Limit Cycle Control And Its Application To The Animation Of Balancing And Walking

Joseph Laszlo† Michiel van de Panne Eugene Fiume

Department of Computer Science and
†Department of Electrical and Computer Engineering
University of Toronto

ABSTRACT

Seemingly simple behaviors such as human walking are difficult to model because of their inherent instability. Kinematic animation techniques can freely ignore such intrinsically dynamic problems, but they therefore also miss modeling important motion characteristics. On the other hand, the effect of balancing can emerge in a physically-based animation, but it requires computing delicate control strategies. We propose an alternative method that adds closed-loop feedback to open-loop periodic motions. We then apply our technique to create robust walking gaits for a fully-dynamic 19 degree-of-freedom human model. Important global characteristics such as direction, speed and stride rate can be controlled by changing the open-loop behavior alone or through simple control parameters, while continuing to employ the same local stabilization technique. Among other features, our dynamic "human" walking character is thus able to follow desired paths specified by the animator.

Keywords: control, limit cycles, physically-based modeling, locomotion, human animation

1. INTRODUCTION

As with any modeling endeavor, Nature has much to tell us about modeling motion. When observing the running motion of a charging bull, before we take cover we can clearly see that motion is a product of both physics and muscular action. Physically based animation mimics Nature by modeling both.

While techniques for simulating the basic physics of motion are well known, less is known about how to provide the necessary control over the muscles or actuators in order to produce a desired motion [1]. By analogy, for most adults, walking is a seemingly effortless task. We know from watching toddlers that the apparent ease of walking, running, and maintaining balance is deceiving. It comes as little surprise, then, that walking has proven to be a difficult motion to model. Indeed the most successful approach has been literally to "watch Nature" by capturing motion data from real walks and mapping this data to computer generated characters. The work of [16, 5, 30, 35] is a good survey of recent techniques aimed at making the most flexible use of existing motion-capture and keyframe data. The advantages of motion capture are obvious: the immediate generation of realistic human motions. However, motion capture does not provide us with sufficient understanding to create more general walking motions, especially when conditions are unpredictable, when new motions need to be generated, or when dealing with non-human characters. Developing methods to control physical simulations can potentially provide us with a more general tool, and it is therefore the modeling paradigm we adopt.

{jflaszlo | van | elf}@dgp.utoronto.ca

We gratefully acknowledge the financial support of our research by ITRC (Ontario) and NSERC (Canada). The ongoing benevolence of Alias|Wavefront to our graphics lab is also appreciated.

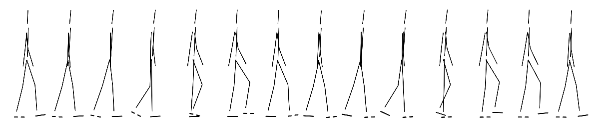

Figure 1. A 3-D dynamic walk using limit cycle control.

This paper proposes a solution for the control of periodic, unstable motions. The technique automates the addition of feedback to otherwise unstable motions, such as walking and running (see Figure 1). We also present an application of our technique by demonstrating a human model capable of a variety of walking styles, as well as a second, imaginary creature.

2. BACKGROUND

Many solutions to animation control problems have been proposed in the animation literature and elsewhere. In particular, there is a large body of work focussed on the important problem of locomotion control. What follows is a brief summary of some of these techniques, with an emphasis on algorithms for locomotion, including human walks and runs.

Procedural methods have been popular for generating motion [10], especially for human walking [2, 3, 4]. These methods directly generate walking motions by using a series of constraints, which are based on empirical data or on kinematic relationships. The work of [3] uses a mixed kinematic/dynamic model. A positive feature of such systems is that they give the animator direct control over useful gait characteristics, such as stride length, pelvic list, etc. The motions produced are typically parameterized in a way that is directly meaningful to an animator.

In physical approaches to animation, one must solve for the control actions that will, upon simulation, produce a desired motion. One approach has been to use a type of underconstrained inverse dynamics with automatic addition and removal of constraints which can be hand-tailored to yield stable 3D walks [29]. Several other approaches have treated control directly as a search problem, employing a particular choice of control representation and a choice of search algorithm, including genetic algorithms [24, 28], and simulated annealing [23, 12, 31]. The results indicate that such techniques are surprisingly adept at finding novel modes of locomotion. Unfortunately, it is less clear that such global search techniques are an efficient means of finding control strategies for motions requir-

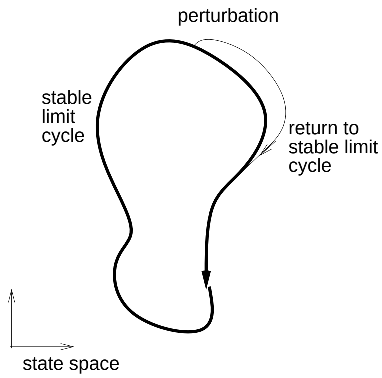

Figure 2. A passively-stable limit cycle.

ing fine control and feedback, such as human walking and running. One would think that to create this type of controller, we should be able to do better than to explore a large solution space at random. Toward this end, some methods, including ours, draw upon previous work in robotics and control.

The hopping and running control strategies presented in [14, 26] represent a more methodical approach to arriving at a control strategy. The basic idea stems from earlier robotics work [25] and is a powerful control technique for hopping and running motions, given that certain assumptions are fulfilled. The key is an elegant decomposition of the control problem which arises from assuming the legs are lightweight with respect to the body mass. This assumption can be overcome to some extent as shown in the work of [14]. A more analytical examination of the remarkable robustness of this and a class of related strategies can be found in [17].

Other specific controller designs have also met with success. The simulated cockroach in [22] is a good example. Strategies proposed for walking robots also have possible applications for animation [34, 9, 11, 23]. Demonstrations of passively-powered walking down a slight incline show that active walking may only require small amounts of energy [20, 21].

