

# Animating Human Athletes

J. K. Hodgins and W. L. Wooten

College of Computing, Georgia Institute of Technology Atlanta, GA (USA)

email: [jkh|wlw]@cc.gatech.edu

## Abstract

*This paper describes algorithms for the animation of male and female models performing three dynamic athletic behaviors: running, bicycling, and vaulting. We animate these behaviors using control algorithms that cause a physically realistic model to perform the desired maneuver. For example, control algorithms allow the simulated humans to maintain balance while moving their arms, to run or bicycle at a variety of speeds, and to perform two vaults. For each simulation, we compare the computed motion to that of humans performing similar maneuvers. We perform the comparison both qualitatively through real and simulated video images and quantitatively through simulated and biomechanical data.*

## 1 Introduction

People are skilled at perceiving the subtle details of human motion. We can, for example, often identify friends by the style of their walk when they are still too far away to be recognizable otherwise. If synthesized human motion is to be compelling, we must create actors for computer animations and virtual environments that appear realistic when they move. We use dynamic simulation coupled with control algorithms to produce natural-looking motion.

In particular, this paper describes algorithms that allow a rigid-body model of a man or woman to stand, run, and turn at a variety of speeds, to ride a bicycle on hills and around obstacles, and to perform gymnastic vaults (Fig. 1). The rigid-body models of the man and woman are realistic in that their mass and inertia properties are derived from data in the biomechanics literature and the degrees of freedom of the joints are chosen so that each behavior can be completed in a natural-looking fashion.

Although the behaviors are very different in character, the control algorithms are built from a common toolbox: state machines are used to

enforce a correspondence between the phase of the behavior and the active control laws, synergies are used to cause several degrees of freedom to act with a single purpose, limbs without required actions in a particular state are used to reduce disturbances to the system, inverse kinematics is used to compute the joint angles that would cause a foot or hand to reach a desired location, and the low-level control is performed with proportional-derivative control laws.

We have chosen to animate running, bicycling, and vaulting because each behavior contains a significant dynamic component. For these behaviors, the dynamics of the model constrain the motion and limit the space that must be searched to find control laws for natural-looking motion. This property is most evident in the gymnastic vault. The gymnast is airborne for much of the maneuver, and the control algorithms can influence the internal motion of the joints but not the angular momentum of the system as a whole. The runner, on the other hand, is in contact with the ground much of the time and the joint torques computed by the control algorithms directly control many of the details of the motion. Because the dynamics do not provide as many constraints on the motion, much more effort went into tuning the motion of the runner than into tuning the motion of the gymnast.

Many computer animations and interactive virtual environments require a source of human motion. The approach used here, dynamic simulation coupled with control algorithms, is only one of several options. Two alternatives, motion capture and keyframing, have been demonstrated to be practical and are widely available in commercial systems. The difficulty of designing control algorithms has prevented the value of simulation from being demonstrated for systems with internal sources of energy such as humans. However, simulation has several potential advantages over other approaches. Given robust control algorithms, simulated motion can easily be computed to produce similar but different motions while



Figure 1: Images of an athlete running on a quarter mile track in the 1996 Olympic Stadium, comparison of the simulated runner and a runner on a treadmill, a bicyclist jumping off the curb, a group of bicyclists riding around a corner during a race, a gymnast performing an arabian vault, and comparison of an handspring vault and a simulated performance. In each case, the spacing of the images in time is equal with the runners at intervals of 0.066 s, the single rider and the group of bicyclists at 0.33 s, and the gymnasts at 0.5 s.

maintaining physical realism (running at 4 m/s rather than 5 m/s for example). Real-time simulations also allow the motion of an animated character to be truly interactive, an important property for virtual environments in which the actor must move realistically in response to changes in the environment and in response to the actions of the user. And finally, when the source of motion is dynamic simulation we have the opportunity to use multiple levels of simulation to generate either secondary motion such as the movement of clothing and hair or higher-level behaviors such as obstacle avoidance and grouping.

## 2 Background

Research in three fields is relevant to the problem of animating human motion: robotics, biomechanics, and computer graphics. Researchers in robotics have explored control techniques for legged robots that walk, run, balance, and perform gymnastic maneuvers. While few robots have been built with a complexity similar to that of the human body, control strategies for simpler machines provide basic principles that can be used to design control strategies for human models.

Raibert and his colleagues built and controlled a series of dynamic running machines, ranging from a planar machine with one telescoping leg to three-dimensional machines that ran on two or four legs. These machines walked, jumped, changed gait, climbed stairs, and performed gymnastic maneuvers ([22], [12], [13], [11], [23], and [21]). The control algorithms for human running described in this paper build on these control algorithms by extending them for a system with many more controlled degrees of freedom and more stringent requirements on the style of the motion.

Biomechanics provides the data and hypotheses about human motion required to ensure that the computed motion resembles that of a human performing similar maneuvers. The biomechanics literature contains motion capture data, force plate data, and muscle activation records for many human behaviors. These data were used to tune the control algorithms for running, bicycling, vaulting, and balancing. Cavagna presents energy curves for walking and running as well as studies of energy usage during locomotion[5]. McMahon provides graphs of stance duration, flight duration, and step length as a function of forward speed[18]. Gregor surveys biomechanical studies of bicyclists[10]. Takei

presents biomechanical data of elite female gymnasts performing a handspring vault and relates the data to the scores that the gymnasts received in competition[27].

Many researchers in computer graphics have explored the difficult problems inherent in animating human motion. The Jack system developed at the University of Pennsylvania contains kinematic and dynamic models of humans based on biomechanical data[1]. It allows the interactive positioning of the body and has several built-in behaviors including balance, reaching and grasping, and walking and running behaviors that use generalizations of motion capture data[15]. Jack has been used extensively for ergonomic analysis and human factors engineering as well as distributed simulation applications.

