

Simulation of Human Motion Data using Short-Horizon Model-Predictive Control

M. da Silva, Y. Abe, and J. Popović

Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Abstract

Many data-driven animation techniques are capable of producing high quality motions of human characters. Few techniques, however, are capable of generating motions that are consistent with physically simulated environments. Physically simulated characters, in contrast, are automatically consistent with the environment, but their motions are often unnatural because they are difficult to control. We present a model-predictive controller that yields natural motions by guiding simulated humans toward real motion data. During simulation, the predictive component of the controller solves a quadratic program to compute the forces for a short window of time into the future. These forces are then applied by a low-gain proportional-derivative component, which makes minor adjustments until the next planning cycle. The controller is fast enough for interactive systems such as games and training simulations. It requires no precomputation and little manual tuning. The controller is resilient to mismatches between the character dynamics and the input motion, which allows it to track motion capture data even where the real dynamics are not known precisely. The same principled formulation can generate natural walks, runs, and jumps in a number of different physically simulated surroundings.

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

1. Introduction

Many data-driven animation techniques are capable of producing high quality motions of human characters. These approaches extend the usefulness of captured motions by allowing applications to adapt existing motions to meet different needs. Applications can create motions that satisfy new user constraints while maintaining the input motion style [AFO03, KGP02] or exhibit new styles while preserving content [HPP05]. Interactive applications such as games can respond to user input and synthesize new results in real-time.

Few techniques, however, are capable of generating motions that are consistent with physically simulated environments. The implicit assumption made by all kinematic synthesis approaches is that the performance environment is the same as the capture environment. This assumption is invalid when motions are performed in physically simulated environments. In a physical simulation, the character can en-

counter new or unpredictable circumstances such as being hit by a ball or standing on a shaky platform. Ignoring these interactions leads to physically inconsistent motion.

In contrast, physically simulated character motions are automatically consistent with the environment but are often unnatural because they are difficult to control. Recorded motions provide an intuitive control specification but simulating any such motion remains a difficult problem. Human characters, in particular, have many degrees of freedom (dofs) subject to non-smooth, non-linear dynamics. This makes it hard to find the forces that reproduce a desired motion, particularly in new environments.

We present a controller, McSim (motion capture in simulation), that yields natural motions by guiding simulated humans toward real motion data. McSim can be categorized as an instance of model-predictive control (MPC). In MPC, the controller predicts a control signal that achieves a desired change in system state based on the current system state




Figure 1: An overview of McSim’s design. At 40 to 100 Hz, the predictive model solves a quadratic program for the joint and external forces that track the reference motion. These forces are adjusted by a low gain proportional derivative component that computes feedback forces at every simulation time step (1-10kHz). Our controller can work with any black-box simulator.

and a model of the system’s dynamics. Our controller uses a predictive component (§4) based on a linearized model of linked rigid body and contact dynamics. The linear dynamics model is used as a constraint in a quadratic program (QP) that solves for the joint and external forces that track the provided input motion for a short window of time into the future.

McSim combines the predictive component with a low gain proportional-derivative (PD) component (§5) as depicted in Figure 1. The predictive component’s control has errors due to high latency and modeling assumptions. The PD component compensates for these errors. The PD component also provides a low-latency response to unexpected perturbations. For certain motions, robustness can be further improved by adapting the input motion according to heuristic feedback rules (§6).

McSim is fast enough for application in interactive systems such as games and training simulations. It can adapt to differences between the character dynamics and the input motion allowing it to track motion capture where the character model can only be estimated. With no precomputation and little manual tuning, McSim is able to produce walking, running, and jumping motions similar to the reference motion while also adapting to new physical surroundings (§7) at interactive rates.

2. Related Work

Most prior online control techniques in the graphics literature have been based on manually designed PD controllers [Rai86, RH91, HWBO95, FvdPT01, YLvdP07, SKL07]. These approaches are typically sensitive to gain parameters and not intuitively directed. In contrast, off-line authoring tools based on continuous optimization leverage the benefit of time to search for physical motions that are optimal according to some metric and satisfy user constraints [WK88, Coh92, PW99, FP03, SP05]. The predictive component of McSim is inspired by these off-line approaches but

sacrifices global optimality for computation speed by restricting the search to a short amount of time into the future. The predictive component of our controller allows the PD component to use relatively small gain parameters, resulting in more stable simulations and more natural motion [YCP03].

