant Hybrid Zero Dynamics

To recover part of the energy required to sustain cyclic walking or running motions in legged robots and to ensure safe interaction with the ground surface, compliant elements in the form of mechanical springs have been incorporated in the legs of many such platforms; in the context of robotic bipeds, Thumper (Hurst and Rizzi, 2008), MABEL (Grizzle et al., 2009), and ATRIAS (Hubicki et al., 2016) are just few examples of robots in this family. The role of elastic energy storage in compliant elements becomes more prominent in running motions (McMahon and Cheng, 1990; Alexander, 1990). However, the inclusion of physical springs in a robot's structure poses additional
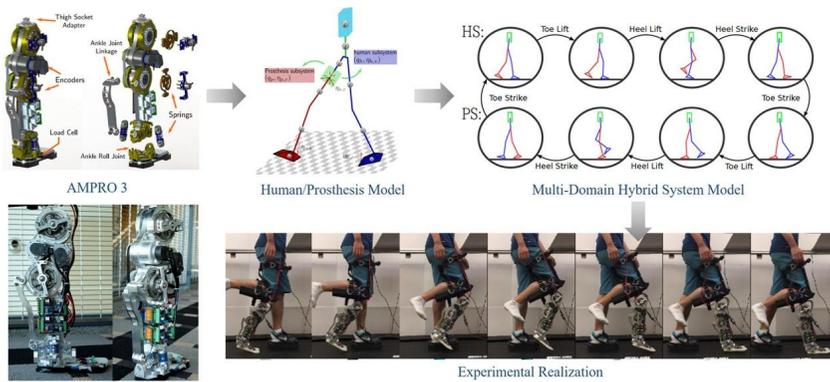
FIGURE 4.7.7 Illustration of the application of HZD methods to prostheses; in this case, in the context of the third generation AMPRO device utilized to achieve multidomain walking (Zhao et al., 2016a, 2017).

challenges to control design. More specifically, in the pursuit of closed-loop stability, the control action must actively exploit open-loop compliance instead of – as is usual in the control of flexible mechanisms – working to replace it. The concept of *compliant hybrid zero dynamics* introduced in Poulakakis (2008) extends HZD controllers so that open-loop compliance is "preserved" in the closed-loop system and determines its behavior. To avoid complexity, we first discuss the main ideas of this method in the context of a simplified hopping model – namely, the *asymmetric spring-loaded inverted pendulum (ASLIP)* (Poulakakis and Grizzle, 2007a, 2009b) – and then provide some information on its application to the control of walking and running motions in MABEL (Sreenath et al., 2011, 2013).

The Asymmetric Spring-Loaded Inverted Pendulum

The ASLIP shown in Fig. 4.7.8 was originally proposed in Poulakakis and Grizzle (2007a) as an intermediate model to bridge the gap between point-mass SLIP-like models and monopedal robots with significant torso pitch dynamics. The ASLIP includes a torso nontrivially coupled to the leg motion,[6] an issue not addressed in the widely studied SLIP, or in its straightforward extensions in which the torso COM coincides with the hip joint. As in the SLIP, the ASLIP features a massless leg and the contact between the leg end and the ground is modeled as an unactuated pin joint.

---

6. Along the same lines with the ASLIP, the Virtual Pivot Point (VPP) model was introduced in Maus et al. (2010) as a template for studying torso stabilization in running; see Subchapters 2.3 and 3.6 for more details.
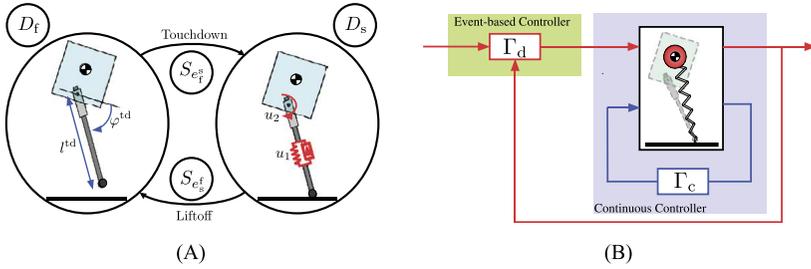
(A)   (B)

FIGURE 4.7.8 (A) The asymmetric spring-loaded inverted pendulum (ASLIP). Running comprises stance and flight phases, separated by touchdown and liftoff events. (B) The structure of the SLIP embedding controller.