Bruderlin and Calvert used a simplified dynamic model and control algorithms to generate the motions of a walking human[4]. The leg model included a telescoping leg with two degrees of freedom for the stance phase and a compound pendulum model for the swing phase. A foot, upper body, and arms were added to the model kinematically, and were made to move in an oscillatory pattern similar to that observed in humans. Pai programmed a walking behavior for a dynamic model of a human torso and legs in a high-level fashion by describing a set of time-varying constraints, such as, “maintain ground clearance during leg swing,” “lift and put down a foot,” “keep the torso vertical,” and “support the torso with the stance leg”[20].

None of these approaches to generating motion for animation are automatic because each new behavior requires significant additional work on the part of the researcher. In recent years, the field has seen the development of a number of techniques for automatically generating motion for new behaviors and new creatures. Witkin and Kass[31], Cohen[7], and Brotman and Netravali[3] treat the problem of automatically generating motion as a trajectory optimization problem. Another approach finds a control algorithm instead of a desired trajectory ([30], [29], [19], [25], [26] and [16]). In contrast, the control algorithms described in this paper were designed by hand, using a toolbox of control techniques, our physical intuition about the behaviors, observations of humans performing the tasks, and biomechanical data. While automatic techniques would be preferable to hand design, automatic techniques have not yet been developed that can find solutions for systems with the number of controlled degrees of freedom needed for a realistic model

of the human body. Furthermore, although the motion generated by automatic techniques is appealing, much of it does not appear natural in the sense of resembling the motion of a biological system. We do not yet know whether this discrepancy is because only relatively simple models have been used or because of the constraints and optimization criteria that were chosen.

### 3 Dynamic Behaviors

The motion of each behavior described in this paper is computed using dynamic simulation. Each simulation contains the equations of motion for a rigid-body model of a human and environment (ground, bicycle, and vaulting horse), control algorithms for balancing, running, bicycling, or vaulting, a graphical model for viewing the motion, and a user interface for changing the parameters of the simulation. The user is provided with limited high-level control of the animation. For example, the desired velocity and facing direction for the bicyclist and runner are selected by the user. During each simulation time step, the control algorithm computes desired positions and velocities for each joint based on the state of the system, the requirements of the task, and input from the user. Proportional-derivative servos compute joint torques based on the desired and actual value of each joint. The equations of motion of the system are integrated forward in time taking into account the internal joint torques and the external forces and torques from interactions with the ground plane or other objects. The details of the human model and the control algorithm for each behavior are described below.

The human models were constructed from rigid links connected by rotary joints with one, two or three degrees of freedom. The dynamic models were derived from the graphical models shown in Fig. 2. The mass and moment of inertia of each body part was calculated by computing the moment of inertia of a polygonal object of uniform density[17] using density data measured from cadavers[8]. The controlled degrees of freedom of the models are shown in Fig. 2. Each internal joint of the model has a very simple muscle model, a torque source, that allows the control algorithms to apply a torque between the two links that form the joint. The equations of motion for each system were generated using a commercially available package[24].

The points of contact between the feet and the ground, and the gymnast’s hands and the vault are modeled using constraints. Interaction of the




Figure 2: The controlled degrees of freedom of the human model. The gymnast represented in the figure has 15 body segments and a total of 30 controlled degrees of freedom. The runner has 17 body segments and 30 controlled degrees of freedom (two-part feet with a one degree-of-freedom joint at the ball of the foot and only one degree of freedom at the ankle), The bicyclist has 15 body segments and 22 controlled degrees of freedom (only one degree of freedom at the neck, hips, and ankles). The directions of the arrows indicates the positive direction of rotation for each degree of freedom. The polygonal models were purchased from Viewpoint Datalabs.

feet with the ground use six constraints for each foot: two keep the metatarsus and heel above the surface of the ground, two prevent the foot from sliding on the ground, and two prevent the foot from rolling or twisting about the center of mass of the foot (Fig. 3). The linear and rotational acceleration of the contact point of the foot with respect to the ground is the constraint error. The penetration of the foot into the ground and the velocity of the foot relative to the ground stabilize the constraint error[2]. Friction is infinite in our model. The ground contact forces computed with the constraint matrix are applied to the foot at the contact points to prevent the foot from penetrating or slipping on the ground. Torques are applied about the center of mass of the foot to prevent rolling or twisting. To allow the feet to leave the ground, the force is applied only when the foot has penetrated the ground and the vertical velocity of the foot is negative. Contact between the hands and the environment are modeled in the same manner as the feet with contact points at the base of the palm and the tips of the fingers.

The points of contact between the bicycle wheels and the ground are also modeled using constraints. The constraints allow rolling with-




Figure 3: Six constraints are used for each foot. Two ground reaction force constraints prevent the foot from penetrating the ground. Two constraints prevent the foot from sliding on the plane in  $x$  and  $y$ . Two more constraints are used to prevent the foot from rolling or yawing about its center of mass.

out slipping. Yaw damping is applied to the wheel about the  $z$  axis and is proportional to the yaw rotation rate of the wheel up to a friction limit.

### 3.1 Running

Running is a cyclic behavior in which the legs swing fore and aft and provide support for the body in alternation. Because the legs perform different functions during the phases of the locomotion cycle, the muscles are used for different control actions at various times in the cycle. When the foot of the simulated runner is on the ground, the ankle, knee, and hip provide support and balance. During the flight phase, a leg is swung forward in preparation for the next touchdown. These distinct phases and corresponding changes in control actions make a state machine a natural tool for selecting the control actions that should be active at a particular time. The state machine and transition events used for the simulation of running are shown in Fig. 4. The main states correspond to the points of contact with the ground: flight, heel contact, heel and toe contact, and toe contact. Transitions between states are based on events that can be easily determined from the variables of the simulation (contact of the foot with the ground or movement of the knee joint, for example). The control laws for each state are described in Fig. 4.