McSim can be guided by an arbitrary input motion. Recently, both off-line and online physically based character animation have used data to produce life-like animations, though the role of data differs for each approach. Since the goal of an off-line approach is to produce a new motion with new content, data is used to restrict the search space of possible solutions [SHP04], to model simplified equations of motion [BJ05, TLP06], and to learn parameters of motion style [LHP05]. In online control, the goal is often to simply track a provided input motion in a dynamically simulated environment [ZH02]. Recent approaches, however, have been limited to special cases of motion such as cyclic motions [YLvdP07] or standing [AdSP07, ZH02]. Our approach can track arbitrary motions exhibiting stylistic variations and transitions such as walking to standing.

Many recent approaches to tracking motion data find approximately optimal control policies using off-line precomputation methods such as feedback error learning or simplex methods [SKL07, SvdP05, vdPL95, YLvdP07]. However, these global search methods are not easily applicable to 3D animation where the number of dofs is large. While McSim could incorporate a precomputed feedforward control signal, it produces plausible motions without precomputation. This enables it to be coupled with kinematic motion synthesis techniques [MK05, MP07] to track newly created 2D or 3D motions at run-time.

Among instantaneous optimization approaches, McSim is closely related to Multiobjective Control [AdSP07]. McSim adds the ability to track motions where the contact state changes regularly as in locomotion. Furthermore, we illustrate how the input motion can be modified to improve tracking performance. There are many previous approaches from robotics that propose some form of optimization over a short time horizon to achieve a motion objective [FOK98, WC06, HMPH04, Wie02], each with key differences in the details. In this paper, we propose an alternative formulation of the tracking problem that is capable of handling arbitrary motions and couple it with robust low-latency feedback mechanisms. Others have argued that this form of control is employed by biological systems [YCP03].

3. Method Overview

McSim’s design is guided by three goals. The output motion should be directed by specifying any input motion. It needs to work at interactive rates without requiring expensive precomputation. Finally, it has to work with existing black-box simulators. We would like our controller to work as a plug-in

module with any simulator without any modification to the simulator itself. Achieving these objectives would make the system suitable for tracking kinematically specified motions in interactive applications such as games and training simulations. In the following sections, we describe how McSim achieves these three goals.

At each time step, t , McSim computes a control signal of the form:

$$u(t, x, x_r(t)) = u_f(t, x, x_r(t)) + u_b(t, x, x_r(t)) \quad (1)$$

where u is the control signal, x is the current system state consisting of joint values and velocities, $[q, \dot{q}]$, and x_r is the desired state. The total control signal consists of the predictive component's signal, u_f , added to the PD component's signal, u_b . A predictive dynamics model computes u_f . The PD controller computes u_b which provides low-latency feedback to deal with unexpected perturbations. Stability is achieved by tracking the velocity of the root of the character and modifying the reference motion, x_r as described in Section 6.

4. The Predictive Component

The predictive component's task is to track the reference motion, $x_r(t)$. A long-horizon approach to tracking the reference motion would solve a single optimization for the control forces exerted over the entire motion [WK88]. For human motions, this form of tracking is a high-dimensional, non-linear, non-convex minimization problem. This makes an exact solution impractical at interactive rates. Furthermore, in interactive applications, long-horizon optimal plans are quickly invalidated by changes in the dynamic environment. Rather than plan optimally for situations that may never come to pass, we plan over a small interval into the future using a linearized dynamics model and re-plan at regular intervals, incorporating changes in system state. We call this form of the problem, short-horizon tracking.

4.1. Short-Horizon Tracking

In a physical simulation, a character's motion is determined by integrating a dynamical system forward in time from some initial configuration,

$$x(T) = x(0) + \int_0^T \dot{x}(t) dt. \quad (2)$$

For an active character modeled as a system of linked rigid bodies with actuators between each joint, the equations of motion depend on the current state, $x(t)$, the control signal, $u(t)$, and the external forces, $u_c(t)$. The precise equations can be derived from classical mechanics [FO00] but are summarized here as

$$\dot{x}(t) = f(x(t), u(t) + u_c(t)). \quad (3)$$

A motion that perfectly tracks the reference satisfies

$\ddot{q}(t) = \ddot{q}_r(t)$ for all t , where \ddot{q}_r is the acceleration of the reference motion. The predictive component computes a u_f that tries to reproduce the reference acceleration over a window of size h . In practice, it is usually not possible to achieve the reference acceleration, \ddot{q}_r , exactly due to dynamics constraints of the character and environmental disturbances. As a result, the simulated motion will drift from the reference motion. To correct this drift, feedback terms are added to the reference acceleration to form the desired acceleration, \ddot{q}_d , as described in the next section. Once the desired acceleration is known, a constrained optimization is solved for the joint torques and external forces that achieve it.

