BurnAR: Feel the Heat

[1]Peter Weir* [1]Christian Sandor† [2]Matt Swoboda‡ [3]Thanh Nguyen§ [1]Ulrich Eck¶
[3]Gerhard Reitmayr‖ [1]Arindam Dey**

[1]Magic Vision Lab, University of South Australia, Adelaide, Australia
[2]Fairlight
[3]Institute of Computer Graphics and Vision, Graz University of Technology, Graz, Austria

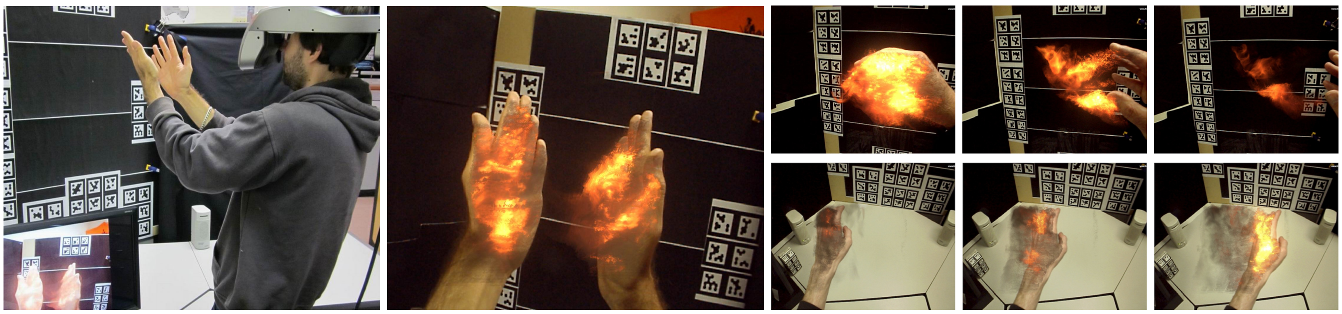

(a) (b)

Figure 1: (a) Our BurnAR demonstration enables users wearing a stereo head-worn display to experience the illusion of seeing their own hands burning, which we achieve by overlaying virtual flames and smoke on their hands. (b) Hand movement affects flames and smoke.

Abstract

Augmented Reality systems that run interactively and in real time, using high quality graphical displays and sensational cues, can create the illusion of virtual objects appearing to be real. This paper presents the design and implementation of BurnAR, a demonstration which enables users to experience the illusion of seeing their own hands burning, which we achieve by overlaying virtual flames and smoke on their hands. Surprisingly, some users reported an involuntary warming sensation of their hands.

Index Terms: H.5.1. [Information Interfaces and Presentation]: Multimedia Information Systems—[Artificial, augmented and virtual realities] H.1.2. [Information Systems]: Models and Principles—[Human factors]

1 Introduction

Advances in technology have allowed the development of sophisticated systems that display real and virtual objects aligned with each other in a real world scene, termed Augmented Reality (AR). AR systems that run interactively and in real time, using high quality graphical displays and sensational cues, can create the illusion of virtual objects appearing to be real.

In October 2011, we presented the BurnAR demonstration at the 2011 International Symposium on Mixed and Augmented Reality. Over two days, more than 100 visitors experienced the illusion of seeing their own hands burning, which we achieved by overlaying virtual flames and smoke on their hands. Surprisingly, around one fifth of users reported experiencing an involuntary sensation of heat in their hands, with one user also experiencing the smell of smoke. At subsequent public demonstrations of BurnAR, a similar proportion of users also reported these same sensations.

2 Design

Our demonstration simulates fire and smoke effects, being responsive to the user's motions, and always visually interesting. The choice of a fire effect over other possible effects is deliberate. Fire is self-illuminating, removing any requirement to match lighting of the generated visuals with the real-world environment—a significant challenge—and it is straightforward to composite. Fire is abstract and noisy in appearance and changes quickly, which allows some margin for error when matching with our computer vision component. Fire is recognizable and well-understood, but seeing one's own hands on fire is something most people do not experience—so providing an experience which is easy to grasp but compelling and unique.