The point of departure for our work is the concept of *periodic limit cycles*. Consider for a moment the motion of a typical mechanical toy, which drives its joints in a repetitive, periodic fashion and is oblivious to its environment. This type of open-loop control is sufficient for many types of animated motion, as has been pointed out in [32]. However, it is insufficient as a control mechanism for unstable dynamic motions such as walking and running. The control technique we propose provides a general method of turning unstable open-loop motions into stable closed-loop motions. The basis of the technique is a process that perturbs the open-loop control actions slightly in order to yield a desired stable, cyclic motion. A broad range of related work on limit cycles and periodic control can be found in the control systems literature [15, 13, 19, 6, 33]. While this paper provides all the essential information for implementing our technique, more details can be found in [18].

3. LIMIT CYCLE CONTROL

A mechanical toy owes its successful motion to a stable limit cycle that arises naturally from the interaction of the toy with its environment. Typically, the environment acts on the toy in such a way that any disturbance to the motion, such as a small external push, is rapidly damped. Figure 2 illustrates this concept using a phase diagram. The phase diagram is a projection of the path a motion traces through state-space over time.[1] Unfortunately, recalling our wobbly toddler, unstable motions such as walking cannot rely solely on passive damping of disturbances

[1] An object's state is the minimal set of parameters necessary to describe the position and velocity of all points on the object.

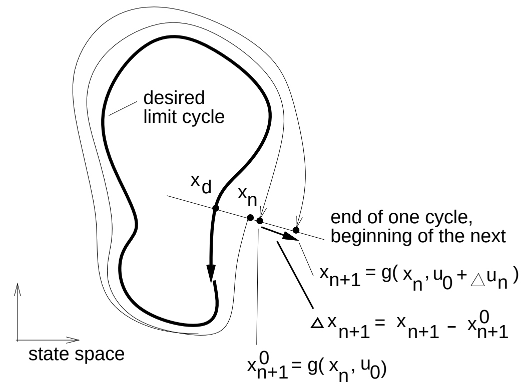

Figure 3. A limit cycle as a discrete dynamical system.

The goal of limit-cycle (LC) control is to actively drive motions that would otherwise be unstable back to a fixed limit cycle. The key point to implementing LC control is to avoid dealing with the complexities of non-linear dynamics by representing the motion using a well-behaved *discrete* dynamical system.

In general, the continuous equations of motion can be expressed by the non-linear differential equation

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$$

where $\mathbf{x}$ is the system state, $\dot{\mathbf{x}}$ its time derivative, and $\mathbf{u}$ is the control input. The discrete dynamical system that we shall deal with has the following iterative form:

$$\mathbf{x}_{n+1} = g(\mathbf{x}_n, \mathbf{u}_0 + \Delta\mathbf{u}_n)$$

where $\mathbf{x}_n$ represents the initial state of the n^{th} periodic limit cycle, $\mathbf{x}_{n+1}$ is the state on the next cycle, $\mathbf{u}_0$ is the periodic open-loop control, and $\Delta\mathbf{u}_n$ is an applied control perturbation. The purpose of $\Delta\mathbf{u}_n$ will be to drive the sequence $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_i$ to a desired value, $\mathbf{x}_d$. Figure 3 depicts an abstraction of the discrete dynamical system.

The advantage of dealing with the discrete dynamical system is that it is relatively smooth and therefore subject to linear approximation. There is no *guarantee* that the discrete system will be smooth, although we have in general found this to be the case for our experiments. For small control perturbations, $\Delta\mathbf{u}_n$, that will be applied on each cycle, the resulting change in state, $\Delta\mathbf{x}_{n+1}$, has a first order approximation of

$$\Delta\mathbf{x}_{n+1} = J\Delta\mathbf{u}_n$$

where J is a Jacobian relating the change in state after one cycle to the control perturbation. A first order approximation of the change of state makes it possible to calculate the control perturbation required to bring a system back onto a desired limit cycle, namely

$$\Delta\mathbf{u}_n = J^{-1}\Delta\mathbf{x}_{n+1}$$

where the desired $\Delta\mathbf{x}_{n+1}$ is calculated as $\Delta\mathbf{x}_{n+1} = \mathbf{x}_d - \mathbf{x}^0_{n+1}$, $\mathbf{x}_d$ being the desired state and $\mathbf{x}^0_{n+1}$ being the state achieved when the nominal control, $\mathbf{u}_0$, is applied.

3.1 Linearity of Control Perturbations

Our evidence for the above linear approximation is empirical, although this type of linearization has been justified more rigorously in the application of control theory to certain types of dynamical systems [8, 13]. Figure 4 provides some experimental evidence that a linear model is a sufficient for modeling the effect of control perturbations over a complete cycle of motion, despite the occurrence of discontinuous events such as foot-falls during the cycle itself.

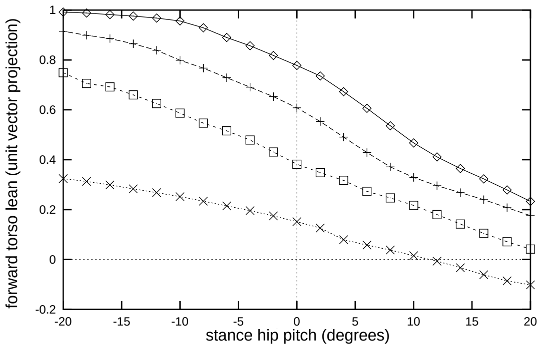

Figure 4. Linearity of perturbation control.

The vertical axis on this graph is a variable we would like to be able to control, namely some aspect of the system state. In this particular example, it is a measure of the forward pitch of the body of a simulated human model performing a walking movement, as measured at the end of a single step. The horizontal axis gives the magnitude of an applied control perturbation, in this case an alteration of the hip pitch angle during a particular part of the step. The graph thus tells us what type of forward pitch our simulated human body will have after having taken one step, using different variations of applied control. Each line shows an example of this relationship for a different step. It is evident in this figure that we can alter the forward pitch of the body as we desire and that a linear approximation is a reasonable model for the effect of a control perturbation.