4.2. The Desired Acceleration

The desired acceleration consists of the reference acceleration and a correction term. It is computed separately for each joint.

$$\ddot{q}_d = \ddot{q}_r + k_{os}d(q_r, q) + k_{od}(\dot{q}_r - \dot{q}). \quad (4)$$

The correction terms act as a damped feedback acceleration on any errors that occur. The function d compares the current joint configuration, q , to the reference configuration, q_r , and computes an angular acceleration that will move q closer to q_r . The scale of this acceleration is determined by the gain parameter, k_{os} . For rotational joints with one degree of freedom (dof), known as pin joints, $d_i(a, b) = a - b$. Three dof joints, known as ball joints, are represented using quaternions. In this case, $d_i(a, b) = \text{vec}_i(a^{-1} \cdot b)$ where \cdot represents quaternion multiplication and vec_i maps the quaternion to the equivalent axis-angle rotation's i 'th component. The last term in Equation 4 corrects for errors with respect to the reference velocity \dot{q}_r obtained from the motion capture data.

With the exception of the root translation, all desired accelerations are computed using the same values of k_{os} and k_{od} . If $k_{os} = c$, then $k_{od} = 2\sqrt{c}$. Errors in the current position of the root are ignored when computing the desired acceleration of the root. Thus, for the root translation, $k_{os} = 0$. This prevents the controller from trying to correct for errors that are unavoidable due to the environment such as the character walking down hill. The velocity gain is not zero, however. This feedback uses the same gain as the other joints, $k_{od} = 2\sqrt{c}$. The controller is fairly insensitive to the particular value of c chosen as shown in section 7.

4.3. Dynamics Constraints

Computing the control input u needed to achieve the desired acceleration just described would be easy if we could simply invert Equation 3. Unfortunately, humans and animals have more degrees of freedom than forces to control them. Simple inverse dynamics algorithms such as those used for robotic arms rely on being rooted to the environment. Humans, however, are not rooted to the ground. They can use their feet

to push, but not pull, on the ground. They must manipulate these unilateral contact constraints while respecting frictional limits to effect their overall motion [Wie02, AdSP07]. These contact constraints are a key component of the dynamic model used by the predictive component.




Figure 2: A friction cone in 2D. Legal contact forces lie within the cone which can be represented using a linear basis. Non-negative combinations of the basis vectors yield forces in the cone.

Contact forces are computed using a polygonal approximation to Coulomb’s model of friction [FP03]. The model is depicted in 2D in Figure 2. Legal contact forces lie within a friction cone at each corner of the foot in contact with the ground. The cone is oriented normal to the contacting surface with a swept angle determined by the coefficient of friction. In 3D, we use a polygonal approximation (4 facets) to this cone which can be described with a linear basis, V . Contact forces are equal to $V\lambda$ with $\lambda \geq 0$. The non-negative bound on λ insures that the ground reaction force resides within the approximation to the friction cone and prevents contacting bodies from pulling on each other. The i th contact force induces generalized torques on the character which are calculated as $J_i^T V_i \lambda_i$ where J is the gradient of the contact point with respect to the joint configuration of the character. The total contact force on the character, then, is $u_c = \sum_i J_i^T V_i \lambda_i$.

For this static contact model to hold, the contact forces must act only on contact points with zero acceleration [Bar89]. This is known as the no-slip condition:

$$J_i \ddot{q} + \dot{J}_i \dot{q} = 0. \quad (5)$$

In addition to constraints on possible contact forces, achievable accelerations are constrained by the dynamics of the character. Since the predictive component plans over a short-horizon, the dynamics of the character are described by a linear relationship between applied forces and resulting accelerations:

$$\ddot{q}(t) = f(q(t), \dot{q}(t), 0) + W(u + u_c) \quad (6)$$

where W is the gradient of f with respect to the control input. Note that the internal torques, u , are limited by bounds on the strength of the character’s actuators, $u_f \in U$, further restricting possible accelerations.

