Eurographics/ ACM SIGGRAPH Symposium on Computer Animation (2007)
D. Metaxas and J. Popovic (Editors)

Multiobjective Control with Frictional Contacts

Yeuhi Abe Marco da Silva Jovan Popović

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Abstract

Standing is a fundamental skill mastered by humans and animals alike. Although easy for adults, it requires careful and deliberate manipulation of contact forces. The variation in contact configuration (e.g., standing on one foot, on uneven ground, or while holding on for support) presents a difficult challenge for interactive simulation of humans and animals, especially while performing tasks in the presence of external disturbances. We describe an analytic approach for control of standing in three-dimensional simulations based upon local optimization. At any point in time, the control system solves a quadratic program to compute actuation by maximizing the performance of multiple motion objectives subject to constraints imposed by actuation limits and contact configuration. This formulation is suitable for interactive animation and it adapts to the proportions of any character model in any non-planar, frictional contact configuration.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three Dimensional Graphics and Realism, Animation

1. Introduction

Dynamic simulation of passive phenomena, such as cloth, fluids, and articulated bodies, is used in many animation systems and has enabled the creation of increasingly complex virtual environments, while reducing demands on talented human animators. This is particularly the case of games and simulations where truly interactive bodies are rapidly replacing static, precomputed motions.

The simulation of active bodies such as humans, robots, and animals has lagged behind, preventing automated animation of characters that act in concert with their dynamic surroundings. Despite previous demonstrations of simulated characters performing impressive actions, including walking, running, diving and swimming [HWBO95, WH96, YLS04], many animation systems still rely on recorded motions or passive ragdoll physics.

The widespread adoption of simulated active bodies is hindered by over-specialization. Most controllers operate correctly only for specific body postures, geometries, or physical properties of the environment. For example, a controller designed to balance a standing character on flat ground many not work on uneven or slippery ground or while leaning against a wall for support. The objective of balancing is conceptually similar in each case, but adapting a controller from one setting to another can be as difficult as developing a new one. Hence, a key advantage of general purpose simulation—automatic motion synthesis under all circumstances—disappears with current control techniques.

We describe an analytic control formulation that solves a local, online optimization to reduce over-specialization in controllers for active bodies. The optimization automatically adapts the control to the frictional properties of the simulation, the mass properties of the character, the posture of the character, and the changing task-specific goals of actions being performed. It also accounts for the constraints imposed by frictional contacts with the environment. Hence the same controller applies to simulation of active bodies in many different situations.

We explore this approach for an important class of actions involving sustained frictional contact with the environment (Figure 1). Such contact occur whenever a character pushes against the environment and uses the resulting force to control its motion. We call this fundamental behavior standing, noting that it is a precursor to locomotion and other complex behaviors. Our approach simplifies the design of standing controllers by decoupling the description of motion from the computation of forces required to accomplish them. A multi-

SCA 2007, San Diego, California, August 04 - 05, 2007
© 2007 ACM 978-1-59593-624-4/07/0008 $ 5.00

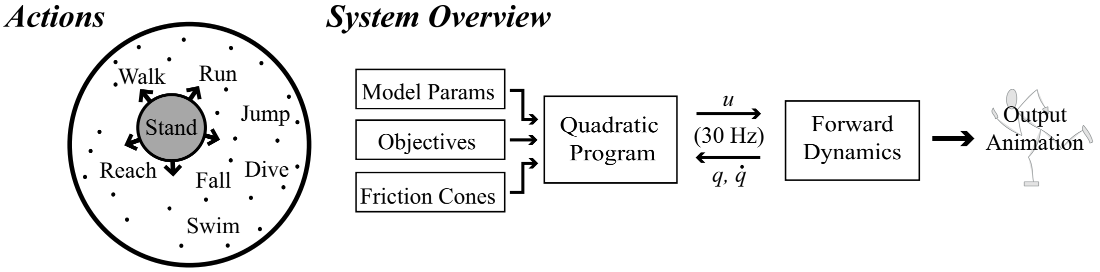

Figure 1: *Previous control systems demonstrate that many human actions can be simulated. Fundamentally, these actions require careful exploitation of external contact forces during periods of sustained contact. We informally refer to these periods as "standing." Multiobjective control ensures robust execution of actions while standing. Given a control strategy and physical properties of the body and environment, our control system uses the current state of the active body* ($\mathbf{q}, \dot{\mathbf{q}}$) *to solve a quadratic program that computes the necessary control torques* $\mathbf{u}$. *This allows us to take a fundamental behavior such as standing and expand its range of application in simulations with different bodies, poses, geometries, and frictional properties.*

objective formulation allows for a compromise between several conflicting motion goals, such as balancing and tracking.

A key component of our approach is a quadratic program (QP) that maximizes instantaneous performance objectives subject to limits on actuation and contact forces. Related formulations have also been proposed in robotic manipulation [CHS88, FOK98, WC06], but current animation techniques still rely on simpler, spring-damper mechanisms. Our experiments reveal that a QP solution improves on the traditional spring-damper techniques by adapting to changing properties of the body, the environment, and the contact. We describe its construction for frictional and non-planar contacts, its use in simulations with large disturbances, and the specifics of accomplishing multiple objectives.

We begin by reviewing the dynamics of active bodies in frictional contact with the environment and highlight the difficulties of controlling such underactuated systems (§2). Next we define multiobjective control with frictional contacts in terms of a local optimization that maximizes instantaneous performance metrics subject to limits on actuation and contact forces (§3). Then we discuss practical strategies needed to accomplish common control objectives in spite of contact variations caused by significant disturbances (§4) and we present lifelike, animations of standing characters in challenging physical environments (§5), all of which were simulated at interactive rates using a standard rigid-body simulator. The results suggest that multiobjective control may be combined with previously proposed control policies for locomotion and other complex behaviors (§6) and used in the design of a new generation of adaptive control systems (§7).