To interact with the animation of the runner, the user specifies desired values for the magnitude of the velocity on the ground plane and the facing direction. The control laws for each state compute joint torques that move the velocity and facing direction toward these desired values while maintaining balance. The simulated runner can

run at speeds between 2.5 m/s and 5 m/s.

We call the leg that is on the ground or actively being positioned for touchdown the *active leg*. The other leg is the *idle leg*. During flight, the active leg is swung forward in anticipation of touchdown. Using the degrees of freedom of the leg in a synergistic fashion, the foot is positioned at touchdown to correct for errors in forward speed and to maintain balance. Forward speed is controlled by placing the average point of support during stance underneath the hip and taking into account the change in contact point from heel to metatarsus during stance. At touchdown, the desired distance from the hip to the heel projected onto the ground plane is

$$x_{hh} = 1/2(t_s \dot{x} - \cos(\theta)l_f) + k(\dot{x} - \dot{x}_d) \quad (1)$$

where  $t_s$  is an estimate of the period of time that the foot will be in contact with the ground (based on the previous stance duration),  $\dot{x}$  is the forward velocity of the runner,  $\dot{x}_d$  is the desired velocity,  $\theta$  is the facing direction of the runner,  $l_f$  is the distance from the heel to the ball of the foot, and  $k$  is a gain for the correction of errors in speed. A similar equation holds for  $y_{hh}$ . The length of the leg at touchdown is fixed and is used to calculate the vertical distance from the hip to the heel,  $z_{hh}$ . The disturbances caused by the impact of touchdown can be reduced by decreasing the relative speed between the foot and the ground at touchdown. This technique is commonly called *ground speed matching*. In this control system, ground speed matching is accomplished by swinging the hip further forward in the direction of travel during flight and moving it back just before touchdown.

The equations for  $x_{hh}$ ,  $y_{hh}$ , and  $z_{hh}$  and the kinematics of the model are used to compute the desired knee and hip angles at touchdown. The angle of the ankle is constant during flight and is chosen such that the toe will not touch the ground at the same time as the heel at the beginning of stance.

The idle leg plays an important role in locomotion by reducing disturbances to the body attitude caused by the active leg as it swings forward and in toward the centerline in preparation for touchdown. The idle leg is bent so that the toe does not stub the ground, and the hip angles mirror the motion of the active leg to reduce the net torque on the body:

$$\alpha_{x_d} = \alpha_{x_{l_o}} - (\beta_{x_d} - \beta_{x_{l_o}}) \quad (2)$$

where  $\alpha_{x_d}$  is the desired rotation of the idle hip with respect to the pelvis,  $\alpha_{x_{l_o}}$  is the rotation

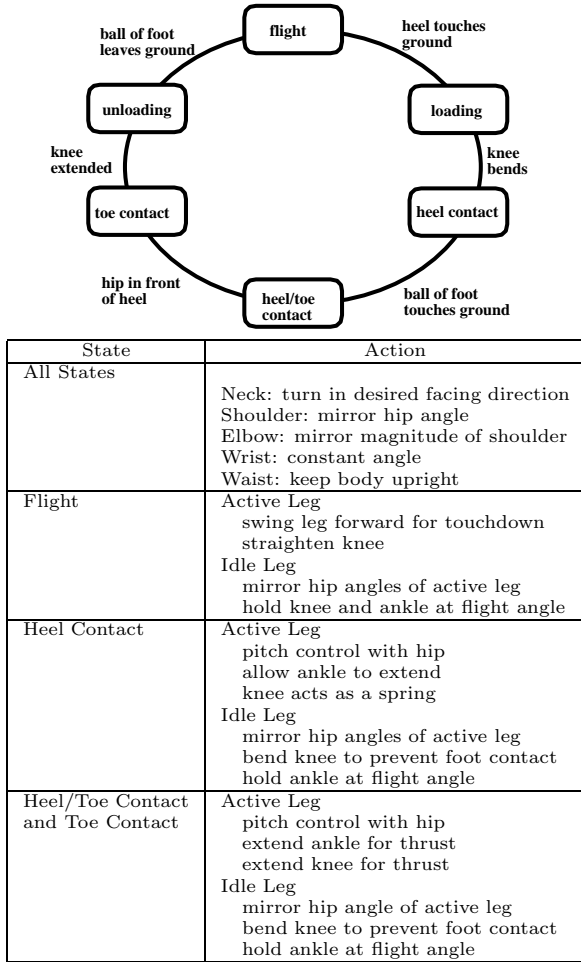


Figure 4: A state machine is used to determine the control actions that should be active for running given the current state of the system. The transition events are computed for the active leg. At liftoff the active and idle legs switch roles.

of the idle hip at the previous liftoff,  $\beta_{x_d}$  is the desired position of the active hip, and  $\beta_{x_{l_o}}$  is the position of the active hip at the previous liftoff. A similar equation holds for the  $y$  component of the hip rotation. The mirroring action of the idle leg is modified by the restriction that the legs should not collide as they pass each other during stance.

During stance, the knee acts as a passive spring to store the kinetic energy that the system had at touchdown. The majority of the vertical thrust is provided by the ankle joint. During the first part of stance, *heel contact*, the toe is passively moved toward the ground because the contact point on the heel is behind the ankle joint. Contact of the ball of the foot triggers the transition from *heel contact* to *heel and toe contact*. The transition from *heel and toe contact* to *toe contact* oc-

curs when the hip has moved a certain distance in front of foot. After the transition, the ankle joint is extended, causing the heel to lift off the ground and adding energy to the system for the next flight phase.

