Mix-and-Match Motion: Animation Virtual Experiences

Michael Gleicher and the UW Graphics Group *

Computer Sciences Department, University of Wisconsin, Madison

As the dangers of our world become more apparent, and our desire to prepare for them increases, there is a growing need for virtual experiences: simulated environments that convey to a participant the sense that they are in a situation, not just a place. A virtual experience can provide an opportunity to explore a situation that is too dangerous to rehearse, too impractical to create, or too inconvenient to attend. They may be used for training, for example to give first hand experience to emergency response personnel in directing panicked crowds; for design, for example to determine how people interact with a space before it is constructed; or for entertainment. The key to a virtual experience is that they combine environment and situation. Simply creating a visually compelling virtual environment is insufficient. Our vision is to use compelling virtual experiences in training, research and entertainment. To be effective in these applications, virtual environments must be authored to meet specific pedagogical and communicative goals. Realizing this vision, therefore, requires technologies for both simulating and rendering the worlds and enabling an author to specify behavior effectively.

Our goal is to populate virtual environments with synthetic characters such that a participant can experience a situation, not just a place. To meet this goal, we need to have characters that are directable such that an author can create a scenarios for the participant to experience, yet behave with enough realism to engage the participant. Satisfying these demands requires us to be able to describe the behavior of crowds of people in terms of their aggregate properties, yet have the movement of these characters have the quality of motion captured animation. Our research program is addressing these challenges by developing new motion synthesis methods that produce motion with quality comparable to motion capture data; crowd simulation methods that allow for aggregate specification of behavior; control mechanisms that couple these high- and low-level simulations; and rendering methods that display many characters with the required visual fidelity.

To achieve our goal of providing high-quality, directable, efficient crowds for virtual environments, we are pursuing several interrelated projects:

Motion Graphs — Our strategy for creating high-quality character motion is to synthesize it based on examples. We generate new motions by carefully assembling pieces of existing motions.

Our method analyzes a database of motion capture clips to identify opportunities for transitions between the clips. These connections are encoded into a *motion graph,* a directed graph with character poses as nodes and movements between these nodes as arcs. By only placing arcs that are known to be good motions into the graph, we can insure that any traversal of the graph produces a high-quality motion. The problem of producing a high-quality motion is reduced to that of finding traversals of the graph that meet desired goals.

We have already demonstrated the ability to synthesize high-quality locomotion that meets specified goals. Presently we are extending the method to allow for better control, a wider range of applicability, and better integration with real-time systems.

Our advances in example-based motion synthesis will lead to more realistic character motion that will enable more compelling virtual experiences.

Behavior Paint — Our strategy for describing aggregate behavior for crowds is to focus on authoring environments, rather than on specifying individual characters' actions. We will allow authors to place cues in the environment that guide characters' behaviors, effectively painting the virtual environment with paints that the characters can interpret.

The behavior paint approach will lead to a new generation of authoring tools that will enable the creation of effective virtual experiences.

Statistical Crowd Control — Our strategy for guiding large numbers of characters to meet specified directions is based on the observation that for crowds, the individuals do not matter, but rather the aggregate behavior does. This suggests an approach where direction is specified as desired distributions to be observed, and these distributions are transformed into probability distributions on the characters actions. Using control-theoretic principles to dynamically fine tune these distributions eliminates the need to determine them precisely in advance, and to have them respond to dynamic events.

Another unique aspect of our crowd control approach is its focus on the authors goals of aggregate behavior. Goals own characters, whom they direct, rather than the more traditional model of each character thinking independently.

Our advances in crowd control will empower authors to specify aggregate behaviors of large numbers of characters, enabling the creation of virtual experiences for urban training.

Model Sheets — To provide good visual appearance for characters, we are looking to the techniques traditionally employed by animators to define the connection between a character's appearance and movement. A *model sheet* is a set of examples that character designers use to specify this connection. We are developing an analog of model sheets for 3D characters, as well as methods for compiling these connections to run on graphics hardware.

The model sheets project will provide for better visual appearances for characters, enabling more compelling virtual worlds.

Our success at developing these technologies will allow us to populate virtual worlds with characters that behave realistically, and to specify the behavior of these characters as an aggregate crowd. This is the essential ingredient in creating virtual experiences for training applications for populated environments.

*gleicher@cs.wisc.edu, http://www.cs.wisc.edu/graphics
This research is funded in part by the National Science Foundation, grants CCR-9984506 and CCR-0204372, and gifts from Intel and Microsoft. Lucas Kovar is supported through an Intel Foundation Fellowship.