With the system that we have developed for the BurnAR demonstration, the user has a stereoscopic view of the real world. Being able to see the real-time kinaesthetic motion of their hands and lower arms provides a highly believable system that supports a number of sensi-motor contingencies. The user is presented with a video display of the real world and their own hands interacting with the virtual flames and smoke, complemented with sound effects of burning.

3 Implementation

The BurnAR demonstration comprises several components based on a closed layer architecture, where data flow is only allowed between adjacent layers. The bottom MR Platform [5] layer drives the stereo head-worn display, streaming a pair of video images, and camera poses to the computer vision layer. In the computer vision

*e-mail: weirpc@gmail.com
†e-mail: chris.sandor@gmail.com
‡e-mail:mattswoboda@yahoo.co.uk
§e-mail:thanhnguyen.cs@gmail.com
¶e-mail:ueck@net-labs.de
‖e-mail:reitmayr@icg.tugraz.at
**e-mail:arindam.dey@unisa.edu.au

IEEE International Symposium on Mixed and Augmented Reality 2012
Science and Technology Proceedings
5 - 8 November 2012, Atlanta, Georgia

layer, the video images and hands, segmented by color, are processed to reconstruct a dense estimation of the hand surface. The resulting 3D point cloud, 3D scene flow, and a pair of video images are streamed to the computer graphics layer. The fire effect uses the 3D point cloud to initialize the particle system, which is overlaid onto the camera images. The distance d to the hand is used to control the volume of the fire sound, using the inverse square law for sound intensity: $1/(d^2)$.

The experimental platform was implemented on a 3.0 GHz Intel Core i7 quad-core desktop computer with 4 GB RAM, an nVIDIA GeForce GTX 570 graphics card, stereo sound card with external speakers, and a Canon VH-2007 video see-through head-worn display (HWD).

3.1 Online 3D Hand Reconstruction

To drive the interactive fire simulation, the system requires a dense estimate of the hand's surface and current motion. Ideally, we compute a 3D location and 3D velocity per hand pixel (3D scene flow) as seen through the HWD. Furthermore, the resulting 3D reconstruction must be available at frame-rates to support the interactive experience. The input to the reconstruction component are left and right images from the HWD cameras converted to grayscale, and raw left and right hand segmentation masks generated by MR Platform through color segmentation of skin color. Both the grayscale images and segmentation masks are rectified using the stereo calibration of the HWD cameras. This step results in image pairs, where scan lines lie on common epipolar lines, allowing more efficient processing in the following steps.

For depth map generation, we employ shape-from-silhouettes techniques, emerged since the early 1990s. Our method focuses on small regions of interest around the border. Using Tukey estimation for per pixel weighting takes only 10ms per frame, compared to the OpenCV grab-cut implementation at 5000ms per frame. After refining the hand segmentation regions, a disparity map is computed per scan line, similar to [1]. First, the disparity for the left border pixel and right border pixel is computed between left and right frame. Then, disparity values for each interior pixel are linearly interpolated between the border pixels and across rows. Finally, the depth map and a 3D point cloud for the segmented hand pixels are computed from the resulting disparity map.

The 3D scene flow of the hands is computed from intermediate 2D optical flow from the previous left frame to the current left frame. The algorithm passes on the following information to the computer graphics layer: camera images, refined left and right hand masks, 3D point cloud and 3D flow of the current frame indexed by pixel.

3.2 Computer Graphics

The graphics layer is implemented using the Demolition Engine from Fairlight, running on DirectX 9 in Windows and using HLSL shaders. It is responsible for generating a fire effect around the user's hands from the given computer vision input. The vision system provides it with a pair of depth/color/velocity images for the left and right eyes. The graphics system uses these to generate 3D voxel data from which to spawn fire and smoke. In order to do this, particles are first spawned by sampling visible points in the color, velocity and depth images and back-projecting into world space by the current camera transforms. Particle positions are randomized in view space depth based on assumed thickness of the objects being processed - typically hands, where reasonable guesses can be made. Points which, after randomization and reprojection, fall outside the visible mask are culled. A point has a position, color and velocity.

