Squidball: An Experiment in Large-Scale Motion Capture and Game Design

Authors (in alphabetical order):
Christoph Bregler, Clothilde Castiglia, Jessica DeVincezo, Roger Luke DuBois, Kevin Feeley, Tom Igoe, Jonathan Meyer,
Michael Naimark, Alexandru Postelnicu, Michael Rabinovich, Sally Rosenthal, Katie Salen, Jeremi Sudol, Bo Wright
NYU Technical Report TR2005-869, more info at http://squidball.net

Figure 1: Electronic Theater audience playing Squidball

Abstract

This paper describes a new large-scale motion capture based game that is called Squidball. It was tested on up to 4000 player audiences last summer at SIGGRAPH 2004. It required to build the world's largest motion capture space, the largest motion capture markers (balls), and many other challenges in technology, production, game play, and social studies. Our aim was to entertain the SIGGRAPH Electronic Theater audience with a cooperative and energetic game that is played by everybody together, in controlling real-time graphics and audio, while bouncing and batting multiple large helium filled balloons across the entire theater space. We detail in this paper all the lessons learned in producing such a system and game, and argue why we believe Squidball was a great success.

1 Introduction

Squidball is a large-scale, real-time interactive game that uses motion capture technology and computer graphics to create a unique and energetic experience for mass audiences. Using the world's largest calibrated motion capture volume with participating player audiences of up to 4,000 people, the game debuted on August 12th, 2004, at the Los Angeles Convention Center as pre-show entertainment for the SIGGRAPH Electronic Theater.

This paper describes the design criteria and technology behind this venture. It also explores the adventures and challenges that a production team of over 30 people had to overcome, the lessons learned and the world record the team had to break.

SIGGRAPH audiences experienced a similarly interactive Electronic Theater pre-show over a decade ago when the Cinematrix System was introduced in 1991 [Carpenter 1993]. Cinematrix was an interactive entertainment system that allowed members in the audience to control an onscreen game using red and green reflective paddles. Other interactive entertainment systems have been tested on audiences in the hundreds to thousands, which is described in greater detail in section 2.

The success of Cinematrix was the original inspiration for our work, and we initiated our project to bring back this style of pre-show entertainment to SIGGRAPH. Although the 2004 Electronic Theater was our first public test, we envision *Squidball* being deployed in other large audience events for entertainment, social studies, team building exercises and other potential applications.

Developing and testing such a system was a very unique and high-risk venture with many challenges. Other games, graphics and interactive systems are usually designed for single user or a small group, and go through several test cycles. However, for the Squidball project, we were dealing with many factors orders of magnitude larger than standard environments, including a gathering of 4000 people, the construction of a system using a $240 \times 240 \times 40$ feet motion capture volume and a huge screen. Furthermore, the system had to work the first time, without the benefit of any full-scale testing.

During our initial brain-storming sessions, we decided to create a game that is played by bouncing and batting a number of balls that would be used as wireless joystick/mouse inputs to a game, across the entire audience. We also decided to track the balls using 3D motion capture technology and to use this data to drive real-time graphics and an audio engine.

We set out to design a game that:

- Requires no explanation of the rules – people must be able to pick it up and start playing;
- Is even more fun than just hitting a beach ball around an auditorium (we already know this is fun);

- Is fundamentally about motion-capture and takes full advantage of the capabilities of this technology;
- Can be played by 4,000 people simultaneously using a small number of input devices;
- Can be played by people standing, sitting or holding a beer in one hand (there was a cash bar in the Electronic Theater); and
- Involves people hitting balls AND looking at a projection screen.

In the following sections, we first summarize what related interactive experiments have been done in larger audiences, in section 3 we describe in detail the challenges and our solutions on building the large scale 3D motion capture volume, and in section 4 we do the same for the game design and game testing challenges.