2. Contact Dynamics

Motion of a body in contact with the environment is more complex than unencumbered motion in free space. This is due to the presence of reaction forces that push on the body at each contact point. For the common case of sustained contact, however, control can exploit the linear relationship between joint torques, reaction forces, and joint accelerations. This relationship can be computed at interactive rates and used to control active bodies. In this section we establish our notation by reviewing contact mechanics and the equations of motion for active articulated bodies [CHS88, Wie02].

2.1. Contact Mechanics

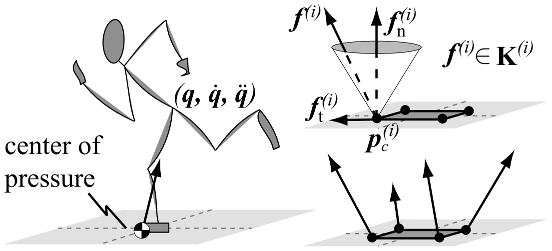

Figure 2: *Contact dynamics expresses the relationship between the motion* $(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}})$ *of an articulated body, its internal torques, and external forces. We model the contact between two surfaces with a set of point contacts* $\mathbf{p}_c^{(1)} \ldots \mathbf{p}_c^{(m)}$ *and the matching contact forces* $\mathbf{f}^{(1)} \ldots \mathbf{f}^{(m)}$. *Each contact force is restricted by a convex cone* $\mathrm{K}^{(i)}$ *according to the standard Coulomb's model of friction.*

Contacts with environment, as shown in Figure 2, restrict the relative velocity of each contact point $\boldsymbol{p}_c^{(i)} \in \mathbb{R}^3$, for $i = 1 \ldots m$. In the case of a non-slipping contact, the relative velocity is zero: $\dot{\boldsymbol{p}}_c^{(i)} = \mathbf{0}$. This condition can also be expressed in terms of joint velocities $\dot{\boldsymbol{q}} \in \mathbb{R}^n$ by using the Jacobian matrix $\boldsymbol{G}^{(i)} \in \mathbb{R}^{3 \times n}$ to compute the body velocity at the point of contact:

$$\boldsymbol{G}^{(i)}\dot{\boldsymbol{q}} = \dot{\boldsymbol{p}}_{\mathrm{c}}^{(i)} = \mathbf{0}. \tag{1}$$

A point contact yields a frictional contact force $\boldsymbol{f}^{(i)} \in \mathbb{R}^3$ that prevents geometric overlap by pushing back on the body. Unlike the forces in a joint linkage (bilateral contact), a contact force does not pull the body in case of separation (unilateral contact) implying that its normal component must be positive: $f_{\mathrm{n}}^{(i)} \geq 0$. Coulomb's model of friction limits the tangential component of the contact force: $\|\boldsymbol{f}_{\mathrm{t}}^{(i)}\| \leq \mu f_{\mathrm{n}}^{(i)}$, where $\mu > 0$ is a coefficient of friction at the contact point. We collect these limits into a friction cone $\mathrm{K}^{(i)}$ that restricts the direction and magnitude of the contact force:

$$\boldsymbol{f}^{(i)} \in \mathrm{K}^{(i)} = \left\{ \mathbf{x} \mid \|\mathbf{x}_{\mathrm{t}}\| \leq \mu \mathbf{x}_{\mathrm{n}} \right\}. \tag{2}$$

By the principle of virtual work, a linear map $\boldsymbol{G}^{\top}\boldsymbol{f}$ determines the total joint torque by aggregating all contact forces and all Jacobian matrices into one vector $\boldsymbol{f} \in \mathbb{R}^{3m}$ and matrix $\boldsymbol{G} \in \mathbb{R}^{3m \times n}$.

2.2. Active Body Dynamics

Conservation of momentum dictates that the total sum of contact forces equals the total change in linear and angular momentum. In the absence of contact forces, it is impossible for an active body to control the location of its center of mass (COM). An active body propels itself using joint torques $\boldsymbol{u} \in \mathbb{R}^{n-6}$. These torques affect only internal joints $\boldsymbol{q}_1 \in \mathbb{R}^{n-6}$ leaving the global position and orientation of the body $\boldsymbol{q}_2 \in \mathbb{R}^6$ as unactuated degrees of freedom. Using this same separation on equations of motion produces two sets of equations, with and without actuation:

$$\boldsymbol{M}_1(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{n}_1(\boldsymbol{q},\dot{\boldsymbol{q}}) + \boldsymbol{G}_1^{\top}(\boldsymbol{q})\boldsymbol{f} = \boldsymbol{u} \tag{3}$$

$$\boldsymbol{M}_2(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{n}_2(\boldsymbol{q},\dot{\boldsymbol{q}}) + \boldsymbol{G}_2^{\top}(\boldsymbol{q})\boldsymbol{f} = \mathbf{0}. \tag{4}$$

The first two terms in both equations combine the inertial and gravitational forces on the body. The two equations summarize the main challenge of active body control with frictional contacts: the dimension of the quantity we need to control $\boldsymbol{q}$ exceeds the dimension of torques $\boldsymbol{u}$ at our disposal. Careful manipulation of contact forces $\boldsymbol{f}$ is the only way to accomplish a specific objective, and yet they are restricted by the friction cone: $\boldsymbol{f} \in \mathrm{K} = \mathrm{K}^{(1)} \times \cdots \times \mathrm{K}^{(m)}$.

3. Multiobjective Control

Our multiobjective control computes the joint torques that drive the motion of an active body in simulation. These joint torques are chosen to try and satisfy several objectives at once. Each objective describes a different facet of the desired motion: one objective may track motion data, another may command the location of the center of mass, and yet a third may force the hands to a specific destination. At each instance in time, the conflicts and trade-offs between different objectives are managed by a fast optimization that respects the dynamics of the current contacts and automatically accounts for the physical properties of the active body. Since speed is a primary requirement of online control, we express all constraints and objectives in a quadratic program that can be solved quickly.