It is important to note that the effects of control perturbations depend on the initial state, $\mathbf{x}_n$. This is clear from Figure 4, which shows that the relationship between control perturbation and change in state varies for different steps. What this means in practice is that for each cycle of the motion, we need to recompute the Jacobian J, which defines the relationship between applied control perturbations and the resulting changes in state after one cycle.

3.2 Regulation Variables

For many models, it is possible to make further simplifications and still effect proper LC control. It is typically neither necessary nor practical to work with a complete state vector in producing a controlled, stable limit cycle, as has thus far been implied. Instead, it can be sufficient to work with a small number of *regulation variables* (RVs). Ensuring that these regulation variables are controlled to follow a limit cycle is sufficient to stabilize the limit cycle for the motion as a whole. The use of RVs instead of the complete state vector could also be considered advantageous for animation, as it potentially leads to a less-constrained, freer motion. This is loosely related to the use of *reduced-order models* in control theory [8, 13]. Along with choosing the regulation variables, one must provide a choice of desired target values for these variables to take.

So as to make the notion of regulation variables more concrete, we introduce some of the possible choices that we know (through experiments) work well for a human walking model. A well chosen set of regulation variables should give a meaningful projection of the system state over a large range of possible states.

Figure 5 shows two possible choices for sets of regulation variables, each based upon the definition of a particular type of vector. The simplest of the two is the *up-vector*, so we begin by explaining this choice. The up-vector is a fixed vector of unit length, defined in the coordinate frame of the pelvis, which can be used to measure the forward lean and sideways tilt of the pelvis. The regulation variables are the forward and lateral components of the up-vector, i.e., the projection of the up-vector onto a horizontal ground plane. Controlling the values of these two scalar variables using LC

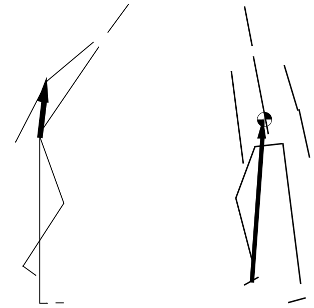

Figure 5. Regulation variables for use in walking. From left to right: (a) Up vector, (b) Swing-COM vector.

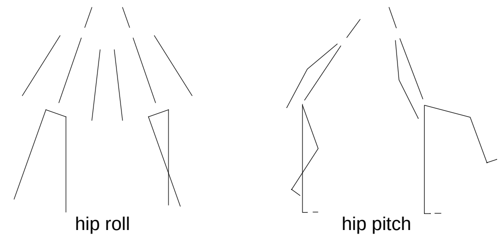

Figure 6. Control perturbations used for walking.

control is sufficient to yield a balanced, fully three-dimensional walking motion for our human model.

Other choices of regulation variable are also possible. The swing-center-of-mass (swing-COM) defines a vector from the swing foot to the center-of-mass, and thus measures where the center-of-mass will lie with respect to the future point of support. Once again, the two scalar components of the projection of the vector onto the ground plane form the set of regulation variables. In the presentation of the results, we primarily demonstrate the use of the up vector.

3.3 Control Perturbations

Given that it is often sufficient to work with a limited number of regulation variables, we need to determine the type and number of control perturbations to effect the necessary control. For two regulation variables, as is the case for our walking example, we shall require two appropriately-chosen control perturbations in order to yield a well-formed Jacobian J.

The control perturbations we work with for our human model are twofold. First, one can use changes to the stance-hip pitch and roll, effected over a particular portion of the walk cycle. Figure 6 illustrates the effects of these perturbations in an exaggerated fashion. The two figures on the left demonstrate stance-hip roll, while those on the right demonstrate stance-hip pitch. Second, the use of alterations to swing-hip pitch and roll can be similarly used as a suitable pair of control perturbations. In this case, the resulting control can be thought of as a type of foot-placement strategy. As with the choice of regulation variables, the choice of control perturbations is not unique, although they must be chosen to span the space of desired changes to the regulation variables

4. MODEL DESCRIPTIONS

We choose human walking as our primary example to illustrate the limit cycle control technique for two reasons. The first is that it is typical of motions for which the open-loop control actions are relatively easy to construct and can thus benefit immediately from LC control in order to “close the loop.” The second reason is that the control of a dynamic human walking model demonstrates the effec-

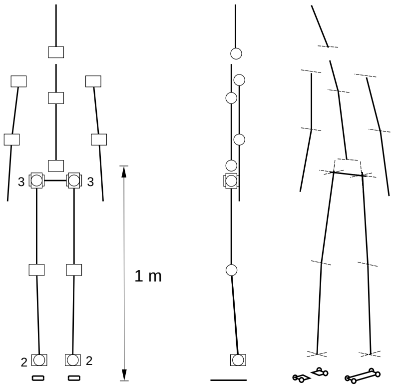

Figure 7. Construction of the human model. Joints are 1 DOF except where indicated.

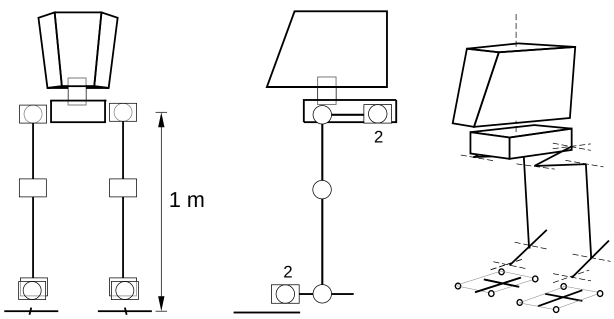

Figure 8. Construction of the robot model. Joints are 1 DOF except where indicated.

tiveness and scalability of this technique. The dynamic control of models with many degrees-of-freedom (DOF) of the type used in animation is problematic for many control techniques.