Throughout stance, proportional-derivative servos are used to compute torques for the hip joint of the stance leg that will cause the attitude of the body (roll, pitch, and yaw) to move toward the desired values. The desired angle for roll is zero except during turning. The desired angle for pitch is inclined slightly forward, and the desired angle for yaw is set by the user. The  $x$  and  $y$  joints of the waist maintain the body attitude by allowing the upper part of the body to be more vertical than the pelvis. The  $z$  joint of the waist allows the upper body to turn to face in the direction specified by the user.

Throughout the running cycle, the shoulder and elbow joints swing the arms fore and aft in a motion that is synchronized with the motion of the legs:  $\gamma_{y_d} = k\alpha_y + \gamma_0$  where  $\gamma_{y_d}$  is the desired fore/aft angle for the shoulder,  $k$  is a scaling factor,  $\alpha_y$  is the fore-aft hip angle for the leg on the opposite side of the body, and  $\gamma_0$  is an offset. The other two degrees of freedom in the shoulder ( $x$  and  $z$ ) and the elbows also follow a cyclic pattern with the same period as  $\gamma_{y_d}$ . The neck is used to turn the head to look at the ground in front of the runner in the desired facing direction. The motion of the upper body is important in running because the counter oscillation of the arms reduces the yaw oscillation of the body caused by the swinging of the legs. However, the details of the motion of the upper body are not strongly constrained by the dynamics of the task and amateur athletes use many different styles of arm motion. Observations of human runners were used to tune the oscillations of the arms to produce a natural-looking gait.

The control laws compute desired values for each joint. Proportional-derivative servos are used to control the position of all joints. For each internal joint the control equation is

$$\tau = k(\theta_d - \theta) + k_v(\dot{\theta}_d - \dot{\theta}) \quad (3)$$

where  $\theta$  is the angle of the joint,  $\theta_d$  is the desired angle, and  $k$  and  $k_v$  are the proportional and derivative gains. The desired values used in the proportional-derivative servos are computed as trajectories from the current desired value of the joint to the new desired value computed by the control laws. Eliminating large step changes in the errors used in the proportional-derivative servos smoothes the simulated motion.




Figure 5: The four degrees of freedom of the bicycle model. The direction of the arrows indicates the positive direction of rotation for each degree of freedom. The polygonal model is a modification of a model purchased from Viewpoint Data-labs.

To generate running motion for a specific virtual environment or animated sequence, the user specifies a path through the environment and a desired velocity along that path. As the simulated runner moves along the path, the control system computes a desired facing direction that will cause the runner to face along a line between the runner’s current location and the point on the path that he will reach in 2 seconds.

When the path is curved, the control algorithms are modified to allow the runner to turn more tightly. The timing and magnitude of thrust are functions of the yaw velocity in order to decrease the thrust from the inside leg and increase the thrust from the outside leg. This modification to the control laws maintains an approximately constant flight duration for the two legs. The idle leg is swung further to the outside to avoid a collision with the active leg during stance and the desired roll angle is computed to cause the torso to lean into the curve.

### 3.2 Bicycling

The bicyclist controls the facing direction and speed of the bicycle by applying forces to the handlebars and pedals with his hands and feet. The rider is attached to the bicycle by a pivot joint between the bicycle seat and the pelvis (Fig. 5). Spring and damper systems connect the hands to the handlebars, the feet to the pedals, and the crank to the rear wheel. The connecting springs are two-sided and the bicyclist is able to pull up on the pedals as if the bicycle were equipped with toe-clips and a fixed gear (no freewheel). The connection between the crank and the rear wheel includes an adjustable gear ratio. The bicycle wheels have a rolling resistance proportional to

the velocity.

The control algorithm adjusts the velocity of the bicycle by using the legs to produce a torque at the crank. The desired torque at the crank is  $\tau_c = k(v - v_d)$  where  $k$  is a gain,  $v$  is the magnitude of the bicyclist’s velocity, and  $v_d$  is the desired velocity. The force applied by each leg depends on the angle of the crank because we assume that the legs are most effective at pushing downwards. For example, the front leg can generate a positive torque and the rear leg can generate a negative torque when the crank is horizontal. To compensate for the crank position, the desired forces for the legs are scaled by a weighting function between zero and one that depends on the crank position,  $\theta_c$ :

$$w = \frac{\sin(\theta_c) + 1}{2}. \quad (4)$$

$\theta_c$  is zero when the crank is vertical and the right foot is higher than the left. If  $\tau_c > 0$ , the force on the pedal that the legs should produce is

$$f_l = \frac{w\tau_c}{l} \quad (5)$$

$$f_r = \frac{(1-w)\tau_c}{l} \quad (6)$$

where  $f_l$  and  $f_r$  are the desired forces from the left and right legs respectively, and  $l$  is the length of a crank arm. If  $\tau_c$  is less than zero, then the equations for the left and right leg are switched. An inverse kinematic model of the legs is used to compute hip and knee torques that will produce the desired pedal forces.