3.1. Optimization

Given the current pose $\boldsymbol{q}$ and velocity $\dot{\boldsymbol{q}}$ for the body, the optimization computes joint torques $\boldsymbol{u}$, joint accelerations $\boldsymbol{a} \in \mathbb{R}^n$, and contact forces $\boldsymbol{f}$ that maximize performance of several objectives $g^{(1)} \dots g^{(\ell)}$:

$$\min_{\boldsymbol{a},\boldsymbol{f},\boldsymbol{u}} \quad \{g^{(1)}, \dots, g^{(\ell)}\}$$

$$\text{subject to} \quad \boldsymbol{M}\boldsymbol{a} + \boldsymbol{n} + \boldsymbol{G}^{\top}\boldsymbol{f} = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \boldsymbol{u} \tag{5a}$$

$$\boldsymbol{f} \in \mathrm{K}, \quad \boldsymbol{u} \in \mathrm{L} \tag{5b}$$

$$\boldsymbol{G}\boldsymbol{a} + \dot{\boldsymbol{G}}\dot{\boldsymbol{q}} = \mathbf{0} \tag{5c}$$

In the above, Equation (5a) restricts the solution to be consistent with the instantaneous contact dynamics of active articulated bodies. This is a linear constraint on the vector unknowns because the remaining quantities $\boldsymbol{M}$, $\boldsymbol{n}$, and $\boldsymbol{G}$ are constant for the current pose and velocity. Equation (5b) limits the contact forces and control torques according to current friction cones Kand constant-bound torque limits L. Lastly, Equation (5c) ensures that accelerations remain compatible with the no-slip contact condition [Bar89] in Equation (1).

3.2. Quadratic Program

Our implementation approximates the general multiobjective formulation with quadratic programming. This requires choosing quadratic objectives whose trade-offs are weighed either by strict prioritization or through a combined weighted-sum objective. We also approximate the nonlinear friction cone constraint with a conservative polygonal approximation, noting that, if needed, interior point methods could also manage the conical convex constraint in its original form [BV04].

3.2.1. Quadratic Objectives

Control strategies include one or more quadratic objectives. For example, objectives can be used to simultaneously track a desired posture while commanding the position of hands and feet. Quadratic objectives regulate the values of such kinematic quantities $\boldsymbol{x}(\boldsymbol{q})$ by choosing their accelerations $\ddot{\boldsymbol{x}}(\boldsymbol{q})$ at each time step.

The value of each objective $g^{(i)}$ measures the difference between the current $\ddot{\boldsymbol{x}}^{(i)}$ and desired $\boldsymbol{d}^{(i)}$ acceleration:

$$g^{(i)} = \|\ddot{\boldsymbol{x}}^{(i)} - \boldsymbol{d}^{(i)}\| = \|\boldsymbol{J}^{(i)}\boldsymbol{a} + \dot{\boldsymbol{J}}^{(i)}\dot{\boldsymbol{q}} - \boldsymbol{d}^{(i)}\|, \tag{6}$$

where the Jacobian matrix $\boldsymbol{J}^{(i)}$ describes the linear relationship between joint velocities and the velocities of regulated kinematic quantities: $\dot{\boldsymbol{x}}^{(i)} = \boldsymbol{J}^{(i)}\dot{\boldsymbol{q}}$.

For example, we can incorporate recorded motion trajectories $\boldsymbol{m}(t)$ by computing the desired accelerations to encourage critically damped tracking:

$$\boldsymbol{d} = k_s\big(\boldsymbol{m}(t) - \boldsymbol{x}\big) + 2\sqrt{k_s}\big(\dot{\boldsymbol{m}}(t) - \dot{\boldsymbol{x}}\big) + \ddot{\boldsymbol{m}}(t), \tag{7}$$

where t is the current simulation time and k_s is the tracking stiffness. We use the same equation to track a single posture by choosing a constant value for $\boldsymbol{m}(t)$, setting $\dot{\boldsymbol{m}}(t)$ and $\ddot{\boldsymbol{m}}(t)$ to zero.

Stiffness values will depend on desired motion. High stiffness will produces animations that follow motion data despite external disturbances. Low stiffness will produces more realistic animations that react to external disturbances. Note that low-stiffness tracking was more difficult to achieve with previous techniques [ZH02, YCP03, YN03].

3.2.2. Control Trade-Offs

Multiobjective control seeks a compromise among various, often conflicting, objectives. One approach to conflict resolution is to identify strict priority levels. A sequence of quadratic programs can then recursively optimize each objective. First, we optimize the most important objective. Next, we constrain its value to the computed optimum in the optimization of the second most important objective. Strict priorities ensure that some objectives (e.g., reaching) are minimized before others (e.g., posture) are even considered. However, our experiments show that strict priorities should not be used for control of standing because balance tasks usually interfere with other tasks such as tracking recorded motions, which leads to less realistic motions.

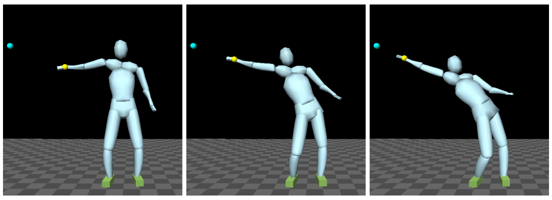

Figure 3: *The weight of a reaching objective is gradually increased, pushing the character to a more precarious stance. In the accompanying video, the reach objective weight is increased to the point where it outweighs the balance objective, and the character falls over. Objective weights from left to right are: 0.01, 0.05, and 0.15.*

We attain a more flexible scheme by using a weighted-sum objective g to strike a compromise between different control objectives:

$$g = w_1 g^{(1)} + w_2 g^{(2)} + \cdots + w_n g^{(n)}. \tag{8}$$

The weights w_i determine the relative importance of each objective $g^{(i)}$ and account for scaling differences in the units of measure. We show the effect of different weights in Figure 3. At first, the reaching objective is given zero weight and the arm does not move. As the importance of the reach increases, the body progressively departs from its balanced stance until the importance of balance is outweighed by the emphasis on reach and the body falls over. This example illustrates that choosing weights is not a burden, but a vital aspect of any control strategy. Balance, for example, may be a top priority for athletes until they have the opportunity to dive for a ball. The weighting in the objective function determines the precise manner in which such motions are accomplished.

