Exploiting the potential of image based crowd rendering

M. Izquierdo and A. Beacco and N. Pelechano and C. Andujar

Universitat Politècnica de Catalunya

Abstract

Per-joint impostors have been used to achieve high performance when rendering thousands of agents, while still allowing us to blend animation. This provides interactively animated crowds and reduces the memory footprint compared to classic impostors. In this poster we exploit the potential of per joint impostors to further increase both visual quality and performance. The CAVAST framework for crowd simulation and rendering has been used to quantitatively evaluate our improvements with the profiling tools that it provides. Since different applications will have different requirements in terms of performance vs. visual quality, we have extended CAVAST with a new user interface to ease this process.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

There have been several approaches in the past attempting to render thousands of realistic looking humanoids in real-time. Rendering impostors instead of geometry has proven to greatly increase performance, at the cost of decreasing visual appearance. By using LOD and switching between representations it is possible to have thousands of simulated agents in real-time without noticeable artifacts. In this poster we focus our attention on image-based techniques, and propose improvements over per-joint approaches to further enhance their visual appearance. We performed a perceptual evaluation to test adequate switching distances between our four representations: geometry, classic impostors [TC00], relief impostors [BSAP11], and flat impostors [BAPS12].

2. Per-joint relief impostors (RI)

Each character is encoded through a small collection of textured boxes storing color and depth values. At run-time, each box is animated according to the rigid transformation of its associated bone and a fragment shader is used to recover the original geometry using a dual-depth version of relief mapping [BSAP11]. Several improvements have been made to enhance the quality of the impostors while keeping a low memory footprint and having a negligible effect over performance. The first improvement consists on having **adaptive texture resolution** for different body parts. The overall representation is split into two VBOs, one for the head and another one for the rest of the body. Using a resolution of 64x64 for the body and 128x128 for the head, we obtain higher visual quality with just 1.3x more memory, as opposed to the 4x increment that would imply using higher resolution for the whole body.

To reduce the per-fragment cost, we have included an **adaptive number of linear and binary steps** into the relief mapping algorithm. As the agents move further away from the camera, the number of steps is reduced. Figure 1 shows the quality observed for different resolutions. We achieved a speed up of 3.7x on average with linear search steps ranging from 16 to 128, and binary from 4 to 10. The final improvement consists of replacing the binary search with the **secant method**, which iteratively bounds the intersection with the view ray and the geometry using a segment to segment test. This reduces the number of steps compared to binary search (order of 1.5 times faster).

3. Per-joint flat impostors (FI)

Instead of using six orthogonal relief maps for each joint, flat impostors are created by sampling each joint from multiple view directions [BAPS12]. A spherical Voronoi map is computed for the desired number of samples, and a cube map is built by projecting the Voronoi cells onto the cube faces. At run-time, a single texture lookup is enough to retrieve the fragment color, which is one order of magnitude faster than relief mapping. Quality enhancements have also

DOI: 10.2312/egp.20151031

Figure 1: *Improvements over relief impostors.*

focused on improving the upper part of the body. In this case we have also included **adaptive texture resolution** to have higher resolution for the head than for the rest of the body. This solution achieves higher visual quality with just 1.15x more memory, as opposed to the the 4x increment that would imply using higher resolution for the whole body. We have also performed **adaptive number of views** to further sample the most noticeable part of the body, lessening the popping artifact that appears when changing views. Figure 2 shows the effects of these improvements. Previous work created masks to define the border of textures based on 3 DOF and all possible rotation angles. For this improved version, only **real joint rotations** obtained from biomedical studies [Her07] are used and thus the mask associated with each joint matches more naturally the real movements found during animation. This leads to a reduction of artifacts due to gaps or overlapping around joints.

Figure 2: *FI with adaptive resolution and views.*

4. CAVAST framework

For this work we have used the CAVAST framework (https://sites.google.com/site/cavastproject/), which offers simulation, animation, and rendering of large numbers of characters in real-time [BP14]. CAVAST offers a modular architecture, allowing the programmer to add new controllers. For the purpose of evaluating our results, a new controller implementing classic impostors [TC00] has been integrated in the rendering module.

To further increase the variety of the crowd, we use color modulation. Using the alpha channel of the textures to identify body parts, the fragment shader modifies the HSV color based on the value stored in the alpha channel. To avoid unnatural colors for skin and hair, we fixed a predefined set of colors. Figure 3 shows a crowd with only two models and color modulation in the CAVAST framework.

In order to give the user the flexibility to choose values for the different improvements presented in this poster (e.g. resolution, number of search steps, etc) we have also included a new interface in CAVAST, for the user to interactively adjust those values to the needs of each application.

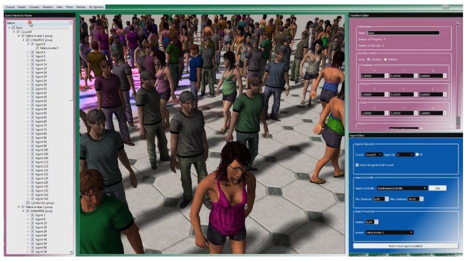

Figure 3: *The CAVAST framework for crowds.*

4.1. Results and Conclusions

For our perceptual evaluation study, we recorded 12 videos with a 500 agent crowd, for 3 configurations: classic, relief and flat impostors, 2 view points (aerial and street) and 2 crowd densities. Each test consisted of two videos side by side, where one side showed all agents with 3D geometry, and the other side started with geometry but slowly switched towards one of the impostor representations. The goal was to find out for what distance would the participants notice a drop in quality. We had 22 participants, watching 36 videos each, in random order. The results show that classic impostors were detected at 15.5m away, flat impostors at 10.7m and relief at 8.9m. We also counted how often people got the wrong answer which was 6% for classic, 9% for flat and 15% for relief.

References

[BAPS12] BEACCO A., ANDÚJAR C., PELECHANO N., SPANLANG B.: Efficient rendering of animated characters through optimized per-joint impostors. *Journal of Computer Animation and Virtual Worlds 23*, 2 (2012), 33–47. 1

[BP14] BEACCO A., PELECHANO N.: Cavast: The crowd animation, visualization, and simulation testbed. In *CEIG Spanish Conference on Computer Graphic* (2014). 2

[BSAP11] BEACCO A., SPANLANG B., ANDUJAR C., PELECHANO N.: A flexible approach for output-sensitive rendering of animated characters. *Computer Graphics Forum 30* (2011). 1

[Her07] HERMAN I. P.: *Physics of the Human Body*. Springer, Berlin, 2007. 2

[TC00] TECCHIA F., CHRYSANTHOU Y.: Real-time rendering of densely populated urban environments. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, UK, 2000), Springer-Verlag, pp. 83–88. 1, 2