2 Related Work

As we mentioned, the inspiration for Squidball came from the Cinematrix system [Carpenter 1993] at the SIGGRAPH 1991 Electronic Theater. Cinematrix was shown at several other events, including SIGGRAPH 1994, 1998, Ars Electronica, ACM '97 Expo Conference, trade shows, corporate events and several permanent installations. With this game, every audience member had red and green reflective paddles. A vision-based system could determine, for each person, if they were holding the red or the green side of the paddle towards the camera. Many different games were tested, including a voting system, Pong and a Flight Simulator. In the voting schema, the system counted how many red vs. green paddles were shown. In Pong, the left side of the audience played against the right side, and the position of the paddle controlled the ratio of red and green paddles on each side of the audience. The Flight Simulator was one of the more complex games where the audience controlled the roll and pitch of a flight simulator using paddles. It was surprising how quickly the audience learned to control the games and to jointly coordinate the mix between red and green paddles. Of course, the yelling and excitement of a large audience was also part of the show.

Another set of similar interactive techniques were studied at student theater screenings at CMU [Mayenes-Aminzade et al. 2002]. Three different input techniques were tested on large audiences: 1) Vision based determination of left vs. right leaning of audience members, 2) shadow tracking of a beach ball, and 3) laser pointer tracking. The input technique most closely related to Squidball is the 2D beach ball shadow tracking, where the location of the shadow could be used as a cursor in several 2D games.

D'CuCKOO (a music band that uses various kinds of new technological instruments) designed a gigantic beach ball that creates music as the audience bats it around. The MIDI-Ball [Blaine 2000], a wireless 5-foot sphere, converts radio signals into MIDI commands that trigger audio samples and real-time 3-D graphics with every blow. The MIDIBall debuted at the Grateful Deads Mardi Gras show at the Oakland Coliseum in 1992.

There have been other systems reported, that track small groups of people as they perform interactive music and dance activities [Ulyate and Bianciardi 2004], or are used for home video games [Freeman et al. 1996] but none of them have been tested on thousands of players.

3 Large-Scale Motion Capture

Here we describe the challenges and experiments of building a large-scale motion capture space and how this ties into the Squidball game engine and game testing. In section 4 we describe additional details about the game design itself.

Our target venue was Hall K in the Los Angeles Convention Center, which was converted into a 4000-seat presentation environment to screen the Electronic Theater for the 2004 SIGGRAPH conference. The total space was 240×240 feet. We needed to build a motion capture volume that covered the entire seating area and allowed enough height above it to throw the balls up in the air: a capture volume of $190 \times 180 \times 40$ feet. To the best of our knowledge, no motion capture space of this size had been built before. One of the larger reported spaces was built for the Nike commercial by Motion Analysis Corp and Digital Domain [MotionAnalysisStudios 2004]. It had dimensions of $50 \times 50 \times 10$ feet. It used 50 cameras to track six football players.

One of our design constraints was tracking multiple (up to 20) balls in 3D in real-time. We had a state-of-the-art Vicon motion capture system [Vicon 2004] with 22 MCAM2 cameras, and each camera had a field of view of 60 degrees (12.5mm lens) and 1280×1024 pixel resolution. In its intended use, the system can track standard motion capture markers (0.5 inch) in a capture distance of up to 25 feet. The markers are made of retro-reflective material. Visible light illuminators placed around the camera lens shine light out, and almost all light energy is reflected back into the camera. This makes the retro-reflective markers appear significantly brighter than any other object in the camera view, and image processing (thresholding and circle fitting) is used to track those markers in each view. Triangulation of multiple camera views results in very accurate and robust 3D marker tracking. Despite all of the advances in vision based tracking, retro-reflective markers and multi-camera triangulation is, to this day, the most accurate and robust 3D tracking technique.

We determined that the only way to utilize the Vicon motion capture system in a significantly larger space with the same number of cameras was to scale up each aspect of the system. The cameras' view scale up in an approximately a linear fashion; in other words, a marker 100 times larger in diameter and 100 times further away looks the same to the camera. Of course, because light intensity falls off with the square of the distance traveled, much greater illumination is necessary. With experimentation, we found that halogen stage lighting provided sufficient illumination for the Vicon tracker.

Three other challenges in scaling up the system were 1) producing the larger markers, 2) dealing with camera placement constraints, and 3) calibrating the space. All of them appear simple in theory but, in practice, these became critical production issues.