4. Practical Standing Control

In simulation, contact between two objects is neither perfectly detected nor perfectly maintained. Numerical errors due to integration can create variations in detected contact points at almost every time step. External disturbances are even more disruptive. Applying multiobjective control in such an environment requires addressing two major challenges. First, we must complement our general theoretical treatment with practical strategies that account for frequent contact variation. Second, we must devise strategies that guide the body to positions from which it is capable of accomplishing control objectives such as standing upright to avoid falling.

4.1. Stabilizing Contacts

Our theoretical model of contact forces assumes that contacts are maintained. However, numerical errors in the dynamics and the kinematics will often create unintentional contact changes even before external disturbances are introduced. When contacts break, the control must adapt or it will fail.

To prevent contacts from breaking in the first place, we ensure that contact forces computed by the control are strictly positive. The friction cone $\mathrm{K}^{(i)}$ is modified so that the magnitude of contact forces $\boldsymbol{f}^{(i)}$ are above a conservative, threshold (> 50 N in our experiments, for a 70 kg character). This directs the QP solution to compute torques that push on each contact point and hence discourage incidental changes, or, in the case of small separation, re-establish the contacts in just a few simulation steps. Higher values yield better contact stabilization, but extremely large thresholds may generate unrealistic motion or infeasible QP problems. The best practice is to scale the threshold with the weight of the character.

We also use conservative estimates of the location of each contact point, $\boldsymbol{p}_{\mathrm{c}}^{(i)}$. First, we define a contact region to be the closed, polygonal surface defined by a set of contact points (e.g., one contact region exist for each foot or hand used for support). The conservative estimate assumes a modified version of the contact region, shrunken to the strict interior of the actual region. This increases the likelihood that an aggregate contact force (i.e., the sum of all contact forces at the contact points surrounding a contact region) originates within the strict interior of the actual contact region. The QP is modified by relocating each contact point, $\boldsymbol{p}_c^{(i)}$, (used in the construction of K) to be closer to the center of the associated contact region. A larger reduction will produce a more conservative solution, but also restrict the possible movement of the character. In our experiments we found that reducing the size of the contact region by 30 percent was a nice compromise, though many different reductions from 0 percent to 90 percent also worked.

Despite these modifications, external disturbances will inevitably cause contacts to break. We measure the scale of a contact disturbance by how far a contact point is from the surface it should contact. For small disturbances ($< 1e-4$ m, for a normal human sized character) we ignore the break, reasoning that the strictly positive contact forces will re-establish contact soon. If the disturbance becomes medium sized ($< 2e-2$ m) we collapse the friction cone $\mathrm{K}^{(i)}$ (setting the coefficient of friction to zero) to disallow tangential contact force and encourage immediate recovery without tangential slipping. However, if the disturbance becomes large ($> 2e-2$ m), we remove the contact point from the QP formulation. In that case, we add a new motion objective that guides the former contact point toward its projection on the external contact surface as a last-ditch attempt to re-establish the contact.

In our experiments, control outcomes were not overly sensitive to precise values of threshold parameters. We chose the reported values to increase the stability of our control for larger disturbances. Most other settings worked well with smaller disturbances.

4.2. Maintaining Balance

Everybody falls on occasion and Equation (4) shows us why: the global position and orientation of a body is not directly controlled by joint torques. Humans adapt to this limitation with both anticipatory and reactive movements. Anticipatory movements, such as bracing for the motion of a bus by leaning in the direction of its motion, generally require sophisticated motion planning. Our low-level control does not handle anticipatory aspects of balance, but it does provide a reactive mechanism to maintain an upright posture while accomplishing motion objectives that are nearly balanced. This is done by regulating the horizontal position of the center of mass.

Once the COM is too far from a desired, upright position, it may become impossible to return to that position due to underactuation. Contact dynamics gives us a precise condition for knowing when a given return trajectory is feasible (§2). The motion of a body $\boldsymbol{q}(t)$ is feasible if and only if there are contact forces $\boldsymbol{f} \in \mathrm{K}$ that satisfy Equation (4):

$$\boldsymbol{M}_2(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{n}_2(\boldsymbol{q},\dot{\boldsymbol{q}}) + \boldsymbol{G}_2^\top(\boldsymbol{q})\boldsymbol{f} = \boldsymbol{0}. \tag{9}$$

This a strict generalization of the often used criteria that the zero-moment point (ZMP) should remain within the support polygon [Wie02]. Whereas the ZMP criteria assumes an infinite coefficient of friction and planar contact with the ground, the above condition accounts for the friction cone and applies to any three-dimensional contact configuration. As with the ZMP criteria, direct application of this condition to plan an optimal recovery from a given dynamic disturbance is beyond the computational budget of online control systems. (The equations are no longer linear because the state of the body changes over time.) For example, efficient implementations of optimal control [SP05] require several seconds of computation time, at least an order of magnitude too slow for online control.

Instead of computing an optimal recovery strategy, our control implements a simple heuristic strategy based on guiding the center of mass toward a more stable configuration for most disturbances. The COM for a human standing on flat ground, for example, is usually above the mid-point between the two footprints. In general, this is a stable configuration for many disturbances because, for small deviations, the COM is fully controllable in the horizontal direction and can be brought back to the desired position.

Since the COM is a kinematic quantity $\boldsymbol{x}(\boldsymbol{q})$, we can use multiobjective control to direct its motion by providing a desired acceleration, $\boldsymbol{d}$, in Equation (6). We found in our experiments that a simple first order damped approach works well for many motions:

$$\boldsymbol{d} = k_s(\boldsymbol{x}_d - \boldsymbol{x}) - k_d\dot{\boldsymbol{x}}, \tag{10}$$

where $\boldsymbol{x}_d$ is the desired horizontal position of the COM and k_s and k_d are manually tuned constants. However, to recover from larger deviations from the desired position we found that this simple strategy was not sufficient. For high values of k_s the strategy tended to be overly forceful, resulting in instability. For lower values of k_s the strategy would fail to reach the goal. Instead, we found that we could achieve better results by varying k_s with the distance from the desired COM position. In our experiments we scaled k_s with the inverse square-root distance: $1/\sqrt{||\boldsymbol{x}_d - \boldsymbol{x}||}$.