4.1 The Human Model

The physical model we use has a mass and inertia distribution comparable to those of a real human. The parameters are identical to those used in [36] and were originally obtained from [7]. The model has 13 joints and 19 DOF, as shown in Figure 7. The hip joints have three rotational DOF, while the ankle joints have two rotational DOF. All other joints have one DOF. The equations of motion are calculated and integrated using a commercially-available simulation package [27]. The ground is modelled using a penalty method. Stiff springs and dampers exert forces on a set of four points on the feet whenever they penetrate the ground. Each point is allowed to slip independently when the ratio of its applied horizontal and vertical component forces exceeds a user supplied threshold. The ground model thus uses no artificial constraints to hold the foot in place, which have in the past been used to simplify the simulation (and to some extent the control) of human motion.

4.2 The Robot Model

As a second test case, we consider the robot shown in Figure 8. This figure has 11 joints and 15 rotational DOF. The lateral base of support is much wider than that of the human model, yielding a very different type of motion. Mass and inertia parameters for this model are shown in Table 1.

5. APPLYING LC CONTROL

In the following sections we present the details of the limit cycle control algorithm as applied to our human walking model. The discussion first looks at the open-loop control before proceeding on to examine how the LC-control algorithm can be superimposed on it.

Link	Mass (kg)	Moment of Inertia (x, y, z kgm^2)		
body	14.72	0.47	0.26	0.40
head	20.9	1.11	0.92	0.72
upper leg	1.2	0.001	0.01	0.01
mid-leg	1.6	0.023	0.0013	0.023
lower leg	2.2	0.057	0.0018	0.057
ankle	1.4	0.0012	0.015	0.015
foot	2.52	0.0098	0.024	0.016

Table 1. Robot model mass and inertia parameters. The x, y and z axes are the forward, vertical and lateral axes respectively.

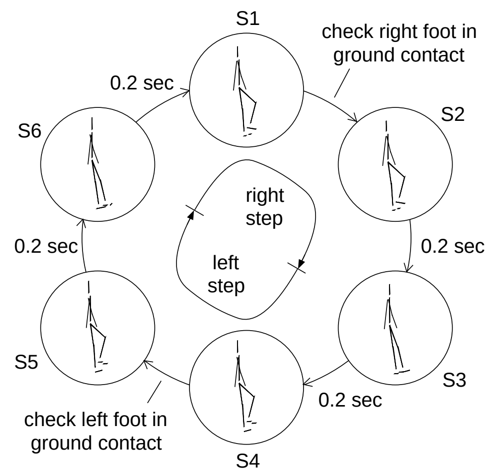

Figure 9. Finite state machine employed for walking.

5.1 Open-Loop Control

Finite state machines (FSMs) combined with proportional-derivative (PD) controllers are a common control mechanism in both physically-based animation [14, 22, 32] and robotics. The finite-state machine used as a basic controller for the walking motions is shown in Figure 9.

Each state in the FSM provides a fixed set of desired angles to the individual PD joint controllers. In our FSM, the desired angles change as a step function when proceeding from one state to the next. The PD controllers calculate a torque according to

$$\tau = k_v(\theta_d - \theta) - k_v\dot{\theta}$$

where θ_d is the desired joint angle, θ is the actual joint angle, $\dot{\theta}$ is the angular velocity of the joint and k_p and k_v are gain constants which serve to define the strength of the joint.

The state transitions in the finite state machine are time-based, with the exception of the transitions exiting states S1 and S4. These latter transitions are sensor-based and perform the simple job of ensuring that the proper stance foot is on the ground before completing the current step.

A basic open-loop motion can be constructed by defining the poses in states S2 and S3, where a pose consists of the set of desired joint angles to be used in a state. The pose for state S1 is identical to the pose for state S2. The poses for states S4, S5, and S6 are the same as the poses for states S1, S2, and S3, respectively, with the left and right sides exchanging roles. In typical operation, state S2 (S5) raises and advances the swing leg and state S3 (S6) straightens it in anticipation of ground contact. Normally, the foot contacts the ground some time after entering state S3 (S6) and the remaining time in the state is spent in double-stance phase[2]. Since the next stance foot is already on the ground, the transition out of state S4 (S1) occurs immediately after entering it, essentially skipping the

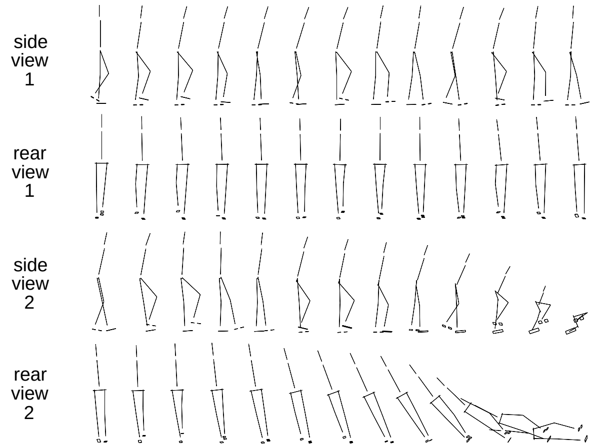

Figure 10. Walking with open-loop control (front/side views).

state. The cycle then repeats for the other leg. The sensor-based transitions serve only to make for a more robust motion. They effectively provide a way for the controlling FSM to remain synchronized with the actual motion and are typically only necessary during startup or when an FSM is dynamically altered to obtain a different motion.

Note that strictly speaking, the FSM does not provide true open-loop control since the desired joint angles are realized using local PD controllers. Nevertheless, the motion is open-loop in the sense that no system-wide feedback is used to drive it towards the desired trajectory. The result of using the "open-loop" FSM of Figure 9 for walking control is shown in Figure 10. It produces a motion which takes several steps and then falls over.

A certain amount of trial-and-error parameter tuning is required to produce open-loop motions which can be balanced successfully. Although tedious, this process is relatively straightforward. Tasks might include ensuring that toes do not stub the ground and that the basic motion can produce movement in the desired directions. Once a good open loop controller has been generated, it can be used to produce a wide variety of motions.