To steer the bicycle and control the facing direction, the control algorithm computes a desired angle for the fork based on the errors in roll and yaw:

$$\theta_f = -k_\alpha(\alpha - \alpha_d) - k_{\dot{\alpha}}\dot{\alpha} + k_\beta(\beta - \beta_d) + k_{\dot{\beta}}\dot{\beta} \quad (7)$$

where  $\alpha$ ,  $\alpha_d$ , and  $\dot{\alpha}$  are the roll angle, desired roll, and roll velocity and  $\beta$ ,  $\beta_d$ , and  $\dot{\beta}$  are the yaw angle, desired yaw, and yaw velocity.  $k_\alpha$ ,  $k_{\dot{\alpha}}$ ,  $k_\beta$ , and  $k_{\dot{\beta}}$  are gains. The desired yaw angle is set by the user or high-level control algorithms; the desired roll angle is zero. Inverse kinematics is used to compute the shoulder and elbow angles that will position the hands on the handlebars with a fork angle of  $\theta_f$ . Proportional-derivative servos move the shoulder and elbow joints toward those angles.

These control laws leave the motion of several of the joints of the bicyclist unspecified. The wrists and the waist are held at a nearly

constant angle with proportional-derivative controllers. The ankle joints are controlled to match data recorded from human subjects[6].

### 3.3 Vaulting

To perform a vault, the gymnast uses a springboard to launch herself toward the horse, pushes off the horse with her hands, and lands on her feet on the other side of the horse. While airborne, the gymnast performs a combination of somersaults or twists. We have implemented two vaults, a handspring and an arabian. In the handspring vault, the gymnast performs a full somersault over the horse while keeping her body extended in a layout position. In the arabian vault, the gymnast performs a half-twist piked and lands facing the horse.

The handspring and arabian vaults are structured by a state machine with seven phases: *compression*, *decompression*, *flight1*, *horse*, *flight2*, *landing* and *balance*. The control actions required during each of the seven phases are presented in Fig. 6. The simulation of each vault begins in a flight state just before touchdown on the springboard. The initial conditions were estimated from video footage and confirmed with data measured from gymnasts[27] (forward velocity is 6.75 m/s and the height of the center of mass is 0.9 m).

The vaulting simulation uses the same control hierarchy as the running and bicycling simulations. A state machine selects the appropriate phase for the vault and middle level controllers calculate new desired angles for the joints. The lowest level of control uses proportional-derivative servos to compute the torque required to generate the desired motion. The next four sections will discuss the implementation of the middle level controllers for vaulting: leaping, tumbling, landing, and balancing.

#### 3.3.1 Leaping

The leaping controller has two phases, *compression* and *decompression* and the control laws are similar to those used during the stance phase of running. During *compression*, the simulated gymnast lands on a springboard that is modeled as a board attached to the ground with a pivot joint. Using a technique called *blocking*, the control system positions the hips forward before the touchdown on the springboard so that some of the horizontal velocity at touchdown is transformed into rotational and vertical velocity at liftoff.

When the springboard reaches maximum deflection, the simulation enters the *decompression*




Figure 7: The control system positions the hands on the line between the shoulder and the horse. The desired shoulder angle is  $\gamma_{y_d}$ .

phase. In this phase the control system extends the knees to push on the springboard and add energy to the system. As the springboard rebounds, the simulation extends the leg muscles rapidly and is launched into the air. When the simulated gymnast's feet leave the springboard, the *flight1* phase begins.

#### 3.3.2 Tumbling

During the *flight1* phase, the control system prepares to put the gymnast's hands on the horse by positioning her hands on the line between the shoulders and the desired hand position on the vault:  $\gamma_{y_d} = \lambda_y - \phi$ , where  $\gamma_{y_d}$  is the desired shoulder angle relative to the body,  $\lambda_y$  is the angle between horizontal and a vector from the shoulder to the desired hand position on the vault, and  $\phi$  is the pitch angle of the body with respect to horizontal (Fig. 7). This control strategy causes the hands to be nearly motionless with respect to the horse at contact, in a similar fashion to the ground speed matching strategy that reduces the impact at touchdown in the simulated runner. The wrists are servoed to ensure that the hands hit palm down and parallel to the surface of the horse:  $\phi_{y_d} = \arccos(V_a \cdot V_x)$ , where  $\phi_{y_d}$  is the desired angle for the wrist,  $V_a$  is a vector from the palm to the shoulder, and  $V_x$  is a vector along the positive  $x$ -axis.

While her hands are in contact with the horse, the gymnast generates the angular momentum required for the maneuver. For the handspring vault, the gymnast maintains a layout position. During an arabian vault, the gymnast enters a pike position before her hands leave the horse in order to generate the required twist. The control system does not apply a torque at the shoulders or wrists for either vault, allowing the gymnast's angular and forward velocity to carry her over the