The ASLIP alternates between stance and flight phases – denoted by "s" and "f", respectively – resulting in a multidomain hybrid system of the type described in Section 4.7.6.2. Let $\Gamma = (V, E)$, where $V = \{s, f\}$ and $E = \{e_s^f, e_f^s\}$, be the oriented graph that captures the contact state of the model; in this notation, e_s^f and e_f^s denote transition from stance to flight and vice versa, respectively. The model consists of two domains D_s and D_f, within which the dynamics $FG = \{(f_v, g_v)_{v \in V}\}$ of the ASLIP evolve until the state intersects the corresponding switching surface $S = \{S_e\}_{e \in E}$; at this point, a switching map $\Delta = \{\Delta_e\}_{e \in E}$ is triggered to provide initial conditions for the ensuing phase. During stance, the ASLIP is controlled by two inputs: the force u_1 acting along the leg and the torque u_2 applied at the hip; $(u_1, u_2) \in U_s$, where U_s is the set of the admissible stance inputs. As in Poulakakis and Grizzle (2007b, 2009b), the leg force u_1 is modeled as a spring in parallel with a prismatic force source. During flight, on the other hand, the assumption of a massless leg implies that the ASLIP follows a ballistic motion. Furthermore, the leg attains its desired configuration $\alpha_f = (l^{td}, \varphi^{td}) \in A_f$ in anticipation to touchdown kinematically, just like the SLIP; see also Fig. 4.7.8. The inherently hybrid nature of the dynamics of the ASLIP can then be represented by a system of the form (4.7.8) as

$$\mathcal{HC}_{\mathrm{ASLIP}} := (\Gamma, D, U, A, S, \Delta, FG) \tag{4.7.52}$$

where $U = \{U_s, \emptyset\}$ and $A = \{\emptyset, A_f\}$ include the inputs available in continuous and in discrete time, respectively. The system (4.7.52) can be brought in the standard form of a system with impulse effects by integrating the flight phase dynamics until touchdown, thereby obtaining a map $\Delta : S_{e_s^f} \to D_s$ that takes the liftoff conditions $x_s^- \in S_{e_s^f}$ together with the desired configuration $\alpha_f \in A_f$ of the leg at touchdown to the initial conditions $x_s^+ \in D_s$ of the next stance phase. The details can be found in Poulakakis and Grizzle (2009b), and the resulting form

is

$$
\mathscr{HC}_{\text{ASLIP}} : \begin{cases} \dot{x}_{\text{s}} = f_{\text{s}}(x_{\text{s}}) + g_{\text{s}}(x_{\text{s}})u_{\text{s}}, & x_{\text{s}} \in D_{\text{s}} \setminus S_{e_{\text{s}}^{\text{f}}}, \\ x_{\text{s}}^{+} = \Delta\left(x_{\text{s}}^{-}, \alpha_{\text{f}}\right), & x_{\text{s}}^{-} \in S_{e_{\text{s}}^{\text{f}}}, \ \alpha_{\text{f}} \in A_{\text{f}}. \end{cases}
\tag{4.7.53}
$$

It should be emphasized that the initial condition $x_{\text{s}}^{+} \in D_{\text{s}}$ of the ensuing stance phase does not only depend on the exit condition $x_{\text{s}}^{-} \in S_{e_{\text{s}}^{\text{f}}}$ of the previous stance phase. It also depends on the parameter α_{f} that determines the configuration of the leg at touchdown, thereby strongly influencing the ensuing stance phase. Clearly, updating α_{f} in an event-based fashion provides a powerful control input.

Embedding the SLIP in the Dynamics of the ASLIP

As was mentioned in Chapter 3, a growing body of evidence in biomechanics indicates that, when running, diverse species tune their musculoskeletal system so that their center of mass bounces along as if it is following the dynamics of a SLIP (Holmes et al., 2006). In the light of this evidence, the SLIP is construed as a canonical model of running, and can be used as a *behavioral control target* for legged robots or robot models. In what follows, we describe a feedback control law that organizes the ASLIP so that its closed-loop dynamics is governed by the dynamics of a variant of the SLIP; namely, the *energy-stabilized SLIP (ES-SLIP)* shown in Fig. 4.7.9. The ES-SLIP is a modification of the standard SLIP that admits exponentially stable[7] hopping motions (Poulakakis and Grizzle, 2007a, 2009b). The dynamics of the ES-SLIP in closed loop with an exponentially stabilizing feedback controller – see Poulakakis and Grizzle (2007a, 2009b) for details – can be written as

$$
\mathscr{H}_{\text{ES-SLIP}} : \begin{cases} \dot{z} = f_{\text{z}}(z), & z \notin S_{\text{z}}, \\ z^{+} = \Delta_{\text{z}}\left(z^{-}\right), & z^{-} \in S_{\text{z}}, \end{cases}
\tag{4.7.54}
$$

where S_{z} corresponds to the stance-to-flight switching surface and the rest of the components of (4.7.54) are defined in a similar fashion to those in (4.7.53); see Poulakakis and Grizzle (2007a, 2009b) for details.

