World Space Servoing for Character Animation under Simulation

Pawel Wrotek
Brown University

Odest Chadwicke Jenkins*
Brown University

Morgan McGuire
Williams College

1 Overview

Real-time, interactive motion brings characters to life in virtual worlds like video games. However, combining playback of motion capture (mocap) data with physical simulation remains a hard problem. Most previous systems either switch between predefined animation and ragdoll (such as in Half-Life 2) or reaction physics [Zordan et al. 2005], or compute parent-space mocap torques for each bone to apply during simulation [Faloutsos et al. 2001; Zordan and Hodgins 2002]. The switching approach to control limits interaction with the environment. The parent space approach suffers from instability and accumulated pose error over long linkages, like the human spine. As a result, creating controllers for executing motion data is nontrivial and requires a heavy burden for tuning motor gains. Neither previous method is an “off-the-shelf” solution for animating characters in highly interactive environments where locomotion, object manipulation, and inter-character interaction is desired.

This sketch describes a novel insight for combining mocap with simulation in real-time. We demonstrate that promoting joint target angles to world space produces stable playback of motion capture in the presence of dynamic constraints and continuous physical simulation. We report that motor gains in world space are more stable and require less manual tuning than in parent space. Our approach also provides physically plausible motion blending due to simulation torque limits. We demonstrate a boxing video game, similar to [Zordan and Hodgins 2002], in which the fighters realistically stagger under the impact of attacks while performing captured animations.

2 Algorithm

The character is a jointed rigid body is a pure rag doll in the absence of torques. We use ODE (http://ode.org) for constrained rigid body simulation at 120 Hz and G3D (http://g3d-cpp.sf.net) for rendering at 60 Hz. In this simulation context, we apply torques as follows.

Mocap specifies each bone’s target orientation (expressible as 3×3 matrix $\mathbf{T}_p$) relative to its parent. We promote the target orientation to world space $\mathbf{T}_w$ by recursively applying its ancestor’s *target* orientations. This is where we differ from the parent-space approach, which could be viewed as performing the same transformation using the *current* parent orientations, and thus accumulating error. Given the current orientation $\mathbf{C}_w$, we compute the discrepancy $\Delta_w = \mathbf{T}_w \mathbf{C}_w^{-1}$, which we convert to axis-angle representation $(\vec{a}\theta)_w$. We then apply a standard P-D torque τ to the servo,

$$\tau_w = k_g (\vec{a}\theta)_w - k_d \omega_w, \tag{1}$$

where ω_w is the current world-space angular velocity, k_g is the servo acceleration gain, and k_d is the servo damping gain. As with previous methods, the k constants are hand-chosen. In world-space, simulation stability appears to be more robust to gain changes than in parent-space methods, which makes the constant-tuning process much easier.

Our world space method has an analogy in robotics. Parent-space control is like measuring the configuration of a joint with a rotational encoders, whereas world-space is like having accelerometers on a bone to sense rotation in the global (Earth) coordinates.

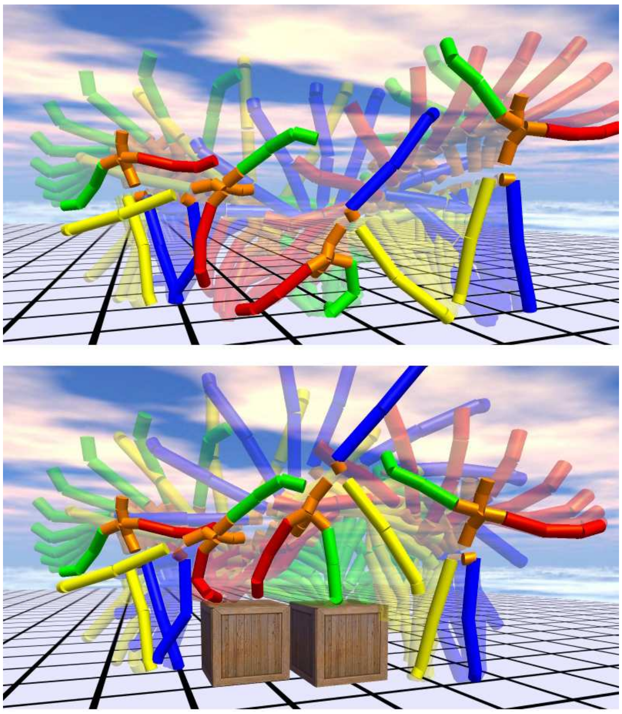

Figure 1: World-space torques naturally incorporate dynamic constraints in real-time. *Top:* a cartwheel animation under simulation; *Bottom:* the animation automatically adapted to an obstruction.

3 Results

The appeal of our approach is simplicity and stability. We do not explicitly encode translations or locomotion from mocap, motion blending, or adaptation to the environment. These important properties all emerge because our target angles are independent of error in the current pose and bones are constrained to realistic motion by the rigid body simulator.

In experiments, characters successfully adapted motions recorded on a flat-floor to novel environments containing obstructions (Figure 1). Motions transition smoothly as a character simultaneously drives all joints to match a new pose under simulation, and even poorly tuned world-space characters are much closer to critically damped (i.e., less ‘jiggly’) and more faithful to mocap than we could achieve with our best-effort parent-space gain tuning.

References

FALOUTSOS, P., VAN DE PANNE, M., AND TERZOPOULOS, D. 2001. The virtual stuntman: dynamic characters with a repertoire of autonomous motor skills. 933–953.

ZORDAN, V. B., AND HODGINS, J. K. 2002. Motion capture-driven simulations that hit and react. In *SCA 2002*, ACM Press, New York, NY, USA, 89–96.

ZORDAN, V. B., MAJKOWSKA, A., CHIU, B., AND FAST, M. 2005. Dynamic response for motion capture animation. *ACM Trans. Graph. 24*, 3, 697–701.

---

*e-mail:cjenkins@cs.brown.edu