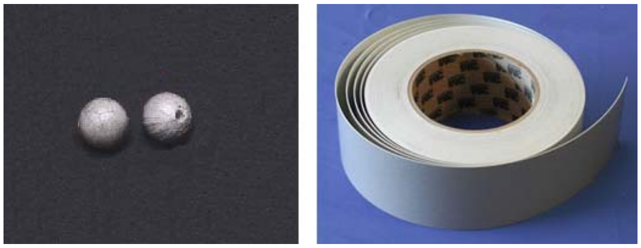

Figure 2: 0.5 inch markers and the 3M retro-reflective tape the markers are coated by.

3.1 How to produce the balls (markers)

Because the Vicon tracker performs circle fitting across multiple cameras, the motion capture system requires spherical markers for effective tracking. Non-spherical shapes are simply discarded or tracked unreliably by the system. Standard motion capture markers are plastic spheres, 0.5 inch in diameter, which are covered with 3M retroreflective tape as seen in figure 2.

In order to stage our event in a radically larger-than-normal space, with larger-than-normal camera view distances (max. 250 ft.) and using the balls as real-time inputs, we had to increase the size of those markers significantly. Initially, we experimented with many different marker objects. We determined that a 16-inch diameter marker was the smallest marker that could be robustly detected at 250 feet. In the final game, for dramatic effect and game-play, we opted for larger markers: 8-foot chloroprene bladders (weather balloons). In order to achieve the right bounciness, we under-inflated the balloons.

Each marker requires a retroreflective coating in order to be tracked by the Vicon motion capture system. Unfortunately, experiments using retroreflective spray-paint failed. The reflective intensity of the sprayed paint was 70% to 90% less than 3M retroreflective tape. Balls covered with this paint were not visible to the Vicon cameras at a distance greater than 50 feet, 200 feet short of our requirements.

After dozens of tests with paint, tapes, and fabrics of varying color, reflectivity and weight, we settled on specific 3M retroreflective fabric (model # 8910). Figure 3 shows the results of these tests. The first test (the ”lemon”), the second iteration (the ”tomato”), and the final version (the ”orange”), which ultimately produced a perfectly round spherical shape.

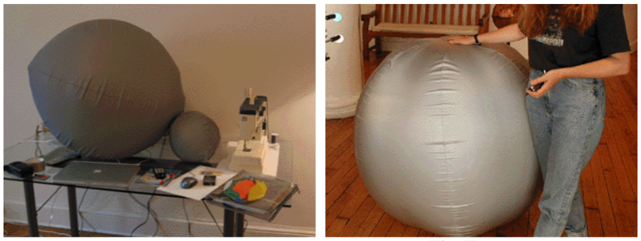

The Lemmon

The Tomato

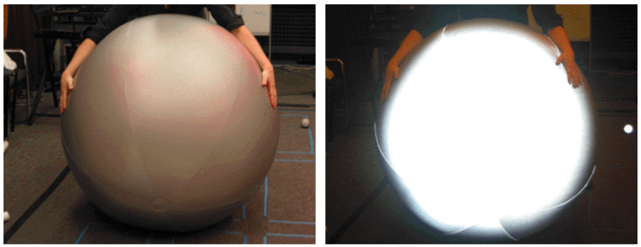

The Orange

The Orange Reflecting

Figure 3: The evolution of our retro-reflective balls.

In order to achieve a perfect spherical shape and to spread the force evenly throughout the surface, the fabric was cut on the bias, in panels like those of a beach ball. At this large scale, any of these shapes were adequate for the Vicon system to track them. The advantages of a perfect sphere were both aesthetic and functional. With the force evenly distributed, one spot is no more likely to rip the fabric than any other spot. Similarly, hitting the ball anywhere has the same predictable result.

The balls were inflated with helium to reduce their weight. Because the fabric was heavy, they did not float away when filled with helium, fortunately.

3.2 Camera Placement

Standard camera placement for a motion capture system is an “iterative refinement process” dependent on several site-specific aspects. For standard motion capture, cameras are usually placed on a rectangle around the ceiling, all facing into the capture space. Sample motion capture markers are distributed over the capture space, and cameras are adjusted so that each marker is seen by as many as possible cameras from as many as possible directions. Additionally, the tracking software is checked for each camera during placement.