The goal of the SLIP embedding controller is to render *any*[8] exponentially stable periodic running orbit of the ES-SLIP exponentially stable in the ASLIP. As Fig. 4.7.8B shows, control action is distributed over continuous and discrete time as follows. The continuous-time feedback law $u_{\text{s}} = \Gamma_{\text{c}}(x_{\text{s}})$ is employed

---

7. The standard SLIP is energy conservative and thus it cannot reject perturbations that shift the total energy of the system.
8. Provided, of course, that the physical constraints associated with ground reaction forces and actuator limitations are respected.
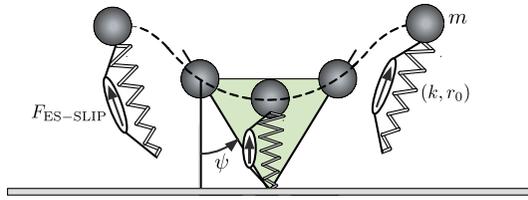
FIGURE 4.7.9 The energy-stabilized SLIP (ES-SLIP), a variant of the standard SLIP which features a prismatic leg actuator in parallel with the spring. The leg actuator can develop nonconservative forces capable or rejecting perturbations that alter the total energy of the system.

during stance with the purpose of (i) creating an *invariant* and *attractive* submanifold Z_s embedded in the stance state space, D_s, and (ii) rendering the restriction of the closed-loop stance dynamics of the ASLIP on Z_s diffeomorphic to the ES-SLIP stance dynamics; formally, $\left(f_s(x_s) + g_s(x_s)\Gamma_c(x_s) \right)\Big|_{Z_s} \cong f_z(z)$. The discrete-time feedback law $\alpha_f = \Gamma_d(x_s^-)$ is employed at transitions from stance to flight with the purpose of updating the leg configuration α_f at touchdown so that (i) Z_s is *hybrid invariant*, i.e., invariant under the closed-loop transition map $\Delta_{cl}(x_s^-) := \Delta\left(x_s^-, \Gamma_d(x_s^-)\right)$ as defined in Section 4.7.5.1, and (ii) the restricted closed-loop reset map $\Delta_{cl}|_{Z_s}$ of the ASLIP is equivalent to the transition map Δ_z of the ES-SLIP; formally, $\Delta_{cl}|_{Z_s} \cong \Delta_z$. More details on the design of the feedback laws (Γ_c, Γ_d) can be found in Poulakakis and Grizzle (2007a, 2009b).

Implications to the Control of Robots

The approach described above essentially combines the practical advantages of compliant reductive models typically used to intuitively tune empirical controllers – the SLIP is a classical example – with the analytical tractability offered by constructive feedback synthesis methods. But, how can we leverage these feedback constructions to introduce a general control synthesis framework for compliant legged robots?

Clearly, the direct implementation of the SLIP embedding controller in legged robots like Thumper and MABEL depicted in Fig. 4.7.1 is far from being a straightforward task. The primary reason is that the ASLIP is based on a number of simplifying assumptions that do not faithfully capture the structural and morphological characteristics of these robots. More specifically, the assumption of a massless leg together with the requirement that the HZD is equivalent – in a strict mathematical sense – to the SLIP severely limit the applicability of the SLIP embedding controller to Thumper and MABEL. Yet, the following lessons learnt from the SLIP embedding controller are important: (i) the HZD is an explicitly compliant system, possessing more than one degrees of free-