When standing on uneven ground with one foot on the pedestal or with one hand on the wall, we still control the COM in the plane perpendicular to the force of gravity. The desired location of the COM is placed near the mid-point of the horizontal projection of all contact points. This ensures that for small deviations from the desired position, the COM is still controllable in the horizontal direction and can be returned to the desired position. The exact location is not important, but different choices will affect the ability to recover from different types of disturbances. For example, a character that places more weight on a forward foot, will be able to recover from larger unexpected forces coming from the front.

The key to understanding our strategy is to observe that it does not prevent falling on its own: the COM can fall to the ground and still be above the horizontal position. Instead, falls are prevented with a combination of this objective and others that prescribe *standing* motions or postures. When the COM wanders significantly outside the support polygon the character falls because it can no longer accomplish the objective of standing. However, we found that our simple strategies worked well even for many significant disturbances.

5. Results

We demonstrate the capabilities of multiobjective control by discussing a few simulations we have created of active characters balancing, tracking motion, and responding to a dynamic disturbances.

5.1. Experiments

The supplementary video includes a few typical runs from the following experiments:

Sobriety. A human-like character balances on a moving platform while reaching for its nose with its hands. Both balancing and reaching are accomplished despite the significant motion of the platform and other disturbances introduced by a human, interactively, during the simulation.

Pelted. A human-like character tracks motion data closely while balancing in response to many collisions with objects. As with all of these simulations, the character balances under its own power and there are no artificial aids preventing the character from falling over.

Platform. A human-like character balances on a moving platform. Varying the coefficient of friction between the feet and the platform cause changes in the balancing motion.

Alien. The human character in the "Pelted" and "Platform" simulations is replaced with a lighter and smaller character. Although the geometry, weight, and proportions differ from those of the human character, the same control works without modification. Since the character has a larger head, however, the posture stiffness of the waist joint was decreased to encourage more waist motion.

Wall. A human character places its hand on a nearby wall for additional support while balancing on a moving platform. The control automatically adapts the strategy of using the hand to provide additional leverage to maintain balance despite severe tipping of the platform. The control adapts easily to the non-planar contact configurations, involving both feet and hands, through the use of friction cone constraints.

Mishap. The character stands with one leg perched on a flimsy table. Friction cone constraints prevent the character from immediately toppling the table. When the table suddenly collapses, the character regains its balance on one foot. The balancing maneuver occurs automatically, though we do direct the free foot to a desired location using an end-effector objective.

We manually modeled the geometry of both characters in our simulations. Their inertial properties were computed automatically using the volume of each limb and standard mass distributions [Win90]. The motions tracked by our control system were recorded with an optical motion capture system. Forward dynamics with frictional contacts were computed with the Open Dynamics Engine (*www.ode.org*), a general purpose rigid body simulator. The QP problems were solved by the MOSEK software system (*www.mosek.com*), which employs the interior point method to solve convex optimization problems [BV04].

The QP control problem is solved 30 times per second of simulation, while we use many more simulation steps in the same interval, between 1000 and 5000. Each solution required around 15 iterations to converge for an average running time of 17 milliseconds. The "Wall" simulation took slightly longer than the others (see Table 1) because of the additional hand contact. All simulations were fast enough to allow the entire system (simulation and control) to run at 30 frames per second, or better, on a 2.8 GHz Intel Pentium 4.

5.2. Direction

Our experiments demonstrate that multiobjective control enables artistic direction of active bodies with two familiar animation mechanisms: control of poses and end effector positions.

In most of our experiments, we track a single recorded posture, but tracking motions is just as easy. Tracking fast motions, such as dodging incoming objects, is accomplished accurately, but "loosely" enough to respond interestingly to collisions with other objects (Pelted).

End-effector objectives are used to control individual limbs: arms, hands, feet, and so on. Our experiments include two simple examples. One directs hands to touch the nose (Sobriety) and the other controls the swing leg to direct the look of a balancing maneuver (Mishap). (§4.2). In both cases, the accelerations of the end-effectors where simply chosen to converge toward a desired goal position. Though tracking of more complex trajectories is possible.

5.3. Friction Cones

Many control systems assume planar contact with flat ground, which limits possible applications. Multiobjective control manages this special case (Pelted, Alien) but it also handles more general contact configurations by using friction cone constraints. Examples of this include uneven footing (Sobriety, Mishap) and hand contact with the wall (Wall).

Since friction cones prevent feet from slipping, the control need not insist on a perfect match between the body and recorded postures. A shorter character, for example, can easily track the recorded trajectory of a full-size human (Alien) without slipping causing a fall.

Friction cones can also be manipulated to create interesting animation. By restricting the allowable friction cones, different balancing motion that respect those limitations emerge automatically(Platform). Friction cones were also instrumental in the Mishap simulation where a very narrow friction cone was used initially on the front foot to instruct the character to apply only vertical forces on the precarious table.

5.4. Balance Objective

Many of our experiments feature a character on a moving platform. Under such conditions, our control system maintains balance by coaxing the COM back toward a conservatively chosen position. We emphasize that treating the control of the COM as a strict priority, above all other objectives, has not produced satisfactory results in our experience. Instead, the corrective motion is weighted against other active objectives, contributing to the quality of the motion. Although our simple balance strategy can be improved with further work, multiobjective control can easily accommodate new strategies once they are available.

5.5. Adaptation

Multiobjective control adapts automatically to external disturbances and physical properties of the character and the environment.