5.2 LC Control for Walking

The first step in implementing LC control is to choose a set of regulation variables and a set of control perturbations. We have experimented with two choices for the former, as shown in Figure 5, and two choices for the latter.

As indicated in Figure 3, the final nature of the limit cycle is defined by the target state to be achieved at the end of the current cycle. It should be noted that LC control is not successful for all choices of target values. Target values should be similar to the values that can be observed during the first few steps of the open-loop motion. This ensures that the generated limit cycle is close to the unstable open-loop limit cycle, thereby limiting the LC control to having to perform relatively small control corrections.

The power of LC control lies in being able to predict the change in values of the regulation variables with respect to the applied control perturbations. Using a linear model for this allows us to easily predict the required perturbation.

The chosen pair of control perturbations for human walking, namely alterations to the desired hip pitch and roll angles, were designed to allow for more-or-less independent control of each of the regulation variables. The hip roll is effectively used to provide balance in the coronal (side-to-side) plane, while the hip pitch provides

[2] The double-stance phase is the part of the walking cycle during which both feet are in ground contact.

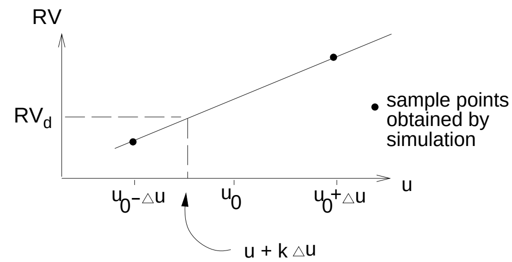

Figure 11. Interpolating for a desired control perturbation.

Figure 12. A falling motion illustrating the torso servo.

balance in the sagittal (front-back) plane. In effect, this corresponds to only requiring the use of the diagonal elements of the Jacobian. We have not yet attempted to make use of the full Jacobian.

Figure 11 illustrates the linear interpolation scheme, which is applied twice, once for sagittal balance, and once for coronal balance.

Carrying out one actual step in a walking motion requires performing five simulations of the step, each slightly different from one another. The first four simulations are used to capture the necessary data to construct a simple model of how control perturbations will affect the state of the body at the end of the step. This model is then used to estimate the necessary control perturbations to achieve the desired target state, and hence the desired limit cycle. The fifth simulation is required to produce the final balanced motion for the current step before proceeding on to the next. The blind reconstruction of the RV-perturbation model each step in this fashion results in a five-fold increase in the required computation time compared to normal forward dynamic simulation. If the local perturbation model can itself be predicted, true closed-loop control can be achieved. We are optimistic that this is possible.

The robot model uses very similar choices to the human model's to achieve a running motion. Stance hip variations provide the control perturbations and the chosen RVs are projections of an up vector attached to the creature's head. The primary difference between the control for the two models is in the open loop FSM. Aside from differences in the particular poses, the robot's transition times are smaller than the human FSM (by about half) since the creature's wide stance makes it difficult to remain on one foot for long.

5.3 Torso Servo

While the limit cycle control mechanism described thus far generates stable walks, the resulting motions exhibit a characteristic bobbing of the torso. This is an artifact of the simple open-loop motion chosen as a point of departure for our walking gait. We implement a simple vertical torso servo which not only smoothes the torso motion (if desired), but also demonstrates the robustness of the limit cycle control in continuing to provide effective balance. For LC control, stabilizing a system which already contains some feedback components such as torso servoing is no different than stabilizing an open-loop motion.

The torso servo consists of a PD controller applied to the one degree-of-freedom waist joint. The applied torque at this joint serves to force the torso to always remain upright with respect to the world coordinate frame. Note that the torso servoing does not prevent the biped from falling because the legs must still ultimately

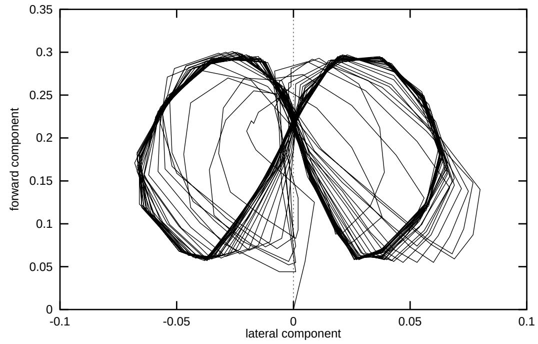

Figure 13. Regulation variable limit cycle (up-vector based).

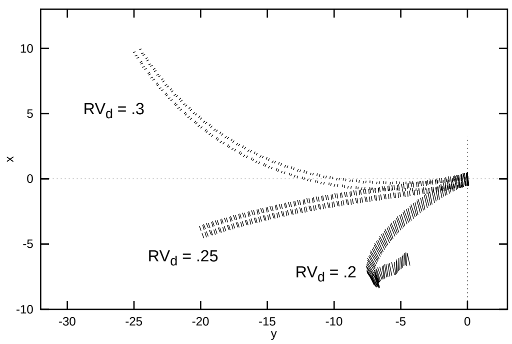

Figure 14. Dynamic walks using the up-vector

provide for balanced support, as illustrated in Figure 12.

6. RESULTS

Limit cycle control has been applied to obtain stable walking gaits for a 19-DOF human model of realistic proportions. As a second example, we control the running motions of a two-legged robot with a bird-like skeleton.

Figure 1 shows a sequence of frames from a typical dynamic walk, resulting from the application of limit cycle control to an open-loop walking motion, with torso servoing enabled.