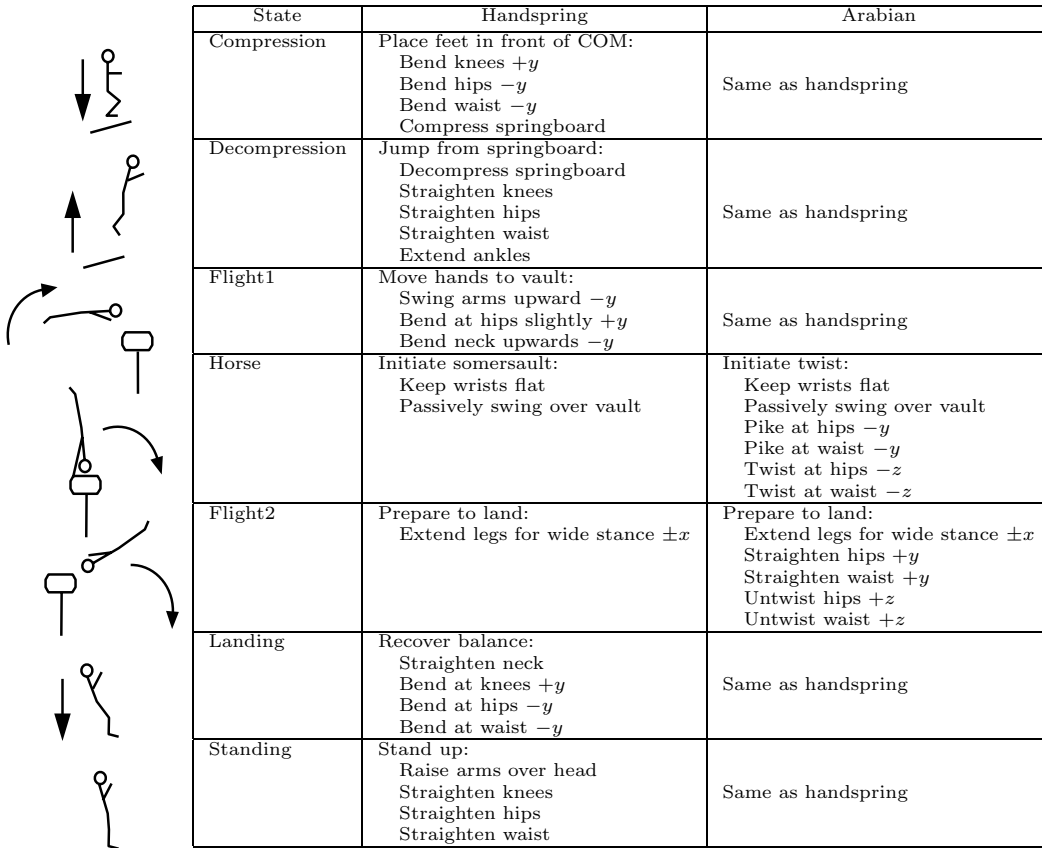


Figure 6: The state machine determines the control laws that are in effect at each phase of the vault.




Figure 8: The gymnast performs a half twist using a *swivel hips* maneuver. As the gymnast leaves the horse she rotates her upper torso as far as possible while unpiking the lower torso. As the hips straighten, the body performs the twist about the  $z$ -axis. The gymnast then “swivels” her hips and ends the maneuver by pulling her arms down while returning to a piked position (figure adapted from Frohlich[9]).

horse.

During the *flight2* phase, the control system completes the specified maneuver and prepares to land. For the handspring vault, the control system maintains the layout position. For the arabian vault, the control system performs a half-twist by moving the limbs in a specified sequence to perform a *swivel hips* maneuver (Fig. 8). This maneuver allows the gymnast to rotate the body  $180^\circ$  about the  $z$ -axis even when the gymnast has little angular momentum. Often gymnasts will

swing their arms to generate more twist, but the control system for the simulated gymnast did not use this strategy and all of the twisting velocity for the arabian vault was generated by piking the hips and twisting the body.

As the simulated gymnast’s feet approach the mat, the control system prepares for landing. The feet are spread slightly to give a larger area of support. The angle of the legs with respect to the ground at touchdown is chosen so that the gymnast’s forward velocity will decrease until it is nearly zero when her center of mass is above the support polygon. This strategy and the blocking strategy used for the springboard contact are essentially the same as the strategy used to control the forward speed of the runner.

### 3.3.3 Landing

When the feet hit the ground, the control system for the *landing* phase must remove the horizontal and rotational energy that was required for the maneuver and must establish an upright, balanced posture. At touchdown, the gymnast is not in a balanced state because the projection of the center of mass is outside the support polygon. But if the angle of the legs at touchdown was chosen correctly, the forward velocity will be near

zero when the center of mass is over the center of the support polygon. Feedback control is used to servo the ankles to adjust the forward velocity of the center of mass as it moves forward. During landing, the knees and waist bend to absorb most of the energy.

The controllers for the handspring vault and arabian vaults are similar except that the orientation of the landing is reversed. We found the handspring vault easier to control than the arabian because it is easier to recover balance when the twisting angular velocities are near zero.

### 3.3.4 Balancing

When the gymnast’s forward velocity is nearly zero, the *balance* phase begins. The balance controller straightens the hips and knees to cause the gymnast to stand up from the crouch used during landing. Changes in the ankle and hip angles compensate for disturbances to the location of the center of mass caused by motion of the upper body and arms.

The balance controller moves the projection of the center of mass of the gymnast on the ground,  $(C_x, C_y)$ , towards the desired center of mass,  $(C_{x_d}, C_{y_d})$ . For balance, the center of mass should be in the center of the polygon formed by the four contact points of the feet, the *support polygon*. The desired angles for the ankle are

$$\rho_{x_d} = -(k_{ax}(C_{x_d} - C_x) - \dot{k}_{ax}\dot{C}_x) \quad (8)$$

$$\rho_{y_d} = k_{ay}(C_{y_d} - C_y) - \dot{k}_{ay}\dot{C}_y \quad (9)$$

where  $k_{ax}$ ,  $k_{ay}$ ,  $\dot{k}_{ax}$ , and  $\dot{k}_{ay}$  are gains for the proportional-derivative controller. The same equations are used for both the left and right ankles. Similar equations are used to compute the desired hip angles, the only difference being the sign and the gains

$$\alpha_{x_d} = k_{hx}(C_{x_d} - C_x) - \dot{k}_{hx}\dot{C}_x \quad (10)$$

$$\alpha_{y_d} = -(k_{hy}(C_{y_d} - C_y) - \dot{k}_{hy}\dot{C}_y) \quad (11)$$

This controller not only allows the gymnast to stand after landing, but also compensates for disturbances resulting from the motion of other parts of the body while she is standing. For example, balance is maintained as the simulated gymnast throws her arms back in a gesture of success after completing the vault.