4.4. Quadratic Programming Optimization

Given all of these constraints, we can now formulate an optimization problem that solves for the joint and external forces, u_f and $J_i^T V_i \lambda_i$, that best achieve the desired acceleration, \ddot{q}_d :

$$\min_{u_f, \lambda_i} \frac{1}{2} \|\ddot{q} - \ddot{q}_d\|^2 \quad (7a)$$

$$\text{subject to } \lambda_i \geq 0 \quad (7b)$$

$$u_f \in U \quad (7c)$$

$$\ddot{q} = f(q, \dot{q}, 0) + W(u_f + \sum_i J_i^T V_i \lambda_i) \quad (7d)$$

$$J_i \ddot{q} + \dot{J}_i \dot{q} = 0. \quad (7e)$$

The predicted acceleration of the character is \ddot{q} . The objective penalizes accelerations different than the desired acceleration, \ddot{q}_d , which was chosen to track the reference motion. This minimization problem can be solved efficiently: it features a quadratic objective with a positive-semidefinite Hessian, and the constraints are linear. This yields a convex quadratic programming (QP) problem. The QP is solved at a much slower rate than the simulation. At time steps where it is not solved, the previously calculated forces are used.

5. Proportional-Derivative Component

Solving the QP in the predictive component is fast but not immediate. The drawback of this latency is that the predictive component cannot adapt to disturbances in between updates to its control signal. We resolve this problem with a PD control that adjusts the QP solution at each simulation step. The PD control guides the character through contact transitions and provides immediate responses to disturbances.

McSim’s PD component computes u_b in Equation 1 at each step of the simulation. It is implemented using a critically-damped proportional-derivative (PD) controller [RH91]. The form of this control varies according to the particular joint. Since the root joint of the character is unactuated, no feedback forces are computed for the root dof’s. Pin joints are computed using a standard critically damped feedback law

$$u_b = k_s(q_r - q) - 2\sqrt{k_s}\dot{q}. \quad (8)$$

To compute the feedback forces of a ball joint, we first compute the composite rotational inertia of all of its child links in world coordinates:

$$I_{c,j} = \sum_{l \in c(j) \cup j} R_l I_l R_l^T. \quad (9)$$

The resulting feedback force is then computed as

$$u_b = k_s I_c d(q_r, q) + 2\sqrt{k_s I_c}(\dot{q}_r - \dot{q}). \quad (10)$$

Note that the term $\sqrt{k_s I_c}$ means taking the square root of each element of the matrix $k_s I_c$. Multiplying by the world-space inertia matrix insures that the feedback force is scaled by the appropriate amount relative to the actual current distribution of mass supported by the joint. The resulting force, u_b , is added to the current predictive force u_f to give the total force at each time step.

6. Maintaining Balance

McSim maintains balance by tracking the input motion with forces that are consistent with the current contact environment. Other works employ a similar approach by using formulations specific to static contact [ZH02] or infinite friction and planar contacts [KKI02, HMPH04, VB04]. In contrast, McSim uses a model of contact dynamics that can account for more general geometric and frictional properties of the contacting surfaces [Wie02, AdSP07].

In certain cases, heuristic methods can adapt the input motion directly to improve tracking robustness. For example, one could track a parameterized family of motions rather than a single motion [WC06] or adapt the center of mass motion through a feedback [AdSP07]. For some of the 2D motions presented in the results section, we employed a feedback scheme similar to the heuristic used in the SIMBICON system [YLvdP07]:

$$\theta_d = \theta_{d0} + c_d d + c_v v \quad (11)$$

where θ_d is the desired angle of the swing hip, θ_{d0} is the value of the swing hip in the reference motion, d is the horizontal distance between the root link and the support foot, and v is the horizontal velocity of the root link. Contrary to SIMBICON's approach, we do not change the gains, c_d and c_v , with changes in contact state. They are fixed for a particular motion.

McSim is largely insensitive to the particular choice of the gains. Normally, McSim tracks the input motion even when the gains are set to zero. However, adding this form of balance feedback improved the robustness of a character walking on a moving platform and allowed the controller to track a run cycle indefinitely. A drawback to using this particular form of balance feedback, however, is that it is specific to walking and running motions. Similar methods of adapting the input motion have been applied to other motions [Woo00].