The allure of physically based animation is clearly demonstrated by a rich diversity of interactions characters can have with their environment. However, this is only possible if characters can adapt naturally to physical disturbances. We present a couple of examples (Pelted, Mishap) that suggest definite progress in this direction. Complex motions, including natural but counter-intuitive balance recoveries, such as lunging in the direction of the fall (Mishap), emerge without explicit modeling.

Multiobjective control also adapt to bodies with different inertial parameters (Alien). Our shorter character is capable of withstanding significant disturbances by using a general control strategy initially tested on a taller and heavier character. This highlights a key advantage of our control system: it decouples the description of control strategies from the computation of required torques. Hence, the objectives are independent of mass distribution, model geometry, and contact dynamics.

Simulation	**Vars.**	**Avg. QP Time**	**Avg. Iterations**
Platform	140	13ms	14.5
Pelted	140	13ms	14.5
Sobriety	146	16ms	13.8
Mishap	143	14ms	15.8
Wall	172	29ms	17.7

Table 1: *The number of variables, average optimization time, and average number of iterations for the multiobjective QP per simulation.*

6. Related Work

Active body control has a long history in graphics and robotics. We review some of this work in the context of our main design decisions: the choice of local optimization for control, the form of the quadratic objective function, and the choice of constraints for modeling general frictional contacts.

6.1. Active Body Control

Multiobjective optimizations complement strengths of previous approaches in computer animation but reduces the need for manual adjustment of control parameters. Raibert and Hodgins [RH91] relied on spring-damper mechanisms to compute torques for online control, leading to some of the most dramatic simulations of active bodies [HWBO95, Woo98, FvdPT01]. Similar control strategies now appear in commercial software systems for computer animation (*www.naturalmotion.com*). The specifics of these commercial systems are unknown, but they likely require tuning of individual rest lengths and spring constants for most characters, tasks, and simulation environments, similar to the predecessors [HP97, FvdPT01].

Manual tuning can be reduced to some extent with dynamic scaling laws and automated search, but it reaches its limits when adapting to new environments, mass distributions, and other variations [HP97]. Our control system enables modular specification of general control policies for the case of sustained frictional contact. Instead of designing and tuning spring-based dampers for each joint, we divorce the specification of control policies from the computation of required control torques. Hence, our control system adjusts more easily to new situations and different environments.

Limit-cycle control tracks periodic motions by computing control perturbations needed to return the present motion back to the desired limit cycle (i.e. limit-cycle control strategy or periodic motion data) [LvdPF96]. Instead of relying on explicit models of contact dynamics, it approximates the Poincaré return map. The advantage of such an approach is that it also incorporates the effect of collisions into execution of control policies: a difficult problem that we do not address in this paper. A potential liability is that approximated return maps might be valid only for small perturbations from limit-cycle trajectories.

Other tracking alternatives have also been proposed to create dynamically responsive motions from kinematic or preplanned trajectories [ZH02, YCP03, ZMCF05]. These approaches scale spring constants by inertial parameters or feed-forward torques magnitudes to reduce the difficulty of tuning parameters, but none explicitly account for contact dynamics. Some techniques have pursued a hybrid alternative instead, accounting for some dynamic parameters with the goal of generating dynamically feasible *motions* instead of controls [GM85, SC89, Pai90, SC92, YN03, KP06]. Our multiobjective control was designed to integrate with any general purpose simulator of rigid bodies. As a result, we can easily animate complex interactions with many moving objects using any rigid-body simulator.

6.2. Quadratic Objective

Our multiobjective approach was inspired by prioritized control of articulated bodies [KSPW04]. The principal advantage of such an approach is automatic coordination of multiple objectives, which makes it easy to combine task-specific objectives with the less specific postural objectives gleaned from motion data [AP06]. Our experimentations, however, showed that prioritization of balance control interferes with posture tracking, which makes it difficult to combine the two in a life-like manner. Instead, we rely on the quadratic weighted-sum objective in all of our experiments to attain necessary tradeoffs between tracking, balance, and other tasks.

Earlier applications of local optimization proposed similar quadratic objectives without incorporating a general model of contact dynamics [SC89, SC92]. Inaccurate modeling of contact dynamics is also a potential drawback of most prioritized control systems: they often assume the existence of bilateral contact constraints, as if the bodies were pinned at contact points. As illustrated in Figure 4, bilateral contacts lead to unrealistic control strategies. Recent developments in prioritized control suggest an iterative active-set solution [SK06], but this approach is less robust and harder to implement than our QP-based method.

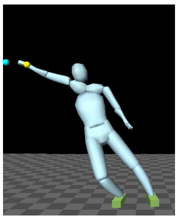

Figure 4: *This illustration underscores the importance of incorporating ground contact constraints into any control formulation. Ignoring contact dynamics, a character can reach for the object as if his feet were pinned to the ground. With proper contact dynamics and multiobjective control, the character strikes a compromise between reaching and not falling as seen in Figure 3.*

6.3. Friction Cones

Friction-cone constraints generalize the zero-moment point (ZMP) constraint, which is often used in local optimizations as an alternative to bilateral contact constraints [HMPH04, KKI02]. The ZMP is a criterion of physical feasibility for bodies in contact with the ground plane [VB04]. For example, its position outside the contact polygon indicates a physically infeasible motion. The ZMP criterion is sometimes incorrectly defined as a measure of dynamic stability in both graphics and robotics literature. Instead, the ZMP criterion enables successful tracking of controllable trajectories by ensuring physically realizable control policies [SKG03, HHHT98]. Hofmann and colleagues, for example, use quadratic programming to restrict the ZMP to remain within the contact polygon [HMPH04]. This approach assumes planar contact configurations (e.g., standing on flat ground) with infinite friction. For example, control systems based on ZMP constraints may lead to slipping and falling in simulations with realistic frictional properties. In contrast, we use friction-cone constraints that are more accurate and valid for general three-dimensional (e.g., standing on uneven ground) contacts with friction [Wie02].