dom, thus capturing not only the progression of the leg's sweep angle – as in the classical HZD method (Westervelt et al., 2003) – but also leg compression and decompression in a way that respects compliance; (ii) control authority is available *within* the zero dynamics – which is now *controlled* – thus allowing the development of additional feedback action to realize compliance "shaping" and active force control for greater flexibility; and (iii) although keeping the torso at a constant angle – which is in fact a necessary condition for embedding the SLIP in the ASLIP (Poulakakis, 2010) – is restricting, commanding zero pitch velocity during the late stage of the stance phase ensures that the angular momentum associated with the torso is small when the system switches to flight so that excessive pitching during flight is eliminated. These three considerations, which underlie the SLIP embedding controller, can be encoded in a set of suitably parametrized outputs of the form (4.7.9) and enforced on the dynamics of Thumper and MABEL though feedback linearization as in Section 4.7.4.1; see Poulakakis (2008, Chapter VI) and Poulakakis and Grizzle (2009a) for the development of the method. Skipping details, we only note that, similar to the block diagram of Fig. 4.7.8B, the continuous-time control action introduces a set of parameters which are updated in discrete time using event-based feedback. A refined version of this method was implemented in Sreenath et al. (2011) to generate experimentally dynamically stable, fast and efficient walking motions on MABEL at top sustained speeds 1.5 m/s. Beyond walking, the notion of compliant hybrid zero dynamics is at the core of stabilizing running on MABEL (Sreenath et al., 2013). Running presents unique challenges due to the presence of substantial flight phases that limit control authority over the system. Addressing these challenges calls for active force control within the compliant HZD as detailed in Sreenath et al. (2013). This method resulted in MABEL running at an average speed of 1.95 m/s and a peak speed of 3.06 m/s.

4.7.7 SUMMARY

Most traditional legged locomotion control approaches heavily rely on heuristic methods which do not provide stability and performance guarantees, thus hindering the use of legged robots in real-life applications. The hybrid zero dynamics (HZD) method described in this chapter has been proposed as a general framework for the synthesis of feedback control laws that induce provably stable, fast, and reliable walking and running motions in legged robots. At the core of the method is the idea of encoding desired locomotion behaviors via a set of suitably parametrized virtual constraints, which effectively coordinate the higher-dimensional robot plant into a lower-dimensional hybrid subsystem – namely, the HZD – that governs the robot's locomotion behavior. This chapter briefly discussed the main concepts as well as key implementation aspects

underlying the applications of the method, pointing to the relevant literature for detailed accounts. Beyond its theoretical value, perhaps the most impressive feature of the HZD method is its versatile nature. This feature supports implementation on robots with different structural and morphological characteristics, ranging from rigid walking to compliant running bipeds and to prostheses.

REFERENCES

Ackermann, M., 2007. Dynamics and Energetics of Walking with Prostheses. PhD thesis. University of Stuttgart.

Alexander, R., 1988. Elastic Mechanisms in Animal Movement. Cambridge University Press, Cambridge.

Alexander, R., 1990. Three uses for springs in legged locomotion. Int. J. Robot. Res. 9, 53–61.

Ames, A.D., 2012. First steps toward automatically generating bipedal robotic walking from human data. In: Robot Motion and Control 2011. Springer, London, pp. 89–116.

Ames, A.D., 2014. Human-inspired control of bipedal walking robots. IEEE Trans. Autom. Control 59, 1115–1130.

Ames, A.D., Cousineau, E.A., Powell, M.J., 2012a. Dynamically stable bipedal robotic walking with NAO via human-inspired hybrid zero dynamics. In: Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control. ACM, pp. 135–144.

Ames, A.D., Galloway, K., Grizzle, J.W., 2012b. Control Lyapunov functions and hybrid zero dynamics. In: Proc. 51st IEEE Conf. Decision and Control.

Ames, A.D., Galloway, K., Sreenath, K., Grizzle, J.W., 2014a. Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics. IEEE Trans. Autom. Control 59, 876–891.

Ames, A.D., Grizzle, J.W., Tabuada, P., 2014b. Control barrier function based quadratic programs with application to adaptive cruise control. In: IEEE 53rd Annual Conference on Decision and Control (CDC). IEEE, pp. 6271–6278.

Ames, A.D., Holley, J., 2014. Quadratic program based nonlinear embedded control of series elastic actuators. In: IEEE 53rd Annual Conference on Decision and Control (CDC). IEEE, pp. 6291–6298.

Ames, A.D., Powell, M., 2013. Towards the unification of locomotion and manipulation through control Lyapunov functions and quadratic programs. In: Control of Cyber-Physical Systems. Springer International Publishing, pp. 219–240.

Ames, A.D., Vasudevan, R., Bajcsy, R., 2011. Human-data based cost of bipedal robotic walking. In: Proceedings of the 14th International Conference on Hybrid Systems: Computation and Control. ACM, pp. 153–162.

Ames, A.D., Xu, X., Grizzle, J.W., Tabuada, P., 2016. Control barrier function based quadratic programs with application to automotive safety systems. arXiv:1609.06408v1.

Borrmann, U., Wang, L., Ames, A.D., Egerstedt, M., 2015. Control barrier certificates for safe swarm behavior. IFAC-PapersOnLine 48, 68–73.