7. Results

McSim produces life-like character motion similar to a provided input motion. In the following section we highlight results that demonstrate McSim's ability to adapt motion capture data to new physical environments and track a variety

of input motions. We also explore the sensitivity of the approach to various modeling errors and discuss the quality of the results. Finally, we provide implementation details.

7.1. New Environments

An exciting application of McSim is adapting motion data to new physical environments. For example, a motion recorded on flat ground can be adapted to walk up or down an inclined ground plane. In our experiments, successful walks were created for uphill slopes as large as five degrees and downhill slopes as large as 10 degrees. Simple kinematic playback of the motion would walk through the ground or into the air [dSAP08].

In a physical simulation, the environment can change dynamically and a character must react to maintain plausibility. Our controller allows motion data to adapt to its environment. We present several results where the character is perturbed by flying balls or obstructed by blocks. The ground too can evolve dynamically as evidenced by simulations of the character walking over a moving platform and a seesaw [dSAP08].

7.2. Tracking

McSim is capable of tracking a wide range of motions in 2D and 3D including walking, running, and jumping motions [dSAP08]. These motions exhibit variations and transitions between modes such as from standing to walking and walking to standing.

A key feature of McSim is that there are few parameters that require tuning. To generate the results, two parameters were tuned manually: the optimal feedback gain used in Equation 4 and a scale factor on the intrinsic joint stiffness parameters used in Equations 8 and 9. In most cases, it was not difficult to find a satisfactory setting of these parameters as a large range of values led to satisfactory results as explained in Table 1. Even across different types of motion, identical parameter values lead to good results.

Though McSim does not satisfy any optimality criterion, it achieves good tracking results in practice. In the absence of large disturbances to the physical system or large errors in the physical character model, McSim will succeed in tracking the input motion. The plots in Figure 3 depict the squared tracking error (squared Euclidean distance between the actual state vector and the desired state vector) over time for selected motions. The plots illustrate several interesting features of the tracking system. First, the beginning of the walk motion is a period of standing. The system has little trouble tracking this portion of the motion. More energetic motions lead to more error. The spikes in the error curves coincide with changes in contact state suggesting that the predictive model could be improved by accounting for mismatches in the current contact state and the contact state in the reference motion.

Motion	k	k_{os}
2D Punchy	0.02	1000
Downhill	0.02	300
Walk Wave	0.02	1000
Sneaky*	0.005	500
2D Jump	0.05	1000
Run	0.2	600
Backwards	0.05	1000
Soldier	0.01	600
March	0.05	1000
Limp	0.08	1000

Table 1: This table lists the relevant parameters used to generate selected results. k is a scale factor that multiplies the intrinsic joint stiffness parameters of the character listed in Tables 2 and 3 which are then used in the PD feedback component of the system. k_{os} is a gain used to calculate a modification to the acceleration from the input motion as in Equation 4. These two parameters were tuned manually to achieve a desired tracking result but reasonable results are achieved for a range of settings. For most 2D motions, values of k in the range between 0.005 and 0.5 worked. The setting of k_{os} is also flexible. Values in the range of 300 to 2000 typically work for this parameter. In many cases, the same settings achieved good results for many different motions. Starred motions were simulated using stiff springs at contacts. Despite using a different contact model, McSim tracks these motions well.

7.3. Modeling Errors

The tracking quality of McSim is adversely effected by physical mismatches between the character model and the capture subject. To explore the effect of modeling errors we introduce various modeling changes and measure the change in tracking performance.

One potential source of error in tracking motion capture data is an incorrect physical model of the subject. The mass distribution and inertial properties are often based on statistical models that are often quite different than the actual properties of the recorded subject. This mismatch can make an input motion physically infeasible for the character. To illustrate the sensitivity to errors in mass distribution, we plot the squared error for different versions of the 2D model for the walking motion in Figure 4. The mass of the character was redistributed to create three new versions of the original. One version of the character has a left leg that is twice as heavy as the right leg. In the next version, the upper body’s mass is doubled while the lower body mass is cut in half. Finally, we double the mass of both legs. For walking motions, McSim is more sensitive to errors in the mass properties of the legs.

Contact geometry was modeled using four small spheres placed at the corners of each foot. The controller is somewhat insensitive to the simulator’s contact dynamics. To il-




Figure 3: Shown are plots of squared error over time for four selected motions. The plots illustrate several interesting features of the tracking system. First, the beginning of the walk motion is a period of standing. The system has little trouble tracking this portion of the motion. More energetic motions lead to more error. The spikes in the error curves coincide with changes in contact state suggesting that the predictive model could be improved by accounting for mismatches in the current contact state and the contact state in the reference motion.