## 4 Discussion

This paper presents algorithms that allow an animator to generate motion for several dynamic be-

haviors. Elsewhere we have presented animations of other dynamic behaviors including platform diving, unicycle riding, and pumping a swing ([28, 14]). Taken together with this previous work, the dynamic behaviors described here represent a growing library. While these behaviors do not represent all of human motion or even of human athletic endeavors, an animation package with ten times this many behaviors would have sufficient functionality to be interesting to students and perhaps even to professional animators.

Several open questions remain before the value of simulation as a source of motion for animation and virtual environments can be conclusively demonstrated:

How can we make it easier to generate control algorithms for a new behavior? This paper partially addresses that question by presenting a toolbox of techniques that can be used to construct the control algorithms for a set of diverse behaviors. However, developing sufficient physical intuition about a new behavior to construct a robust control algorithm remains time consuming. We hope that these examples represent a growing understanding of the strategies that are useful in controlling simulations of human motion and that this understanding will lead to the development of more automatic techniques.

What can we do to reduce the number of new behaviors that need to be developed? One idea that has been explored by researchers in the domain of motion capture and keyframing is to perform transitions between behaviors in an automatic or semiautomatic fashion. Such transitions may be much more amenable to automatic design than the design of entire control algorithms for dynamic simulations.

What rules can we add to the system to improve the naturalness of the motion? The techniques presented here are most effective for behaviors with a significant dynamic component because the dynamics constrain the number of ways in which the task can be accomplished. When the gross characteristics of the motion are not constrained by the dynamics of the system, the task can be completed successfully but in a way that appears unnatural. For example, the simulated runner can run while holding his arms fixed at his sides, but an animation of that motion would be amusing rather than realistic. Humans are strong enough and dextrous enough that simple arm movements such as picking up a coffee cup can be completed in many different ways. In contrast, only good athletes can perform an arabian vault and the variations seen in their perfor-

mances are relatively small. When the dynamics do not significantly constrain the task, the control algorithms must be carefully designed and tuned to produce motion that appears natural while matching the key features of the behavior when performed by a human. The tuning process might be aided by using data from psychophysical experiments to provide additional constraints on the motion.

Can human motion be simulated interactively? To be truly interactive, the motion of synthetic actors in virtual environments must be computed in real time (simulation time must be less than wall clock time). Our implementation of the bicyclist runs in real time on a Silicon Graphics Indigo<sup>2</sup> Computer with a R8000 processor. The runner is six times slower than real time on the same processor and the vaulter is twenty times slower than real time. We anticipate that with improved dynamic simulation techniques and the continued increase in workstation speed, a three-dimensional human simulation with a realistic number of degrees of freedom will run in real time within a few years.

Are the behaviors robust enough for use as interactive synthetic actors? The runner can run at a variety of speeds and change direction, but abrupt changes in velocity or facing direction will cause him to fall down. The planning or reactive response algorithms that lie between the locomotion control algorithms and the perceptual model of the simulated environment will have to take in account the limitations of the dynamic system and control system in order to provide robust but agile higher-level behaviors.

One goal of this research is to demonstrate that dynamic simulation of rigid-body models can be used to generate natural-looking motion. Fig. 1 shows a side-by-side comparison of video footage of a human runner and gymnast with images of the simulated runner and gymnast. This comparison represents one form of evaluation of our success in generating natural-looking motion. The simulated runner's leg motion is similar to that of the athlete on the treadmill although the simulated runner has a shorter stride and runs with a more upright body posture. The simulated and human vaults differ in that the simulated vaulter has a straighter position in the layout and less motion on the landing.

A second form of comparison involves the use of biomechanical data to measure how closely the simulated motion resembles natural-looking motion. Data available in the biomechanical literature such as force platform records illustrate some

Variables	Human			Sim.
	Mean	Min	Max	
Mass (kg)	47.96	35.5	64.0	64.3
Height (m)	1.55	1.39	1.66	1.64
Board contact (s)	0.137	0.11	0.15	0.105
First flight (s)	0.235	0.14	0.30	0.156
Horse contact (s)	0.245	0.19	0.30	0.265
Second flight (s)	0.639	0.50	0.78	0.632
Horiz. vel. (m/s)				
Board touchdown	6.75	5.92	7.25	6.75
Board liftoff	4.61	3.97	5.26	4.01
Horse touchdown	4.61	3.97	5.26	4.01
Horse liftoff	3.11	2.48	3.83	2.83
Vert. velocity (m/s)				
Board touchdown	-1.15	-1.54	-.71	-1.13
Board liftoff	3.34	2.98	3.87	3.81
Horse touchdown	1.26	0.74	2.39	2.13
Horse liftoff	1.46	0.56	2.47	1.10
Aver. vert. force (N)				
Board contact	2175	1396	2792	5075
Horse contact	521	309	752	957

Figure 9: Comparison of velocities, contact times, and forces for a simulated vaulter and human data measured by Takei. The human data was averaged from 24 subjects. The simulated data was taken from a single trial.

ways in which the motion of the simulated runner does not match the motion of a human. Fig. 9 compares data from the handspring vault performed by 24 gymnasts in the 1987 Pan American Games to our simulated handspring vault[27]. Most of the simulated data fell within the range of the measured data.

A final form of evaluation would be a Turing test with direct comparison between simulated motion and human motion. The question we would like to ask is the following: if simulated data and human data were represented using the same graphical model, would the viewer sometimes choose the simulated data as the more natural motion? Unfortunately, the capture techniques for human data have characteristic flaws making this test more difficult than it might at first appear.