An illustration of the typical limit cycle which is achieved is shown in Figure 13. The path indicates the continuous-time projection of the unit up vector onto the horizontal plane. This figure is the real analog of the earlier abstraction shown in Figure 3. For walking, the limit cycle consists of two roughly symmetric halves. This occurs because a single cycle consists of a left step and a right step, each forming half of the complete limit cycle for a stride. The desired values for the regulation variables lie at the center of this diagram, at $RV_{forward} = 0.25$, $RV_{lateral} = 0$. Perturbation control is applied on each step, thus forcing the limit cycle towards the desired point twice on each cycle. The startup phase of the motion is also evident from the figure, with the regulation variables eventually being driven onto a stable limit cycle.

While LC control is a general method of adding balance to a walk, it does not by itself ensure a straight walk. Figure 14 illustrates different paths taken for different target values (RV_d) of the regulation variable controlling the desired forward lean. The figure is a top view of the walking motion, showing only the position of the pelvis, enlarged in order to make the orientation of the body clearly visible during the walk. The correct scale for the walk is given by the axes, indicated in metres.

An example of an alternative choice for the set of regulation variables is the use of the swing-center-of-mass (swing-COM) vector. This also leads to stable walks, but not necessarily straight walks. Figure 15 shows how the path can vary as the target value

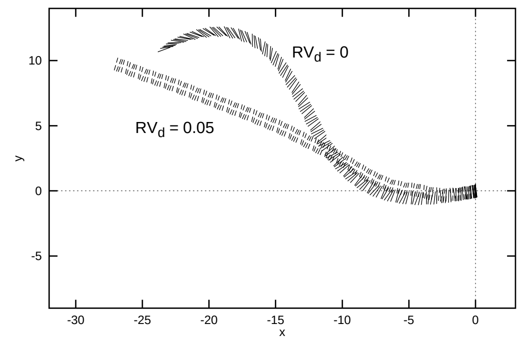

Figure 15. Dynamic walks using the COM-vector.

for the forward component of the vector takes on different values.

The quality of the final motion depends heavily on the open-loop motion on which it is based. The simple four pose FSM used to generate most of the human model walks is quite robust, but produces a motion more akin to a robot marching than to normal human walking. A more human-like walk has been generated by increasing the number of states and changing the timing of the swing leg motions. The use of motion captured data for tuning the open-loop gait for realism presents the possibility of further refining the motion while retaining both the guarantee of realistic motion and the flexible autonomy LC control can provide.

6.1 Additional Animator Control

Control over the speed and direction of a walking gait is of obvious necessity to an animator. Because the stabilization (i.e., balance) of the walk is automated, it is a relatively simple matter to provide the necessary hooks to control these parameters, as well as other possible stylistic variations.

6.2 Speed Control

The speed of the walking gait can be controlled in several ways. One technique we have applied is to alter the underlying open-loop motion to produce particular faster and slower velocity walks and to use interpolations of these *base controllers* to achieve intermediate speeds. Another successful approach makes use the fact that the forward speed is a function of the target values chosen for the regulation variables. As we choose limit cycles for which the pelvis and torso lean forward more, the speed of the gait increases. In this case, a simple form of velocity feedback is used to give consistent, stable steady-state velocities. The RV target value for each step, i, is:

$$RV_d^i = RV_d^* + K(v_d - v)$$

where v and v_d are the actual and desired speeds respectively, K is a proportionality constant, and RV_d^* is a bias term used to relate RV values to speed. Either the bias term or the desired velocity can be used to vary walking speed. Figure 16 illustrates the walking speed for a set of walks obtained using this approach. The startup phase can be recognized during the first several steps and the longest walk demonstrates both positive and negative acceleration phases. All of the motions begin from rest. Only relatively small changes in the final RV target values are required to achieve a reasonable range of speeds. The fastest walk in Figure 16 has $RV_d^i \approx 0.30$ while the slowest walk (nearly stationary) has $RV_d^i \approx 0.25$. These correspond to forward pelvis angles of approximately 17 and 14 degrees from vertical respectively. This technique allows the underlying open loop motion to remain fixed, but tends to reduce motion quality somewhat at higher speeds.

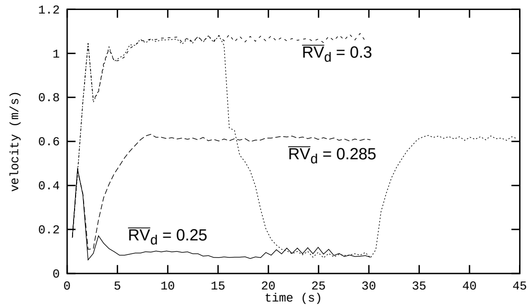

Figure 16. Speed control for a dynamic walk. $\overline{RV}_d$ = steady-state RV_d^i.

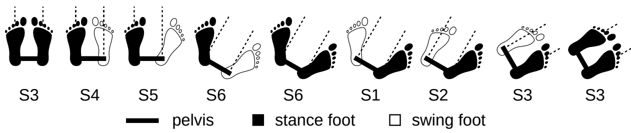

Figure 17. Hip rotations for turning during a walk.

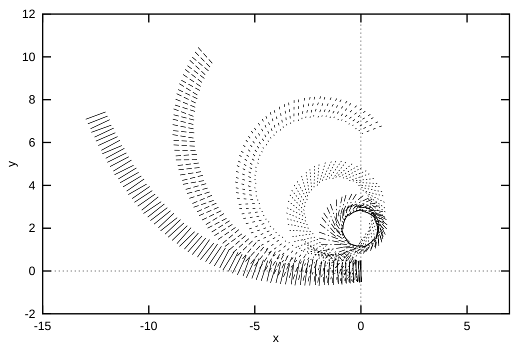

Figure 18. Turning motions for a dynamic walk

6.3 Direction Control

By changing the open-loop motion to include suitable hip twists, controllable turning motions can be achieved. Figure 17 shows a sequence of footprints which illustrate the effect of the hip joint rotations on the orientation of the pelvis for the case of a stationary walk. The figure indicates the hip rotations relative to the pelvis at key points in the walking cycle and the states of the basic FSM (Figure 9) in which they occur. The turning motion works best when torso servoing is applied, in that turns of tighter radius can be performed. Figure 18 shows turns obtained using scaled versions of the twisting motion and stabilized using LC control.