Given a set of 3D points determined to be inside the 3D objects, a voxel set is generated, which is used to feed a fire and smoke simulation. The basis of the simulation is a typical 3D Navier Stokes fluid solver [4] operating on a voxel grid and evaluated on the GPU

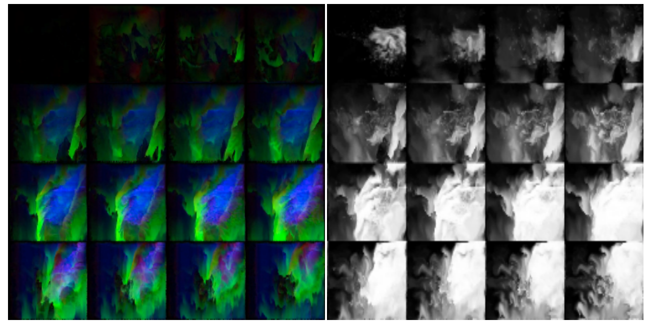

(a) (b)

Figure 2: Visualization of 16 slices of our volumetric fire simulation at one timestep: (a) velocities, (b) densities. Velocities are mapped to RGB from low (blue) to high (red). Densities are mapped to greyscale. Note that the user's hand is visible as a low-velocity region in the two middle rows of (a).

as in [3]. Each grid cell contains a 3D velocity vector, a density and a temperature. Each computation stage is executed serially in standard GPGPU fashion. A 128^3 voxel grid is arranged as a 2D array of 16x8 slices (see Figure 2 for a depiction of 16 slices) in a 2048x1024 render target and a computation stage executes by rendering a 2D quad covering the entire render target with a shader to perform the computation in parallel across the GPU cores. The color of each cell is determined by evaluating a function of the temperature and sampling a color lookup table which ranges from dark smoke to fire. A function of the density is used to set the alpha blend value for the cell [2].

Compositing with the video data is achieved by reading the original mask and depth images. Alpha values of the cells are adjusted by a function of the cell depth against the depth image value. This allows some depth to the composition so a hand may occlude the fire.

4 Conclusions & Future Work

In this paper, we have described the design, and implementation of BurnAR, a demonstration which enables users to experience the illusion of seeing their own hands burning, which we achieve by overlaying virtual flames and smoke on their hands. It is our intention to formally evaluate the stimulation of involuntary heating and smelling sensations generated through experiencing the demonstration.

Acknowledgements Canon Inc., for HWD and MR Platform.

References

[1] B. Dellen and F. Wörgötter. Disparity from stereo-segment silhouettes of weakly-textured images. In *Proceedings of the British Machine Vision Conference*, pages 96.1–96.11, London, UK, Sept. 2009. BMVA Press.

[2] R. Fedkiw, J. Stam, and H. W. Jensen. Visual simulation of smoke. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 15–22, Los Angeles, CA, USA, Aug. 2001. ACM New York, NY, USA.

[3] M. J. Harris. Fast fluid dynamics simulation on the gpu. In *GPU Gems*, pages 637–665. Addison Wesley, 2004.

[4] J. Stam. Stable Fluids. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 121–128, Los Angeles, CA, USA, Aug. 1999. ACM Press New York, NY, USA.

[5] S. Uchiyama, K. Takemoto, K. Satoh, and H. Yamamoto. MR Platform: A Basic Body on Which Mixed Reality Applications Are Built. In *ISMAR '02: Proceedings of the 1st International Symposium on Mixed and Augmented Reality*, pages 246–320, Darmstadt, Germany, Sept. 2002. IEEE Computer Society Washington, DC, USA.