Moving Crowds: A Linear Animation System for Crowd Simulation

M. Pražák[1] L. Kavan[1,2] R. McDonnell[1] S. Dobbyn[1] C. O'Sullivan[1]

[1*] Graphics, Vision and Visualisation Group, Trinity College Dublin
[2†] Disney Interactive Studios.

Figure 1: *A crowd simulation based on our animation system (left, middle). Our middle level-of-detail animation system significantly relieves the computational cost of animation generation for large crowds of characters while avoiding animation artifacts. Another advantage of our method is the ability to combine presimulated cloth animation with articulated figure motion (right).*

Overview

Animation of individual characters is a well-established topic in animation research with a broad range of methods available. However, the animation of many hundreds or even thousands of simultaneously visible characters presents new challenges. First, heterogeneous crowds require a variety of motions, making fully automatic motion pre-processing highly desirable. Second, the run-time component of the animation system needs to be fast enough to allow the real-time rendering of many individually animated characters. For this reason, previous crowd systems tended to employ very simplified motion models.

In this poster we present a linearized blending model based on hierarchy flattening and linear matrix interpolation which is used to build a simple crowd animation system with a medium level-of-detail (LOD). Because we are focusing on crowd animation, we assume a LOD management system is already in place, allowing the characters to be separated into three groups - those closest to the camera (which need to be animated with a highly accurate animation model), middle level-of-detail characters (targeted by our system) and the largest group of very distant characters (which are far enough away to allow extremely simple animation models – e.g., single repeated walk pattern).

Character motion is usually represented as a set of rotations of individual joints plus global translation and orientation. In contrast, we store and interpolate the skinning transformation directly (i.e., bone transformations relative to object space rather than relative to the parent bone) [Lewis et al. 2000], thus avoiding the bottleneck of transformation concatenation at run-time. This enables us to achieve better blending performance than classical animation systems. Even though this approach does not guarantee rigid bones, we show that the non-rigidity is below perceptual limits, and that it can actually provide the means to easily combine mesh-based animation (e.g., clothes) and skeletal animation on a single character.

Based on this blending method, we build a locomotion animation system suitable for crowds. As input, we accept motion-captured clips that reasonably sample the desired locomotion (10-15 clips per motion type). Note that while we focus on gait (as our target application is pedestrian crowds), nothing prevents us from processing more general motion data. Our system's run-time component takes the position and orientation of every agent in space for every frame as input (usually obtained by behavioural simulation of the individual agents). Its output is a set of transformation matrices which can be directly used for rendering skinned characters on graphics hardware.

The central concept of our approach is the *parametric space*, similar in spirit to [Park et al. 2002]. From the mo-cap clips we automatically derive motion parameters (such as the speed and curvature of a walk cycle). We then construct a simplicial complex (triangulation in 2D) that allows us to use linear methods to synthesize a motion with arbitrary parameters, thereby constraining the behaviour as little as possible.

Contributions

- A linearized motion blending model targeted at the requirements of crowd animation, together with a list of assumptions on the input motion, its mathematical description and both theoretical and practical tests of its performance
- A crowd locomotion system built on this model, exploiting the linearization of the blending model
- Completely automatic pre-processing of imperfect input data (i.e., not precisely periodic, not-constant speed)
- A versatile behaviour-animation interface for this model, allowing a range of behavioral simulation systems to be plugged in
- Roughly 3.5 times faster run-time execution compared to other animation methods

References

LEWIS, J. P., CORDNER, M., AND FONG, N. 2000. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. *International Conference on Computer Graphics and Interactive Techniques*.

PARK, S., SHIN, H., AND SHIN, S. 2002. On-line locomotion generation based on motion blending. *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*.

*e-mail:{PrazakM,Rachel.McDonnell,Carol.OSullivan}@cs.tcd.ie
†e-mail:Ladislav.Kavan@disney.com

I3D 2010, Washington, DC, February 19 – 21, 2010.
ISBN 978-1-60558-938-1/10/0002