The animations described in this paper motion can be seen on the WWW at <http://www.cc.gatech.edu/gvu/animation>

## 5 Acknowledgments

Earlier versions of portions of this paper were presented at Siggraph '95 and the 1996 IEEE International Conference on Robotics and Automation. The authors would like to thank Debbie Carlson and Ron Metoyer for their help in developing our simulation and rendering environment and the CAD Systems Department at the Atlanta Committee for the Olympic Games for allowing

us to use models of the Olympic venues. This project was supported in part by NSF NYI Grant No. IRI-9457621, Mitsubishi Electric Research Laboratory, and a Packard Fellowship. Wayne Wooten was supported by a Intel Foundation Graduate Fellowship.

## References

- [1] Badler, N. I., Phillips, C. B., Webber, B. L. 1993. *Simulating Humans*. Oxford: Oxford University Press.
- [2] Baumgarte, J. 1972. Stabilization of Constraints and Integrals of Motion in Dynamical Systems. *Computer Methods in Applied Mechanics and Engineering* 1:1–16.
- [3] Brotman, J. S., Netravali, A. N. 1988. Motion Interpolation by Optimal Control. *Proceedings of SIGGRAPH*, 309–315.
- [4] Bruderlin, A., Calvert, T. W. 1989. Goal-Directed, Dynamic Animation of Human Walking. *Proceedings of SIGGRAPH*, 233–242.
- [5] Cavagna, G. A., Thys, H., Zamboni, A. 1976. The Sources of External Work in Level Walking and Running. *Journal of Physiology* 262:639–657.
- [6] Cavanagh, P., Sanderson, D. 1986. The Biomechanics of Cycling: Studies of the Pedaling Mechanics of Elite Pursuit Riders. In *Science of Cycling*, Edmund R. Burke (ed), Human Kinetics: Champaign, Ill.
- [7] Cohen, M. F. 1992. Interactive Spacetime Control for Animation. *Proceedings of SIGGRAPH*, 293–302.
- [8] Dempster, W. T., Gaughran, G. R. L. 1965. Properties of Body Segments based on Size and Weight. *American Journal of Anatomy* 120: 33–54.
- [9] Frohlich, C. 1979. Do springboard divers violate angular momentum conservation? *American Journal of Physics* 47:583–592.
- [10] Gregor, R. J., Broker, J. P., Ryan, M. M. 1991. Biomechanics of Cycling *Exercise and Sport Science Reviews* Williams & Wilkins, Philadelphia, John Holloszy (ed), 19:127–169.
- [11] Hodgins, J. K. 1991. Biped Gait Transitions. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2092–2097.
- [12] Hodgins, J., Raibert, M. H. 1990. Biped Gymnastics. *International Journal of Robotics Research* 9(2):115–132.
- [13] Hodgins, J. K., Raibert, M. H. 1991. Adjusting Step Length for Rough Terrain Locomotion. *IEEE Transactions on Robotics and Automation* 7(3): 289–298.
- [14] Hodgins, J. K., Sweeney, P. K., Lawrence, D. G. 1992. Generating Natural-looking Motion for Computer Animation. *Proceedings of Graphics Interface '92*, 265–272.
- [15] Ko, H., Badler, N. I. 1993. Straight-line Walking Animation based on Kinematic Generalization that Preserves the Original Characteristics. In *Proceedings of Graphics Interface '93*.
- [16] Laszlo, J., van de Panne M., Fiume, E. 1996. Limit Cycle Control and its Application to the Animation of Balancing and Walking. *Proceedings of SIGGRAPH*, 155–162.
- [17] Lien, S., Kajiya, J. T. 1984. A Symbolic Method for Calculating the Integral Properties of Arbitrary Nonconvex Polyhedra. *IEEE Computer Graphics and Applications* 4(5):35–41.
- [18] McMahon, T. A. 1984. *Muscles, Reflexes, and Locomotion*. Princeton: Princeton University Press.
- [19] Ngo, J. T., Marks, J. 1993. Spacetime Constraints Revisited. *Proceedings of SIGGRAPH*, 343–350.
- [20] Pai, D. 1990. Programming Anthropoid Walking: Control and Simulation. Cornell Computer Science Tech Report TR 90-1178.
- [21] Playter, R. R., Raibert, M. H. 1992. Control of a Biped Somersault in 3D. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 582–589.
- [22] Raibert, M. H. 1986. *Legged Robots That Balance*. Cambridge: MIT Press.
- [23] Raibert, M. H., Hodgins, J. K. 1991. Animation of Dynamic Legged Locomotion. *Proceedings of SIGGRAPH*, 349–356.
- [24] Rosenthal, D. E., Sherman, M. A. 1986. High Performance Multibody Simulations Via Symbolic Equation Manipulation and Kane's Method. *Journal of Astronautical Sciences* 34(3):223–239.
- [25] Sims, K. 1994. Evolving Virtual Creatures. *Proceedings of SIGGRAPH*, 15–22.
- [26] Sims, K. 1994. Evolving 3D Morphology and Behavior by Competition. *Artificial Life IV*, 28–39.
- [27] Takei, Y., 1990. Techniques Used by Elite Women Gymnasts Performing the Handspring Vault at the 1987 Pan American Games. *International Journal of Sport Biomechanics* 6:29–55.
- [28] Wooten, W. L., Hodgins, J. K. 1996. Simulation of Human Diving. *Computer Graphics Forum* 15(1):3–13.
- [29] van de Panne M., Fiume, E. 1993. Sensor-Actuator Networks. *Proceedings of SIGGRAPH*, 335–342.
- [30] van de Panne M., Fiume, E., Vranesic, Z. 1990. Reusable Motion Synthesis Using State-Space Controllers. *Proceedings of SIGGRAPH*, 225–234.
- [31] Witkin, A., Kass, M. 1988. Spacetime Constraints. *Proceedings of SIGGRAPH*, 159–168.