Burden, S., Revzen, S., Sastry, S.S., 2011. Dimension reduction near periodic orbits of hybrid systems. In: 2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC). IEEE, pp. 6116–6121.

Chevallereau, C., Abba, G., Aoustin, Y., Plestan, F., Westervelt, E.R., Canudas, C., Grizzle, J.W., 2003. RABBIT: a testbed for advanced control theory. Control Syst. Mag. 23, 57–79.

Chevallereau, C., Westervelt, E.R., Grizzle, J.W., 2005. Asymptotically stable running for a five-link, four-actuator, planar bipedal robot. Int. J. Robot. Res. 24, 431–464.

Collins, S.H., Ruina, A., 2005. A bipedal walking robot with efficient and human-like gait. In: Proceeding of 2005 International Conference on Robotics and Automation. Barcelona, Spain.

Cousineau, E., Ames, A.D., 2015. Realizing underactuated bipedal walking with torque controllers via the ideal model resolved motion method. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 5747–5753.

Freeman, R.A., Kokotović, P.V., 1996. Robust Nonlinear Control Design. Birkhäuser.

Galloway, K., Sreenath, K., Ames, A.D., Grizzle, J.W., 2015. IEEE Access 3, 323–332.

Goebel, R., Sanfelice, R., Teel, A., 2009. Hybrid dynamical systems. IEEE Control Syst. Mag. 29, 28–93.

Goldstein, H., Poole, C., Safko, J., 2002. Classical Mechanics, third edition. Addison-Wesley, San Francisco.

Goswami, A., 1999a. Foot rotation indicator (FRI) point: a new gait planning tool to evaluate postural stability of biped robots. In: International Conference on Robotics and Automation. IEEE, pp. 47–52.

Goswami, A., 1999b. Postural stability of biped robots and the foot-rotation indicator (FRI) point. Int. J. Robot. Res. 18, 523–533.

Gregg, R.D., Tilton, A.K., Candido, S., Bretl, T., Spong, M.W., 2012. Control and planning of 3-D dynamic walking with asymptotically stable gait primitives. IEEE Trans. Robot. 28, 1415–1423.

Grizzle, J.W., Abba, G., Plestan, F., 2001. Asymptotically stable walking for biped robots: analysis via systems with impulse effects. IEEE Trans. Autom. Control 46, 51–64.

Grizzle, J.W., Chevallereau, C., Ames, A.D., Sinnet, R.W., 2010. 3D bipedal robotic walking: models, feedback control, and open problems. In: IFAC Proceedings, vol. 43, pp. 505–532.

Grizzle, J.W., Chevallereau, C., Sinnet, R.W., Ames, A.D., 2014. Models, feedback control, and open problems of 3D bipedal robotic walking. Automatica 50, 1955–1988.

Grizzle, J.W., Hurst, J., Morris, B., Park, H.W., Sreenath, K., 2009. MABEL, a new robotic bipedal walker and runner. In: American Control Conference, 2009. ACC'09. IEEE, pp. 2030–2036.

Haddad, W.M., Chellaboina, V.S., Nersesov, S.G., 2006. Impulsive and Hybrid Dynamical Systems: Stability, Dissipativity, and Control. Princeton University Press, Princeton, NJ.

Hereid, A., Cousineau, E.A., Hubicki, C.M., Ames, A.D., 2016. 3D dynamic walking with underactuated humanoid robots: a direct collocation framework for optimizing hybrid zero dynamics. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 1447–1454.

Hereid, A., Kolathaya, S., Jones, M.S., Van Why, J., Hurst, J.W., Ames, A.D., 2014. Dynamic multi-domain bipedal walking with ATRIAS through slip based human-inspired control. In: Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control. ACM, pp. 263–272.

Holmes, P., Full, R.J., Koditschek, D.E., Guckenheimer, J., 2006. The dynamics of legged locomotion: models, analyses, and challenges. SIAM Rev. 48, 207–304.

Hubicki, C., Grimes, J., Jones, M., Renjewski, D., Spröwitz, A., Abate, A., Hurst, J., 2016. ATRIAS: design and validation of a tether-free 3D-capable spring-mass bipedal robot. Int. J. Robot. Res. 35, 1497–1521.

Huihua, Z., Yadukumar, S.N., Ames, A.D., 2012. Bipedal robotic running with partial hybrid zero dynamics and human-inspired optimization. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 1821–1827.

Hurst, J.W., 2008. The Role and Implementation of Compliance in Legged Locomotion. PhD thesis. The Robotics Institute, Carnegie Mellon University.