Figure 4: Shown are plots of squared error over time for four versions of the 2D model tracking the walking motion. The modifications are described in the legend. McSim is more sensitive to errors in the mass properties of the legs.

lustrate this, we compared the performance of the controller on a walking motion with varying coefficients of friction in 5. Tracking performance was not greatly effected. Contact dynamics were approximated using a friction cone model with a coefficient of friction ranging from 0.75 to 2.0 or stiff springs as in [SKL07].

The feet present another difficulty when tracking motion capture data. Our motion capture data for the ankle is fairly




Figure 5: In these plots, the squared error is shown for a walking motion where the coefficient of friction in the predictive component is varied from 0.5 to 1.5. The simulator's coefficient of friction was fixed at 1. For walking motions, the error is not greatly effected by the coefficient of friction used in the model. When the predictive model's coefficient of friction exceeds the actual coefficient of friction, performance is worse, but only slightly.




Figure 6: Shown are plots of squared error over time for versions of the 2D model with different sized feet. The big feet were 4 centimeters larger than the standard feet. The small feet are 4 centimeters smaller than the standard while the smallest feet are 8 centimeters smaller. These plots show that, at least for walking, slightly better results can be achieved by shrinking the foot. This might suggest that the actor performing the motion had slightly smaller feet. However, the results indicate that McSim is robust to small discrepancies in foot size.

inaccurate. We offset the ankle angle by a constant so that the character's contact points are flush with the ground while standing. To get a feel for how sensitive McSim is to variations in foot geometry, we varied the foot size of the 2D model and plotted the results in Figure 6. The big feet were 4 centimeters larger than the standard feet used in most of the results in this paper. The small feet were 4 centimeters shorter than the standard while the smallest feet were 8 centimeters shorter. The results indicate that McSim is robust to small discrepancies in foot size.

7.4. Motion Quality

The results of McSim's tracking often look robotic and abrupt. For example, the 3D marching motion makes hard contacts with the ground that are not present in the reference motion. The 2D walk uphill sways a bit unnaturally as well. There are a couple of factors that affect the quality of the results. The first is that the short-horizon approach to tracking is a greedy approach. It applies large torques to try and immediately cancel any errors. These large forces can lead to unnatural accelerations and motion. The other factor affecting quality is the fact that gain parameters are manually set by hand. This was more of an issue for the 3D examples which were more sensitive to the gain parameter settings.

7.5. Experimental Setup

The motion data for this work came from two sources. The 2D examples were downloaded from <http://mrl.snu.ac.kr/research/ProjectSimulBiped/SimulBiped.html>.

Link	k_s	Mass	Inertia
head	3000	3	0.011
upper arm	4000	2	0.022
lower arm	3000	1	0.009
torso	N/A	10	0.176
thigh	4000	7	0.121
shin	4000	5	0.077
foot	4000	4	0.019

Table 2: This table lists the inertial properties of each link in the 2D model and the stiffness of the associated joint. Note that there is no stiffness for the unactuated root joint. It is also important to note that the stiffnesses listed here are not directly used by the PD feedback component. They are first scaled by a single scale parameter that is typically much less than one. This scaled value is used to calculate a critical damping gain. The units are as follows: newtons per radian for the gains, kilograms for the mass, and kilogram meters cubed for the inertias.

This data was converted to 2D from motion capture data as described in [SKL07]. The 3D data was captured and processed using a standard motion capture system.

A prerequisite of simulating character motion is a physical model of the inertial and stiffness properties of the character's limbs and joints. A good model is important as significant errors make the input motion physically infeasible for the model. For the 2D examples, the physical model




Figure 7: The models. A free joint has six degrees of freedom and is represented by a position and a quaternion. A pin joint has one degree of freedom and is represented by an angle of rotation. The center of mass of each link is located at the center of the link. Inertial and joint stiffness properties are listed in Table 2 and 3.

Link	k_s	Mass
trunk	N/A	12.92
thigh	4000	9.0853
shin	4000	3.944
foot	1000	1
toes	4000	0.3
thorax	3000	17.155
clavicle	4000	2.535
upper arm	4000	1.435
lower arm	3000	0.575
hand	3000	0.5

Table 3: This table lists the inertial properties of each link in the 3D model and the stiffness of the associated joint. Again, there is no stiffness for the unactuated root joint.