In our scaled-up system, camera placement was a significant challenge. We could not afford as many trial-and-error cycles in camera placement that would be possible in standard-sized motion capture labs since our time was limited in the final space and each adjustment took a significant amount of time. Other logistical constraints affecting camera adjustment included: A) cooperating with the Union LACC workers schedule to get access to the ceiling and catwalks, B) coordinating between people on the 40-foot high ceiling and people on the floor up through radio-communication for each re-mount and re-alignment of a camera, C) getting live feedback from the Vicon PC station in the control booth to people on the ceiling so they could see the effects of their adjustments, D) camera view limitations - the 60 degree wide-angle lenses did not actually see a full 60 degree angle of view; even with the extra heavy studio lights mounted next to the cameras the visibility of the weather balloons dropped off after 250 feet in the center (and at even shorter distances at the perimeter of the camera view), and E) scale issues - moving balls on the ground takes much longer because of their large size and the distance to be covered. In a standard mo-cap studio, you pick up a marker and lay it down a few seconds later; in this space, we had to move a shopping cart with a ball or drive an electric car across the hall.

3.2.1 3D Simulation in Maya

In anticipation of all those problems, we designed a 3D model in Maya for all the target spaces, including one for the campus theater (our first test), one for the campus sports center (our second, third and fourth test), and one for Hall K at the LACC (Figure 4), our final show. This final model was derived from blueprints we obtained from the building maintenance team.

We also built a 3D model of the “visibility area” to determine the sight lines of the cameras. We experimented in the campus sports center and determined that the cameras could not “see” of center at distances of 250 feet. The further out we moved the balls, as shorter the visibility became. We therefore placed four cameras at one end of the court facing the other end. Moving a ball around on the capture space, we marked the 3D locations where the visibility of the ball vanished relative to each camera. Given this data, we built a 3D Maya model for the camera visibility volume (Figure 4 green). This volume was then used in our Maya building model to simulate several camera placement alternatives. Our goal was that each point in the capture volume should be seen at least by 3 cameras, given the constraints on the lengths of video cables. The final configuration we used was pretty close to our simulation. We settled on mounting evenly all 22 cameras around the left and right catwalk, and along the back-end catwalk and the center catwalk, but not the frontal catwalk. We didn’t want to mount any cameras and high-intenstiy lights above the screen, so the audience would not be distracted.

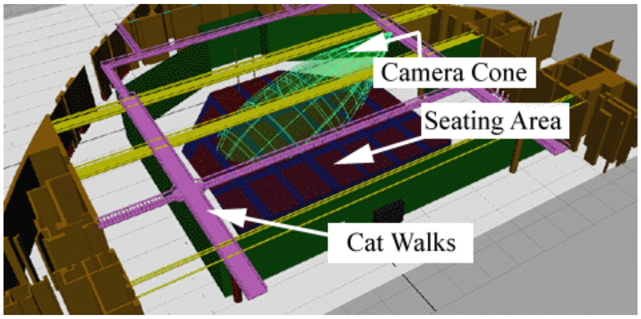

Figure 4: This shows our Maya model of Hall K and the visibility cone (green) for one of the cameras.

3.2.2 Mounting and Networking

We knew we had to set up the system in LACC very quickly, so we ran multiple practice sessions for camera mounting in New York, first for 10 cameras and then 22 cameras.

The setup required careful cabling. Vicon sends camera data first through analog wires to a Datastation, which thresholds video frames and compresses the resulting binary images. The Datastation then sends all 22 Video Streams at 120Hz over gigabit Ethernet to the Vicon PC, which does the real-time 3D tracking. In order to have the shortest possible video cable length, the Datastation had to be close to the cameras. In Hall K at LACC, the Datastation was placed 40 feet above the audience on one of the catwalks. The compressed video data was then sent via gigabit Ethernet down to the control booth in the back of the audience on the floor. The control booth contained all the workstations, including the real-time 3D tracker and the game system.