Hurst, J.W., Chestnutt, J.E., Rizzi, A.A., 2007. Design and philosophy of the BiMASC, a higly dynamic biped. In: Proceedings of the IEEE International Conference of Robotics and Automation. Roma, Italy, pp. 1863–1868.

Hurst, J.W., Rizzi, A., 2008. Series compliance for an efficient running gait. IEEE Robot. Autom. Mag. 15, 42–51.

Isidori, A., 1995. Nonlinear Control Systems, third edition. Springer-Verlag, Berlin.

Kajita, S., Kanehiro, F., Kaneko, K., Fujiwara, K., Harada, K., Yokoi, K., Hirukawa, H., 2006. Biped walking pattern generator allowing auxiliary ZMP control. In: Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 2993–2999.

Kajita, S., Morisawa, M., Harada, K., Kaneko, K., Kanehiro, F., Fujiwara, K., Hirukawa, H., 2003. Biped walking pattern generation by using preview control of zero-moment point. In: IEEE International Conference on Robotics and Automation. Taipei, Taiwan, vol. 2, pp. 1620–1626.

Lamperski, A., Ames, A.D., 2013. Lyapunov theory for Zeno stability. IEEE Trans. Autom. Control 58, 100–112.

Lygeros, J., Johansson, K.H., Simic, S., Zhang, J., Sastry, S., 2003. Dynamical properties of hybrid automata. IEEE Trans. Autom. Control 48, 2–17.

Ma, W., Kolathaya, S., Ambrose, E., Hubicki, C., Ames, A.D., 2017. Bipedal robotic running with DURUS-2D: bridging the gap between theory and experiment. In: Proceedings of the 20th ACM International Conference on Hybrid Systems: Computation and Control. ACM.

Maus, H.-M., Lipfert, S., Gross, M., Rummel, J., Seyfarth, A., 2010. Upright human gait did not provide a major mechanical challenge for our ancestors. Nat. Commun. 1.

McMahon, T.A., Cheng, G.C., 1990. The mechanics of running: how does stiffness couple with speed? J. Biomech. 23, 65–78. Suppl. 1.

Mehra, A., Ma, W.-L., Berg, F., Tabuada, P., Grizzle, J.W., Ames, A.D., 2015. Adaptive cruise control: experimental validation of advanced controllers on scale-model cars. In: American Control Conference (ACC). IEEE, pp. 1411–1418.

Morris, B., 2008. Stabilizing Highly Dynamic Locomotion in Planar Bipedal Robots with Dimension Reducing Control. Ph.D. thesis. University of Michigan.

Morris, B., Grizzle, J.W., 2009. Hybrid invariant manifolds in systems with impulse effects with application to periodic locomotion in bipedal robots. IEEE Trans. Autom. Control 54, 1751–1764.

Morris, B., Powell, M.J., Ames, A.D., 2013. Sufficient conditions for the lipschitz continuity of QP-based multi-objective control of humanoid robots. In: 2013 IEEE 52nd Annual Conference on Decision and Control (CDC). IEEE, pp. 2920–2926.

Morris, B., Westervelt, E.R., Chevallereau, C., Buche, G., Grizzle, J.W., 2006. Achieving bipedal running with RABBIT: six steps toward infinity. In: Mombaur, K., Dheil, M. (Eds.), Fast Motions Symposium on Biomechanics and Robotics. In: Lecture Notes in Control and Information Sciences. Springer-Verlag, Heidelberg, Germany, pp. 277–297.

Morris, B.J., Powell, M.J., Ames, A.D., 2015a. Continuity and smoothness properties of nonlinear optimization-based feedback controllers. In: 2015 IEEE 54th Annual Conference on Decision and Control (CDC). IEEE, pp. 151–158.

Motahar, M.S., Veer, S., Huang, J., Poulakakis, I., 2015b. Integrating dynamic walking and arm impedance control for cooperative transportation. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. Hamburg, Germany, pp. 1004–1010.

Motahar, M.S., Veer, S., Poulakakis, I., 2016. Composing limit cycles for motion planning of 3D bipedal walkers. In: Proceedings of IEEE Conference on Decision and Control, pp. 6368–6374.

Murray, R.M., Li, Z., Sastry, S.S., 1994. A Mathematical Introduction to Robotic Manipulation. CRC Press, Boca Raton.

Nguyen, Q., Hereid, A., Grizzle, J.W., Ames, A.D., Sreenath, K., 2016. 3D dynamic walking on stepping stones with control barrier functions. In: 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, pp. 827–834.