According to a survey by Srinivasa [Sri05], the first control system with an explicit model of contact dynamics appeared in the robotics literature as a solution to multi-fingered manipulation of two-dimensional objects [CHS88]. The control systems proposed in graphics literature, however, did not employ explicit formulations of contact dynamics until Fang and Pollard [FP03] demonstrated their value in *offline* optimal control. We demonstrate the feasibility and importance of this model for *online* control in interactive animations of active bodies.

Two methods in robotics literature have relied on a similar QP formulations to ours for the control of walking bipeds [FOK98, WC06], but without addressing contact variations and significant disturbances. Our work also defines the concept of multiobjective control with frictional contacts and examines its role in animations of standing active bodies. We emphasize the resilient treatment of disturbances, reasoning that locomotion and more complex behaviors can be robust only after standing is more robust.

7. Conclusion

We have presented a multiobjective control formulation based upon local optimization that models and respects frictional contact dynamics. We arrive at three important conclusions about such systems. First, friction-cone constraints, which generalize ZMP constraints, improve control of active bodies in simulation with arbitrary frictional contacts. Second, the compromise between multiple conflicting motion objectives should be accomplished with soft trade-offs rather than strict priorities, especially when actively regulating the center of mass. And third, special care must be taken to stabilize frictional contacts in simulations with large dynamic disturbances.

Our multiobjective control is restricted in scope: it assumes the existence of a proper high-level control policy instead of searching for one. The more difficult problem has been tackled *offline* with many approaches including continuous optimization [WK88], dynamic programming [vdPFV90], genetic algorithms [Sim94, NM93], neural networks [GTH98], and simulated annealing [vdPF93, GT95]. In the future we plan to apply similar ideas to online control by solving appropriate approximations of the offline problem. Another possible direction would be to complements multiobjective control with kinematic methods that ignore physics but learn from data and other studies of natural motion [KKKL94, RSC01, GMHP04, YKH04]. These would be the next steps toward designing adaptive control strategies for walking, running, jumping, and other complex behaviors.

8. Acknowledgments

We are grateful to Eugene Hsu and Daniel Vlasic for motion-capture assistance and acting; Tom Buehler for video production; and Russ Tedrake for comments throughout the project. This work was supported in part by the Singapore-MIT GAMBIT Game Lab.

References

[AP06] Abe Y., Popović J.: Interactive animation of dynamic manipulation. In *Symposium on Computer Animation (SCA)* (2006), pp. 195–203. 8

[Bar89] Baraff D.: Analytical methods for dynamic simulation of non-penetrating rigid bodies. In *Computer Graphics (Proceedings of SIGGRAPH 89)* (July 1989), Annual Conference Series, ACM SIGGRAPH, pp. 223–232. 3

[BV04] Boyd S., Vandenberghe L.: *Convex Optimization*. Cambridge University Press, 2004. 3, 6

[CHS88] Cole A., Hauser J., Sastry S.: Kinematics and control of multifingered hands with rolling contact. In *International Conference on Robotics and Automation (ICRA)* (1988), vol. 1, IEEE, pp. 228–233. 2, 8

[FOK98] Fujimoto Y., Obata S., Kawamura A.: Robust biped walking with active interaction control between foot and ground. In *International Conference on Robotics and Automation (ICRA)* (1998), IEEE, pp. 2030–2035. 2, 8

[FP03] Fang A. C., Pollard N. S.: Efficient synthesis of physically valid human motion. *ACM Transactions on Graphics 22*, 3 (July 2003), 417–426. 8

[FvdPT01] Faloutsos P., van de Panne M., Terzopoulos D.: Composable controllers for physics-based character animation. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Annual Conference Series, pp. 251–260. 7

[GM85] Girard M., Maciejewski A. A.: Computational modeling for the computer animation of legged figures. In *Computer Graphics (Proceedings of SIGGRAPH 85)* (July 1985), vol. 19, pp. 263–270. 8

[GMHP04] Grochow K., Martin S. L., Hertzmann A., Popović Z.: Style-based inverse kinematics. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 522–531. 9

[GT95] Grzeszczuk R., Terzopoulos D.: Automated learning of muscle-actuated locomotion through control abstraction. In *Proceedings of SIGGRAPH 95* (Aug. 1995), Annual Conference Series, pp. 63–70. 9

[GTH98] Grzeszczuk R., Terzopoulos D., Hinton G.: NeuroAnimator: Fast neural network emulation and control of physics-based models. In *Proceedings of SIGGRAPH 98* (July 1998), Annual Conference Series, pp. 9–20. 9

[HHHT98] Hirai K., Hirose M., Haikawa Y., Takenaka T.: The development of honda humanoid robot. In *International Conference on Robotics and Automation (ICRA)* (1998), IEEE, pp. 1321–1326. 8

[HMPH04] Hofmann A., Massaquoi S., Popovic M., Herr H.: A sliding controller for bipedal balancing using integrated movement of contact and non-contact limbs. In *International Conference on Intelligent Robots and Systems (IROS)* (2004), vol. 2, IEEE/RSJ, pp. 1952–1959. 8

[HP97] Hodgins J. K., Pollard N. S.: Adapting simulated behaviors for new characters. In *Proceedings of SIGGRAPH 97* (Aug. 1997), Computer Graphics Proceedings, Annual Conference Series, pp. 153–162. 7

[HWBO95] Hodgins J. K., Wooten W. L., Brogan D. C., O'Brien J. F.: Animating human athletics. In *Proceedings of ACM SIGGRAPH 95* (Aug. 1995), Annual Conference Series, pp. 71–78. 1, 7

[KKI02] Kudoh S., Komura T., Ikeuchi K.: The dynamic postural adjustment with the quadratic programming method. In *International Conference on Intelligent Robots and Systems (IROS)* (2002), pp. 2563–2568. 8

[KKKL94] Koga Y., Kondo K., Kuffner J., Latombe J.-C.: Planning motions with intentions. In *Proceedings of SIGGRAPH 94* (July 1994), Computer Graphics Proceedings, Annual Conference Series, pp. 395–408. 9