Num. Vars.	QP Solve Time (secs)
36	0.0013
44	0.0015
52	0.0023
68	0.003
150	0.007
154	0.0075
158	0.0097

Table 4: Timing results for the QP solver as a function of the number of variables in the QP. The number of variables is a function of the number of degrees of freedom in the character and the current contact state. Note that, for ease of implementation, we used dummy variables for the acceleration of each degree of freedom. This is not strictly necessary and would result in a much smaller QP problem.

(see Figure 7) has the same properties as the one used in [SKL07]. The root link, however, is three dimensional. Its state is represented with a position vector, an orientation quaternion, and linear and angular velocity. The resulting model has 18 dofs. The inertial properties and joint stiffness parameters are presented in Table 2. These stiffness parameters are first scaled uniformly by a gain factor that is smaller than one and then used as the PD gains in 8 and 9. The 3D model has 57 dofs. The parameters for the 3D model are presented in Table 3.

The simulations were executed in DANCE [SFNTH05] using the Open Dynamics Engine (ODE) as the simulator. The step size was 1 ms for the 2D examples and 0.1ms for the 3D examples. We use a smaller step size for 3D examples as ODE was unstable with larger step sizes. A simulator using an implicit or semi-implicit integration scheme could presumably use a larger step size.

The controller implementation sets up the QP problem described in section 4 using the current contact state from the simulation. It uses our C++ implementation of recursive dynamics equations [FO00] to compute various dynamical quantities needed for the optimization such as the inertial matrix of the system and gravitational and centrifugal forces on the system. The QP is solved using SQOPT [GMS97]. Timings for the QP solver on a Pentium 4 2.8 Ghz processor are presented in Table 4. The code for the PD component took roughly 0.4 ms on the 3D character and 0.05 ms on the 2D character.

8. Conclusions

Motion data is an intuitive way to direct the actions of a physically simulated character. Determining the forces that track the motion faithfully while respecting physical

and environmental constraints is a difficult problem. McSim finds these forces at interactive rates making it suitable for the control of characters in interactive applications such as games.

McSim sacrifices optimality for computational performance. This sacrifice impacts the quality of the resulting motions. Quality was also impacted by the manually set parameters of the controller: the gain on desired acceleration and the PD gain. For some 3D motions, it was more difficult to find parameters that produced nice results. Also, there were certain motions that we could not track well such as turning motions. An interesting area of future work would be to apply optimization techniques that automatically tune the manually set parameters. In addition to reducing dimensionality, parameterizing control with our approach may help smooth the energy landscape, making it easier to find solutions.

Tracking a single input motion is not a good strategy for robust and stable control of a physically simulated character. In this paper, we experimented with a simple heuristic that adjusts the desired angle of the swing hip to help stabilize walking and running. In the future, we would like to incorporate long range planning to improve the quality of the output motion and improve the stability of the controller. This would require a good understanding of which aspects of the motion are crucial for stability versus those aspects that can vary.