Nguyen, Q., Sreenath, K., 2015. Safety-critical control for dynamical bipedal walking with precise footstep placement. IFAC-PapersOnLine 48, 147–154.

Nguyen, Q., Sreenath, K., 2016. Optimal robust control for constrained nonlinear hybrid systems with application to bipedal locomotion. In: American Control Conference (ACC). IEEE, pp. 4807–4813.

Park, H.-W., Ramezani, A., Grizzle, J.W., 2013. A finite-state machine for accommodating unexpected large ground-height variations in bipedal robot walking. IEEE Trans. Robot. 29, 331–345.

Poulakakis, I., 2008. Stabilizing Monopedal Robot Running: Reduction-by-Feedback and Compliant Hybrid Zero Dynamics. PhD thesis. Department of Electrical Engineering and Computer Science, University of Michigan.

Poulakakis, I., 2010. Spring loaded inverted pendulum embedding: extensions toward the control of compliant running robots. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 5219–5224.

Poulakakis, I., Grizzle, J., 2007a. Formal embedding of the Spring Loaded Inverted Pendulum in an Asymmetric hopper. In: Proceedings of the European Control Conference. Kos, Greece.

Poulakakis, I., Grizzle, J.W., 2007b. Monopedal running control: SLIP embedding and virtual constraint controllers. In: Proceedings of the IEEE/RSJ International Conference of Intelligent Robots and Systems. San Diego, USA, pp. 323–330.

Poulakakis, I., Grizzle, J.W., 2009a. Modeling and control of the monopedal running robot thumper. In: Proceedings of the IEEE International Conference on Robotics and Automation. Kobe, Japan, pp. 3327–3334.

Poulakakis, I., Grizzle, J.W., 2009b. The spring loaded inverted pendulum as the hybrid zero dynamics of an asymmetric hopper. IEEE Trans. Autom. Control 54, 1779–1793.

Powell, M.J., Hereid, A., Ames, A.D., 2013. Speed regulation in 3D robotic walking through motion transitions between human-inspired partial hybrid zero dynamics. In: 2013 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 4803–4810.

Powell, M.J., Zhao, H., Ames, A.D., 2012. Motion primitives for human-inspired bipedal robotic locomotion: walking and stair climbing. In: 2012 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 543–549.

Raibert, M.H., 1986. Legged Robots that Balance. MIT Press, Cambridge, MA.

Reher, J., Cousineau, E.A., Hereid, A., Hubicki, C.M., Ames, A.D., 2016a. Realizing dynamic and efficient bipedal locomotion on the humanoid robot DURUS. In: 2016 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 1794–1801.

Reher, J.P., Hereid, A., Kolathaya, S., Hubicki, C.M., Ames, A.D., 2016b. Algorithmic foundations of realizing multi-contact locomotion on the humanoid robot DURUS. In: Twelfth International Workshop on Algorithmic Foundations on Robotics.

Sastry, S., 1999. Nonlinear Systems: Analysis, Stability and Control. Springer-Verlag.

Sinnet, R.W., Jiang, S., Ames, A.D., 2014. A human-inspired framework for bipedal robotic walking design. Int. J. Biomechatronics Biomedical Robot. 3, 20–41.

Sinnet, R.W., Powell, M.J., Shah, R.P., Ames, A.D., 2011a. A human-inspired hybrid control approach to bipedal robotic walking. IFAC Proceedings 44, 6904–6911.

Sinnet, R.W., Zhao, H., Ames, A.D., 2011b. Simulating prosthetic devices with human-inspired hybrid control. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 1723–1730.

Spong, M.W., Hutchinson, S., Vidyasagar, M., 2006. Robot Modeling and Control. John Wiley & Sons, New York.

Sreenath, K., 2011. Feedback Control of a Bipedal Walker and Runner with Compliance. PhD thesis. University of Michigan.

Sreenath, K., Park, H.-W., Poulakakis, I., Grizzle, J., 2011. A compliant hybrid zero dynamics controller for stable, efficient and fast bipedal walking on MABEL. Int. J. Robot. Res. 30, 1170–1193.

Sreenath, K., Park, H.-W., Poulakakis, I., Grizzle, J., 2013. Embedding active force control within the compliant hybrid zero dynamics to achieve stable, fast running on MABEL. Int. J. Robot. Res. 32.

van der Schaft, A., Schumacher, H., 2000. An Introduction to Hybrid Dynamical Systems. Lecture Notes in Control and Information Sciences, vol. 251. Springer-Verlag.