[KP06] Kry P. G., Pai D. K.: Interaction capture and synthesis. *ACM Transactions on Graphics 25*, 3 (July 2006), 872–880. 8

[KSPW04] Khatib O., Sentis L., Park J.-H., Warren J.: Whole body dynamic behavior and control of human-like robots. *International Journal of Humanoid Robotics 1*, 1 (2004), 29–43. 8

[LvdPF96] LASZLO J. F., VAN DE PANNE M., FIUME E. L.: Limit cycle control and its application to the animation of balancing and walking. In *Proceedings of SIGGRAPH 96* (Aug. 1996), Annual Conference Series, pp. 155–162. 7

[NM93] NGO J. T., MARKS J.: Spacetime constraints revisited. In *Proceedings of ACM SIGGRAPH 2000* (Aug. 1993), Annual Conference Series, pp. 343–350. 9

[Pai90] PAI D. K.: *Programming Anthropoid Walking: Control and Simulation*. Tech. Rep. TR–90–1178, Cornell University, 1990. 8

[RH91] RAIBERT M. H., HODGINS J. K.: Animation of dynamic legged locomotion. In *Computer Graphics (Proceedings of SIGGRAPH 91)* (July 1991), Annual Conference Series, ACM SIGGRAPH, pp. 349–358. 7

[RSC01] ROSE C. F., SLOAN P.-P. J., COHEN M. F.: Artist-directed inverse-kinematics using radial basis function interpolation. *Computer Graphics Forum 20*, 3 (2001), 239–250. 9

[SC89] STEWART A. J., CREMER J. F.: Algorithmic control of walking. In *International Conference on Robotics and Automation (ICRA)* (1989), IEEE, pp. 1598–1603. 8

[SC92] STEWART A. J., CREMER J. F.: Beyond keyframing: an algorithmic approach to animation. In *Graphics Interface* (1992), pp. 273–281. 8

[Sim94] SIMS K.: Evolving virtual creatures. In *Proceedings of SIGGRAPH 94* (July 1994), Computer Graphics Proceedings, Annual Conference Series, pp. 15–22. 9

[SK06] SENTIS L., KHATIB O.: A whole-body control framework for humanoids operating in human environments. In *International Conference on Robotics and Automation (ICRA)* (2006), IEEE, pp. 2641–2648. 8

[SKG03] SHIN H. J., KOVAR L., GLEICHER M.: Physical touchup of human motions. In *Proceedings 11th Pacific Conference on Computer Graphics and Applications* (2003), pp. 194–203. 8

[SP05] SULEJMANPASIĆ A., POPOVIĆ J.: Adaptation of performed ballistic motion. *ACM Transactions on Graphics 24*, 1 (Jan. 2005), 165–179. 5

[Sri05] SRINIVASA S.: *Control synthesis for dynamic contact manipulation*. PhD thesis, Carnegie Mellon University, 2005. 8

[VB04] VUKOBRATOVIC M., BOROVAC B.: Zero-moment point – thirty five years of its life. *International Journal of Human Robotics 11*, 1 (2004), 157–173. 8

[vdPF93] VAN DE PANNE M., FIUME E.: Sensor-actuator networks. In *Computer Graphics (Proceedings of SIGGRAPH 93)* (Aug. 1993), Annual Conference Series, ACM SIGGRAPH, pp. 335–342. 9

[vdPFV90] VAN DE PANNE M., FIUME E., VRANESIC Z.: Reusable motion synthesis using state-space controllers. In *Computer Graphics (Proceedings of SIGGRAPH 90)* (Aug. 1990), Annual Conference Series, ACM SIGGRAPH, pp. 225–234. 9

[WC06] WIEBER P.-B., CHEVALLEREAU C.: Online adaptation of reference trajectories for the control of walking systems. *Robotics and Autonomous Systems 54*, 7 (July 2006), 559–566. 2, 8

[WH96] WOOTEN W. L., HODGINS J. K.: Animation of human diving. *Comput. Graph. Forum 15*, 1 (1996), 3–13. 1

[Wie02] WIEBER P. B.: On the stability of walking systems. In *International Workshop on Humanoid and Human Friendly Robotics* (2002). 2, 5, 8

[Win90] WINTER D. A.: *Biomechanics and Motor Control of Human Movement*, 2nd ed. John Wiley and Sons, Inc., New York, 1990. 6

[WK88] WITKIN A., KASS M.: Spacetime constraints. In *Computer Graphics (Proceedings of SIGGRAPH 88)* (Aug. 1988), vol. 22, pp. 159–168. 9

[Woo98] WOOTEN W.: *Simulation of Leaping, Tumbling, Landing, and Balancing Humans*. PhD thesis, Georgia Institute of Technology, 1998. 7

[YCP03] YIN K., CLINE M., PAI D. K.: Motion perturbation based on simple neuromotor control models. In *Pacific Conference on Computer Graphics and Applications (PG)* (2003), pp. 445–449. 4, 8

[YKH04] YAMANE K., KUFFNER J. J., HODGINS J. K.: Synthesizing animations of human manipulation tasks. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 532–539. 9

[YLS04] YANG P.-F., LASZLO J., SINGH K.: Layered dynamic control for interactive character swimming. In *Symposium on Computer Animation (SCA)* (2004), ACM SIGGRAPH/Eurographics, pp. 39–47. 1

[YN03] YAMANE K., NAKAMURA Y.: Dynamics filter—concept and implementation of online motion generator for human figures. *IEEE Transactions of Robotics and Automation 19*, 3 (2003), 421–432. 4, 8

[ZH02] ZORDAN V. B., HODGINS J. K.: Motion capture-driven simulations that hit and react. In *Symposium on Computer Animation (SCA)* (July 2002), pp. 89–96. 4, 8

[ZMCF05] ZORDAN V. B., MAJKOWSKA A., CHIU B., FAST M.: Dynamic response for motion capture animation. *ACM Transactions on Graphics 24*, 3 (Aug. 2005), 697–701. 8