References

- [AdSP07] ABE Y., DA SILVA M., POPOVIĆ J.: Multiobjective control with frictional contacts. In *Symposium on Computer Animation (SCA)* (Aug. 2007), pp. 249–258.
- [AFO03] ARIKAN O., FORSYTH D. A., O'BRIEN J. F.: Motion synthesis from annotations. *ACM Transactions on Graphics* 22, 3 (July 2003), 402–408.
- [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.
- [BJ05] BARBIČ J., JAMES D.: Real-time subspace integration for st. venant-kirchhoff deformable models. *ACM Transactions on Graphics* 24, 3 (Aug. 2005), 982–990.
- [Coh92] COHEN M. F.: Interactive spacetime control for animation. In *Computer Graphics (Proceedings of SIGGRAPH 92)* (July 1992), vol. 26, pp. 293–302.
- [dSAP08] DA SILVA M., ABE Y., POPOVIĆ J.: Simulation of human motion data using short-horizon model-predictive control. CSAIL Work Product (<http://hdl.handle.net/1721.1/40091>), Jan. 2008.
- [FO00] FEATHERSTONE R., ORIN D. E.: Robot dynamics: Equations and algorithms. In *International Conference on Robotics and Automation (ICRA)* (2000), pp. 826–834.
- [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.
- [FP03] FANG A. C., POLLARD N. S.: Efficient synthesis of physically valid human motion. *ACM Transactions on Graphics* 22, 3 (July 2003), 417–426.
- [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.
- [GMS97] GILL P. E., MURRAY W., SAUNDERS M. A.: *User's Guide for SQOPT 5.3: A Fortran Package for Large-Scale Linear and Quadratic Programming*. Tech. Rep. NA 97–4, University of California, San Diego, 1997.
- [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.
- [HPP05] HSU E., PULLI K., POPOVIĆ J.: Style translation for human motion. *ACM Transactions on Graphics* 24, 3 (Aug. 2005), 1082–1089.
- [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.
- [KGP02] KOVAR L., GLEICHER M., PIGHIN F.: Motion graphs. *ACM Transactions on Graphics* 21, 3 (July 2002), 473–482.
- [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.
- [LHP05] LIU C. K., HERTZMANN A., POPOVIĆ Z.: Learning physics-based motion style with nonlinear inverse optimization. *ACM Transactions on Graphics* 24, 3 (Aug. 2005), 1071–1081.
- [MK05] MUKAI T., KURIYAMA S.: Geostatistical motion interpolation. *ACM Transactions on Graphics* 24, 3 (Aug. 2005), 1062–1070.
- [MP07] MCCANN J., POLLARD N.: Responsive characters from motion fragments. *ACM Transactions on Graphics* 26, 3 (2007), 6.
- [PW99] POPOVIĆ Z., WITKIN A. P.: Physically based

- motion transformation. In *Computer Graphics (Proceedings of SIGGRAPH 99)* (Aug. 1999), Annual Conference Series, ACM SIGGRAPH, pp. 11–20.
- [Rai86] RAIBERT M. H.: *Legged Robots That Balance*. MIT Press, Cambridge, MA, 1986.
- [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.
- [SFNTH05] SHAPIRO A., FALOUTSOS P., NG-THOWHING V.: Dynamic animation and control environment. In *Proceedings of Graphics Interface (GI)* (2005), pp. 61–70.
- [SHP04] SAFONOVA A., HODGINS J., POLLARD N.: Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces. *ACM Transactions on Graphics* 23, 3 (Aug. 2004), 514–521.
- [SKL07] SOK K. W., KIM M., LEE J.: Simulating biped behaviors from human motion data. In *Computer Graphics (Proceedings of SIGGRAPH 07)* (jul 2007), vol. 26 of *Annual Conference Series*, ACM SIGGRAPH, pp. 107:1–107:9.
- [SP05] SULEJMANPASIĆ A., POPOVIĆ J.: Adaptation of performed ballistic motion. *ACM Transactions on Graphics* 24, 1 (Jan. 2005), 165–179.
- [SvdP05] SHARON D., VAN DE PANNE M.: Synthesis of controllers for stylized planar bipedal walking. In *International Conference on Robotics and Automation (ICRA)* (2005), pp. 2387–2392.
- [TLP06] TREUILLE A., LEWIS A., POPOVIĆ Z.: Model reduction for real-time fluids. In *Computer Graphics (Proceedings of SIGGRAPH 06)* (2006), Annual Conference Series, ACM SIGGRAPH, pp. 826–834.
- [VB04] VUKOBRATOVIC M., BOROVAC B.: Zero-moment point—thirty five years of its life. *International Journal of Human Robotics* 11, 1 (2004), 157–173.
- [vdPL95] VAN DE PANNE M., LAMOURET A.: Guided optimization for balanced locomotion. In *Eurographics Workshop on Computer Animation and Simulation* (1995), pp. 165–177.
- [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.
- [Wie02] WIEBER P. B.: On the stability of walking systems. In *International Workshop on Humanoid and Human Friendly Robotics* (2002), pp. 1–7.
- [WK88] WITKIN A., KASS M.: Spacetime constraints. In *Computer Graphics (Proceedings of SIGGRAPH 88)* (Aug. 1988), vol. 22, pp. 159–168.
- [Woo00] WOOTEN W.L., HODGINS J.: Simulating leaping, tumbling, landing and balancing humans. In *Robotics and Automation, 2000. Proceedings. ICRA '00. IEEE International Conference on* (2000), pp. 656–662.
- [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.
- [YLvdP07] YIN K., LOKEN K., VAN DE PANNE M.: SIMBICON: Simple biped locomotion control. vol. 26, pp. 105:1–105:10.
- [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.