Vasudevan, R., Ames, A., Bajcsy, R., 2011. Using persistent homology to determine a human-data based cost for bipedal walking. In: 18th IFAC World Congress. Milano, Italy.

Vasudevan, R., Ames, A., Bajcsy, R., 2013. Persistent homology for automatic determination of human-data based cost of bipedal walking. Nonlinear Anal. Hybrid Syst. 7, 101–115.

Veer, S., Motahar, M.S., Poulakakis, I., 2015. On the adaptation of dynamic walking to persistent external forcing using hybrid zero dynamics control. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. Hamburg, Germany, pp. 997–1003.

Veer, S., Motahar, M.S., Poulakakis, I., 2016. Local input-to-state stability of dynamic walking under persistent external excitation using hybrid zero dynamics. In: Proceedings of the American Control Conference, pp. 4801–4806.

Veer, S., Motahar, M.S., Poulakakis, I., 2017. Almost driftless navigation of 3D limit-cycle walking bipeds. In: Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems.

Vukobratović, M., Borovac, B., Potkonjak, V., 2006. ZMP: a review of some basic misunderstandings. Int. J. Humanoid Robot. 3, 153–175.

Wang, L., Ames, A., Egerstedt, M., 2016a. Safety barrier certificates for heterogeneous multi-robot systems. In: American Control Conference (ACC). IEEE, pp. 5213–5218.

Wang, L., Ames, A.D., Egerstedt, M., 2016b. Multi-objective compositions for collision-free connectivity maintenance in teams of mobile robots. In: 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, pp. 2659–2664.

Wendel, E., Ames, A.D., 2012. Rank deficiency and superstability of hybrid systems. Nonlinear Anal. Hybrid Syst. 6, 787–805.

Wendel, E.D., Ames, A.D., 2010. Rank properties of Poincaré maps for hybrid systems with applications to bipedal walking. In: Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control. ACM, pp. 151–160.

Westervelt, E.R., 2003. Toward a Coherent Framework for the Control of Planar Biped Locomotion. PhD thesis. University of Michigan.

Westervelt, E.R., Buche, G., Grizzle, J.W., 2004. Experimental validation of a framework for the design of controllers that induce stable walking in planar bipeds. Int. J. Robot. Res. 23, 559–582.

Westervelt, E.R., Grizzle, J.W., Chevallereau, C., Choi, J.H., Morris, B., 2007. Feedback Control of Dynamic Bipedal Robot Locomotion. Taylor & Francis/CRC Press.

Westervelt, E.R., Grizzle, J.W., Koditschek, D.E., 2003. Hybrid zero dynamics of planar biped walkers. IEEE Trans. Autom. Control 48, 42–56.

Xu, X., Tabuada, P., Grizzle, J.W., Ames, A.D., 2015. Robustness of control barrier functions for safety critical control. IFAC-PapersOnLine 48, 54–61.

Zhao, H., Ambrose, E., Ames, A.D., 2017. Preliminary results on energy efficient 3D prosthetic walking with a powered compliant transfemoral prosthesis. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE.

Zhao, H., Hereid, A., Ambrose, E., Ames, A.D., 2016a. 3D multi-contact gait design for prostheses: hybrid system models, virtual constraints and two-step direct collocation. In: 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, pp. 3668–3674.

Zhao, H., Hereid, A., Ma, W.-l., Ames, A.D., 2015a. Multi-contact bipedal robotic locomotion. Robotica, 1–35.

Zhao, H., Horn, J., Reher, J., Paredes, V., Ames, A.D., 2011. First steps toward translating robotic walking to prostheses: a nonlinear optimization based control approach. Auton. Robots, 1–18.

Zhao, H., Horn, J., Reher, J., Paredes, V., Ames, A.D., 2016b. Multicontact locomotion on transfemoral prostheses via hybrid system models and optimization-based control. IEEE Trans. Autom. Sci. Eng. 13, 502–513.

Zhao, H., Kolathaya, S., Ames, A.D., 2014a. Quadratic programming and impedance control for transfemoral prosthesis. In: 2014 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 1341–1347.

Zhao, H., Powell, M., Ames, A., 2014b. Human-inspired motion primitives and transitions for bipedal robotic locomotion in diverse terrain. Optim. Control Appl. Methods 35, 730–755.

Zhao, H., Reher, J., Horn, J., Paredes, V., Ames, A.D., 2015b. Realization of nonlinear real-time optimization based controllers on self-contained transfemoral prosthesis. In: Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems. ACM, pp. 130–138.