During camera placement we operated the Vicon PC in the control booth through a wireless laptop and Remote Desktop. This allowed us to walk from camera to camera on the catwalk, and do all adjustments, while remotely monitoring what the camera “sees” and how the tracking software performs. The wireless bandwith was high enough to do this with good latency. Of course, we had to use radio communications and lots of yelling from the ceiling down to people on the floor (who were moving target marker balls around) in order to adjust the cameras properly. In our NYC locations, this usually worked very well, and we had the cameras up in 1 to 2 hours. Unfortunately, at LACC we encountered quite a few surprises. As we said, complying with the Union workers schedule was a challenge, because they always had to stand next to us while we were mounting up on the catwalk or controlling the motion capture space so that it remained free of “occlusion” . Also, many other parties needing access to Hall K during our setup. Finally, we ran into cross-talk on the wireless network for our laptops from the LACC building antennas and from the exhibition space next-door. Nevertheless, we were able to mount the entire 22 cameras cables betwen 8am and 4pm (including all Union breaks and delays!).

3.3 Calibrating the space

The final challenge for the motion capture setup was camera calibration. Using the Vicon software, the calibration process in a standard space is done by waving a calibration object throughout the entire capture volume. Usually, this is a T-shaped wand that has 2 or 3 retroreflective markers placed on a straight line. (Figure 5) The 2D-tracking data for the calibration object from each camera is then used to compute the exact 3D locations, directions and lens properties of the cameras. This is called the calibration data, which is crucial for accurate 3D tracking.

Of course, the standard T-wand would not be seen by any camera in such a large target space (below pixel resolution). We determined that a 16-inch marker was the smallest marker that could reliably be seen and tracked from 250 feet. To overcome this, we built several “calibration T-wand” versions. Figure 5 shows one version that allowed us to “wave” the calibration object as high as 30 feet. We conducted initial tests on how much time an “exhaustive volume coverage” would take and how physically exhausting it would be using the roof of our lab. In the campus theater space and the campus sports center, we either walked the wand around holding it at several heights or skate-boarded through the space. In the final test at Hall K, we first used a crane and ropes. Ultimately, we ended up using a T-wand constructed out of bamboo sticks lashed together using a traditional Japanese method (which could be maneuvered by a single person due to it’s light weight, durability, and quick alterability) and then drove that around exhaustively at several heights on an electric car. A calibration run took around 30 minutes. We needed several calibration trials until we found a good path to achieve numerically good calibration software results. Tensions ran high during calibration in Hall K, but the process was a success (see video) and we had a spot-on reading of Hall K right up to the periphery of the seating areas. In actuality, we were able to track the balls beyond the boundaries of the game “board”.

Figure 5: Left: The standard sized calibration objects of length 15 inch. Right: Large 15 feet high wand.

Figure 6: Calibration in Hall K using a crane.

4 Software Integration of Motion Capture and Game Engine

Before we could start the various game tests, we needed a rapid prototyping environment connected with the real-time motion capture input, one that generated real-time computer graphics and sound effects.

The real-time visualization system and game engine were written using the Max/MSP/Jitter development environment distributed by Cycling74. The system consisted of five main components:

- A TCP socket communication system, which distributed motion capture data from the Vicon system. The system worked on a frame synchronized protocol, where the main Max computer would request a frame of motion capture data whenever it drew a frame (nominally 60fps).
- A Java-based tracking module that would take the raw motion capture data, filter it using Kalman filters and then extract useful metrics such as object velocity and collision detection. It would also perform continuous sequential numbering of the objects so that the tracking of any given ball would be continuous between frames.
- A game engine written partly in Java and partly as Max patches, which drove the game simulation. Submodules of this system included components that handled the basic game narrative (e.g. which level we were in, how much time remained for the level), the Squidball media files (pictures, textures), the interface to the cv module and tracker, and the graphics engine. The graphics engine used Jitter to render the game using OpenGL commands, with extensive use of texturing and alpha blending to create a nuanced look-and-feel for the game environment.
- An audio subsystem resident as a Max patch on a second computer (and receiving forwarded motion capture information as well as scene control from the main Max computer). This module contained a synthesizer/sampler written in MSP that allowed us to tie quadraphonic sound events and signal-processing parameters to collisions (both between balls and between balls and targets), floor bounces, ball Cartesian position, acceleration, and velocity. The audio subsystem was slaved to the main system and triggered sounds in response to game play on a micro level (individual data frames) as well as a macro level (providing different background music for each level).