Once the turning radius can be controlled, it can be used to produce a path-following algorithm. Figure 19 shows a dynamic walk following a desired trajectory. The algorithm makes use of a *target point* on the trajectory and chooses a turning rate proportional to the current error in direction. When the target point has almost been reached, its position is updated to be further along the desired path.

6.4 Other Variations

Many other walking styles can be implemented by tailoring the open-loop control as desired. In many cases, transitions between different motions can be performed by simple linear interpolation of the underlying open-loop control over a period of a few steps. In some cases, a more gradual transition or more complex control of the desired RV values is necessary to avoid a fall. In either case, the basic LC control mechanism remains the same. Note that because a dynamic simulation is always being used, ground constraints and other physics constraints are always fulfilled, something that is not necessarily the case with direct kinematic interpolation of motion data, especially for large variations.

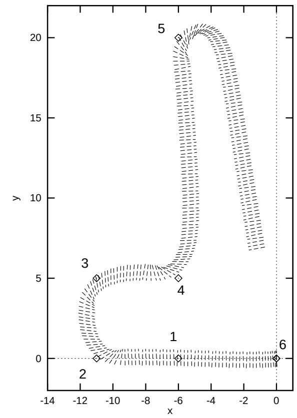

Figure 19. Path following using six target points.

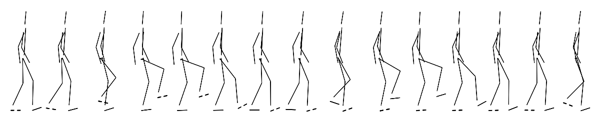

Figure 20. Stylistic variation on a walk.

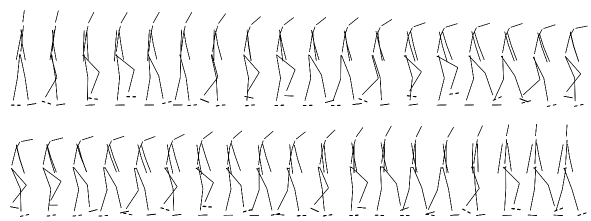

Figure 21. A ducking motion.

Figure 19 illustrates a stylistic variation on a walk with the knees being bent and lifted high on each step. Figure 21 shows a ducking motion obtained by transitioning into and out of a bent-over walk. Stride rate variations can be achieved by changing the duration of time-based state transitions in the open-loop FSM. We have also simulated walking motions into a strong wind, for which the automatic feedback provided by the LC control visibly alters the motion to lean into the wind. In addition to these variations, LC control has proven capable of balancing lateral and backward walking motions obtained using exactly the same open-loop FSM as the forward moving walks. When the human model's initial state has a sufficient backward or lateral velocity component, a balanced walk ensues in the direction of this initial nudge. The application of velocity control to effect transitions into and out of such motions has not yet been attempted but is expected to be relatively straightforward.

As a final example, Figure 22 shows our robot model which can run in a controlled fashion using LC control. This example serves to further illustrate the generality of LC control with respect to significant model and gait variations.

7. CONCLUSIONS

Physically-based animation is difficult because of the lack of general control techniques. Motions such as walking are known to be

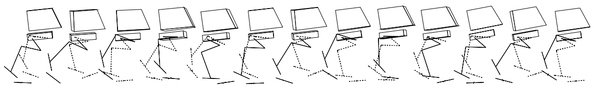

Figure 22. Running motion for the robot.

particularly difficult to control because of their unstable nature. The limit cycle control technique offers an automated way of adding closed-loop control to a basic desired open-loop motion. The open-loop component of the control can be tailored in a variety of ways to produce stylistic variations and useful parameterizations of the motion without any loss of physical realism.

While the human walks obtained are not yet equivalent to motion-capture quality, they are among the first demonstrations that general control techniques can indeed be developed for figures of relatively high complexity performing unstable motions such as walking. For imaginary creatures, physically-based simulation at present provides the best way of ensuring that motions abide by all the laws of physics. Thus, a general method of providing closed-loop control for such simulations is of considerable importance.

In the near future, we foresee integrating the closed-loop motion control developed here with an ever-growing library of other types of skilled motor control in order to produce simulated synthetic actors capable of a truly diverse set of physically-correct behaviors.

8. REFERENCES

[23] J. AUSLANDER ET AL. Further Experience With Controller-based Automatic Motion Synthesis For Articulated Figures. ACM Transactions on Graphics, October 1995.

[1] N. I. BADLER, B. BARSKY and D. ZELTZER. Making them move. Morgan Kaufmann Publishers Inc., 1991.

[2] R. BOULIC, N. M. THALMANN and D. THALMANN. A Global Human Walking Model With Real-time Kinematic Personification. The Visual Computer, 6, 1990, pp. 344-358.

[3] A. BRUDERLIN and T. W. CALVERT. Goal-Directed Animation of Human Walking. Proceedings of SIGGRAPH 89 (1989), In Computer Graphics 23, 4, (1989), pp. 233-242.

[4] A. BRUDERLIN and T. W. CALVERT. Interactive Animation of Personalized Human Locomotion. Proceedings of Graphics Interface (1993), pp. 17-23.

[5] A. BRUDERLIN and L. WILLIAMS. Motion Signal Processing. Proceedings of SIGGRAPH 95 (Los Angeles, August,1995). In Computer Graphics Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 97-104.

[6] H. C. CHANG, ET AL. A General Approach For Constructing The Limit Cycle Loci Of Multiple-Nonlinearity Systems. IEEE Transactions on Automatic Control, AC-32, 9, 1987, pp. 845-848.

[7] W. T. DEMPSTER and G. R. L. GAUGHRAN. Properties Of Body Segments Based On Size And Weight. American Journal of Anatomy, 1965, 120, 33-54.