The prototyping environment proved robust and fast enough to use in the final show, and we continued to tweak the game until the night before the opening.

In the production mode, we ran two duplicate sets of the game and audio computers for redundancy, with a switch, but fortunately this was never needed. The system required five human operators during the shows one on sound, one on the game system, one monitoring the Vicon PC, one watching the audience, and a master show controller who coordinated the team.

5 Large-Scale Game Design

We wanted to design a game that worked well in the poorly-understood dynamics of cooperation and competition of a large-scale group, but we knew we had a limited number of markers that we could track. We also understood that however we used the markers, the game had to run well with no full-scale testing. We wanted everyone to win the game as a single body, not as many small groups. This meant we had to discourage degenerate strategies.

We decided that the rules of the game should be discovered on the fly as people played, rather than through an instruction sheet. Finally, we wanted to create a game that was more rewarding than simply bouncing balls around in a space.

Playtesting was a major challenge we faced. Gathering a 4000-player audiences is difficult and expensive, so we had limited opportunities to test the game at full-scale. This led to a number of creative solutions in testing. Though we were able to use spaces of approximately the same size as the final space for testing, we were not able to test with a full-scale audience. We came up with a number of innovations, including clumping groups in various locations in the testing space and organizing the clumps strategically to make them appear to be a larger audience. However, the first test of the game with a full audience was not until its premiere at SIGGRAPH.

After many discussions and iterations with the prototyping system, we settled on the game rules described below.

Figure 7: Example Screen shot of Squidball game. Please see video for game in action.

5.1 The game

The rules for the game were simple, and had to be discovered by participants through gameplay. The twelve weather balloons in physical space were represented within the digital game space as green spheres on the screen. Players moved the weather balloons around the auditorium (whose space corresponded to a 3D space onscreen), in order to destroy changing grids populated by 3D target spheres. The game had 3 levels of increasing complexity, and each level could be replayed 3 times before a loss condition was reached. The second level introduced an element of time pressure, so players had to complete the game challenge within the allotted period of time. This was the level that really taught players how the game worked, as most audiences failed to clear the level on the first try.

Through repetition and the existence of a loss condition, the players eventually discovered the victory condition, as well as the correspondence between the weather balloons and their representation within the virtual game space. In the third level, players worked to clear special colored spheres that, when activated, revealed a composite image. Players quickly discovered a range of social strategies that emerged from their physical proximity with other players; in each instance of the game (the game was played 6 times over the course of 4 days) the 4,000 or so players came together organically to collaborate in the play of the game as they discovered what the gameplay required of them. It was an interesting first step in designing a kind of game that was extremely simple in its rules and interaction but extremely complex in the forms of social dynamics it spawned.

6 Gameplay in Practice

In the control both, we had a control which we could adjust to alter the sensitivity of the game during play. Turning up the sensitivity made the virtual target sizes larger and therefore gameplay easier. Turning down the sensitivity made the virtual target sizes smaller and gameplay harder.

Figure 8: Squidball Gamers

For good gameplay, we felt it was essential that the players be able to make mistakes and learn from them. So, by default, we set the sensitivity fairly low. However, in some circumstances we increased the sensitivity to temporarily make the gameplay easier, giving the audience a little ”boost”. A person in the control booth was responsible for watching the progress of the game and making these ”group mind” decisions regarding when to adjust gameplay. We believe that similar controls were included in Cinematrix

Even with a sensitivity control, there were some issues. One problem was that the audience at the start of show was less than a full house, which we had not anticipated in our game design. Because of this, some of the game levels proved hard to clear, because virtual targets were located in places where few audience members could reach them. This problem could be addressed by creating multiple configurations for different audience sizes.

Figure 9: Squidball Gamers

A second issue was uneven audience distribution. Some people were in sparse sections of the audience, and did not get to participate as much as others. To address this concern, we enlisted student helpers to help move balls around.

A third issue was that, during the game, people had to divide their attention between the screen and the balls. Some people decided to only watch the screen. Others ignored the screen and simply pushed the balls towards the center of the room. Initially, a relatively small percentage of ”aware” players actually watched both and drove the gameplay forward. The number of ”aware” players increased dramatically towards the end of the game, demonstrating that the game design principles were working. However, the split attention issue remains a challenge for any game design involving thousands of people, balls and a single screen.

Figure 10: Squidball Gamers

One solution might be to place multiple screens on all sides of the audience. However, this introduces another difficulty: coordination. Even with a single screen, players had difficulty coordinating the balls and target locations. The problem is that the player must face one direction, look at a screen over their shoulder, and then punch a ball in a third direction towards a target. Since few people have much practice at this activity, balls were popped left when they should have been popped right, or forward rather than back. Adding more screens would confound this issue.

7 Conclusions

After all the hard work to create and setup Squidball for SIGGRAPH 2004, the roar of the crowd at the end of each a level was gratifying validation of our efforts.

Since Squidball, we have been discussing possible iterations for future games. One option we have discussed is to use spotlights shone onto the crowd as targets, rather than using targets on a virtual screen. This would address some of the gameplay issues we encountered. The audience would have a physical cue showing where they are trying to get the balls to, rather than a virtual cue shown on a screen over their shoulder. Using spotlights, it would be possible to create roving patterns, enabling the spotlights to be moved in a pattern which ensures that everyone gets a chance to participate, taking into account the audience density and distribution. Of course, spotlights introduce a whole new set of technical challenges, though none that are insurmountable. We are considering this and other game design changes for Squidball 2.

And if you wonder what world record the team had to break (as mentioned in the introduction), it was about building the worlds largest motion capture space and producing the largest motion capture markers. We believe we broke it in August 2004.

8 Acknowlegments

This work has been partially supported by NYU, NSF, ACM SIGGRAPH, Vicon Motion Systems Ltd, Apple Computer Inc, Advanced Micro Devices Inc, Alienware, Cycling 74, David Rokeby: very nervous systems, NVIDIA Corp., Segway Los Angeles. The audio was supplied for “Insert Coin” and ”Re-Atair” by Skott, “Superkid” by Max Nix, and “Watson Songs” by The Jesse Styles 3000.

Furthermore we would like to thank especially our friends from AVW - TELAV, specificially Jim Irwin, Gary Clark, John Kennedy, Tom Popielski, Gerry Lusk, Mark Podany and Mike Gilstrap. Without their production efforts this would have been impossible. Also special thanks to Debbi Baum, Robb Bifano, Alyssa Lees, Jared Silver, Lorenzo Torresani, Gene Alexander, Boo Wong, Damon Ciarrelli, Jason Hunter, Gloaria Sed, Toe Morris, Scott Fitzgerald, Ted Warburton, Chris Ross, Eric Zimmerman, Cindy Stark, Brian Mecca, Carl Villanueva, and Pete Wexler for all their great help. And the SIGGRAPH chair Dena Slothower, who let us use their venue for the first big test, and for believing in us.

References

BLAINE, T. 2000. The outer limits: A survey of unconventional musical input devices. In *Electronic Musician*.

CARPENTER, L., 1993. Video imaging method and apparatus for audience participation. US Patent #5210604, #5365266.

FREEMAN, W. T., TANAKA, K., OHTA, J., AND K.KYUMA. 1996. Computer vision for computer games. In *2nd International Conference on Automatic Face and Gesture Recognition, Killington, VT, USA, IEEE*.

MAYENES-AMINZADE, D., PAUSCH, R., AND SEITZ, S. 2002. Techniques for interactive audience participation. In *IEEE Int. Conf. on Multimodal Interfaces, Pittsburgh, Pennsylvania*.

MOTIONANALYSISSTUDIOS, 2004. Largest mocap volume. http://www.motionanalysis.com/about_mac/50x50volume.html.

ULYATE, R., AND BIANCIARDI, D. 2004. The interactive dance club: Avoiding chaos in a multi participant environment. In *Int. Conf. on New Interfaces for Musical Expression*.

VICON, 2004. Sponsor of siggraph electronic theater. http://www.vicon.com.