[8] J. FURUSHO and M. MAUBUCHI. A Theoretically Motivated Reduced Order Model for the Control of Dynamic Biped Locomotion. Journal of Dynamic Systems, Measurement, and Control, 109, 1987, pp. 155-163.

[9] J. FURUSHO and A. SANO. Sensor-Based Control of a Nine-Link Robot. The International Journal of Robotics Research, 9, 2, 1990, pp. 83-98.

[10] M. GIRARD. Interactive Design Of Computer-animated Legged Animal Motion. IEEE Computer Graphics and Applications, 7, 6, June, 1987, pp. 39-51.

[11] C. L. GOLLIDAY and H. HEMAMI. An Approach to Analyzing Biped Locomotion Dynamics and Designing Robot Locomotion Controls. IEEE Transactions on Automatic Control, AC-22, 6, 1970, pp. 963-972.

[12] R. GRZESZCZUK and D. TERZOPOULOS. Automated Learning of Muscle-Actuated Locomotion Through Control Abstraction. Proceedings of SIGGRAPH 95 (Los Angeles, California, August 1995). In Computer Graphics Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 63-70.

[13] H. M. HMAM and D. A. LAWRENCE. Robustness Analysis of Nonlinear Biped Control Laws Via Singular Perturbation Theory. Proceedings of the 31st IEEE Conference on Decision and Control, 1992, pp. 2656-2661.

[14] J. K. HODGINS ET AL. Animating Human Athletics. Proceedings of SIGGRAPH 95 (Los Angeles, August, 1995). In Computer Graphics Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 71-78.

[15] R. KATOH and M. MORI. Control Method of Biped Locomotion Giving Asymptotic Stability Of Trajectory. Automatica, 20, 1984, pp. 405-414.

[16] H. KO and N. I. BADLER. Straight Line Walking Animation Based on Kinematic Generalization that Preserves the Original Characteristics. Proceedings of Graphics Interface '93, 1993, pp. 9-16.

[17] D. E. KODITSCHEK and M. BÜHLER. Analysis Of A Simplified Hopping Robot. The International Journal of Robotics Research, 10, 6, 1991, pp. 587-605.

[18] J. F. LASZLO. Controlling Bipedal Locomotion for Computer Animation, M.A.Sc. thesis, University of Toronto, 1996. URL: <www.dgp.utoronto.ca/~jflaszlo>

[19] J. M. LIN and K. W. HAN. Reducing The Effects Of Model Reduction On Stability Boundaries And Limit-Cycle Characteristics. IEEE Transactions on Automatic Control, AC-31, 6, 1986, pp. 567-569.

[20] T. MCGEER. Passive Dynamic Walking. The International Journal of Robotics Research, 9, 2, 1990, pp. 62-82.

[21] T. MCGEER. Passive Walking with Knees. Proceedings of IEEE International Conference on Robotics and Automation, 1990, pp. 1640-1645.

[22] M. MCKENNA and D. ZELTZER. Dynamic Simulation of Autonomous Legged Locomotion. Proceedings of SIGGRAPH 90 (1990). In Computer Graphics (1991), pp. 29-38.

[23] H. MIURA and I. SHIMOYAMA. Dynamic Walk Of A Biped. International Journal of Robotics Research, Summer 1984, pp. 60-74.

[24] J. T. NGO and J. MARKS. Spacetime Constraints Revisited. Proceedings of SIGGRAPH 93 (1993). In Computer Graphics Proceedings, Annual Conference Series, 1993, ACM SIGGRAPH, pp. 343-350.

[25] M. H. RAIBERT. Legged Robots that Balance. MIT Press, 1986.

[26] M. H. RAIBERT and J. K. HODGINS. Animation Of Dynamic Legged Locomotion. Proceedings of SIGGRAPH 91 (1991). In Computer Graphics, 1991, pp. 349-358.

[27] SYMBOLIC DYNAMICS INC. SD/Fast User's Manual, 1990.

[28] K. SIMS. Evolving Virtual Creatures. Proceedings of SIGGRAPH 94 (Orlando, Florida, July, 1994). In Computer Graphics Proceedings, Annual Conference Series, 1994, ACM SIGGRAPH, pp. 15-22.

[29] A. J. STEWART and J. F. CREMER. Beyond Keyframing: An Algorithmic Approach to Animation. Proceedings of Graphics Interface '92, 1992, pp. 273-281.

[30] M. UNUMA, K. ANJYO and R. TAKEUCHI. Fourier Principles for Emotion-based Human Figure Animation. Proceedings of SIGGRAPH 95 (Los Angeles, California, August, 1995). In Computer Graphics Proceedings, Annual Conference Series, 1995, ACM SIGGRAPH, pp. 91-96.

[31] M. VAN DE PANNE and E. FIUME. Sensor-Actuator Networks. Proceedings of SIGGRAPH 93, (1993). In Computer Graphics Proceedings, Annual Conference Series, 1993, ACM SIGGRAPH, pp. 335-342.

[32] M. VAN DE PANNE, R. KIM and E. FIUME. Virtual Wind-up Toys for Animation. Proceedings of Graphics Interface '94, 1994, pp. 208-215.

[33] H. G. VISSER and J. SHINAR. First-Order Corrections In Optimal Feedback Control Of Singularly Perturbed Nonlinear Systems. IEEE Transactions on Automatic Control, AC-31, 5, 1986, pp. 387-393.

[34] VUKOBRATOVIC ET AL. Biped Locomotion: Dynamics, Stability, Control and Applications, Springer Verlag, 1990.

[35] A. WITKIN and Z. POPOVI´C. Motion Warping. Proceedings of SIGGRAPH 95 (Los Angeles, California, August, 1995). In Computer Graphics Proceedings, Annual Conference Series, 1995 ACM SIGGRAPH, pp. 105-107.

[36] W. L. WOOTEN and J. K. HODGINS. Simulation Of Human Diving. Proceedings of Graphics Interface '95, 1995, pp. 1-9.