A Synthetic-Aperture Camera Array
MIT9904-14

Progress Report: July 1, 1999 December 31, 1999

Leonard McMillan and Julie Dorsey

Project Overview

The goal of the Image Based Synthetic Aperture Rendering Project is to develop new image-based representations for use in computer graphics. Image-based representations are easy to acquire, yet they have similar flexibility and greater realism than traditional computer graphics models. We are working on various low-cost and real-time capture methods, interactive rendering algorithms, and autostereoscopic display devices for these image-based representations. The final system will be a complete capture-to-display system, giving the user imaging capabilities far beyond that of a traditional camera.

Progress Through December 1999

During the past six months, we have made significant advances on our project in the areas of capture, rendering, and display. For capture, we have 1) designed and built a low-cost light field scanner and 2) designed the sensor pod to be used for the video-rate light field camera. For rendering, we have implemented a real-time visualization tool for synthesizing novel views from a captured database. For display, we have built a static light field display that lets Multiple viewers perceive a three-dimensional scene, without any need for special viewing or tracking devices.

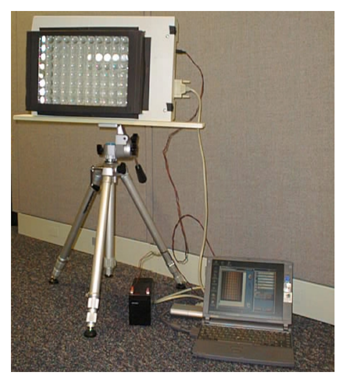

The first capture systems addresses the high costs and lack of portability of standard light field acquisition systems. The core device is an inexpensive flatbed scanner. The particular scanner we are using has a 600x1200 dpi optical resolution at 36 bits of color depth and costs less than $100. Mounted on the glass of the scanner is an 8x11 rectangular array of acrylic one-inch lenses. In order to make the scanner mobile, we use a 12V lead acid battery as a portable power supply and a laptop to interface with the scanner. We have disabled the internal fluorescent bulb so as not to interfere with environment lighting and to cause reflections on the lenses. Because we have used simple, low-cost optics and hardware, the challenges of the project lie in correcting and calibrating the scanned images.

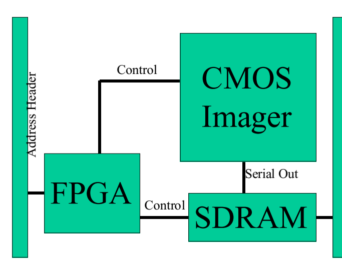

This first capture device, and nearly every light-field capturing device to date, are only able to acquire static scenes. We are working with NTT on a second capture system to extend light-field acquisition and rendering techniques for use with real-time dynamic scenes. To reduce the tremendous amount of bandwidth necessary to transmit a complete image database 30 times a second, we have designed a random access imager which only transmits the pixels that we need to generate the novel view. We use a color CMOS imager with sequential access that is connected to a SDRAM and a field programmable gate array (FPGA). The FPGA is the primary control logic, which 1) drives the imager, 2) manages the transfer of image data between the imager and memory, and 3) provides a dual-port access to the SDRAM memory buffer from the host computer. This memory-buffering scheme provides the illusion of random access necessary to support our rendering algorithms. To this date, the sensor pod FPGA circuitry has been designed.

To produce high-quality novel views at interactive rates, we have built a rendering system running on Microsoft Windows systems using DirectX 7. This renderer allows the user to create images with dynamically controlable position, orientation, focus, and aperture size. This permits synthetic depth-of-field and refocusing in real-time on commonly available rendering hardware. Our renderer uses projective texture-mapping hardware, now available for under $300.

We have developed a transformation that reorganizes light fields and allows them to be directly viewed using a simple display device. Multiple viewers will perceive the display as a three-dimensional scene, without any need for special viewing or tracking devices. Using our reparameterization technique, we can capture a scene using light field capture hardware, reparameterize it, and create a wide range of novel 3-D views that can be redisplayed with an $30 acrylic lens array. The scenes can exhibit considerable depth, both in front and behind the actual display surface.

Research Plan for the Next Six Months

We plan to make significant advances in the next six months in the areas of capture, rendering, and display. We plan on producing light fields from the low-cost scanner, building the sensor pods and a test bed for the real-time camera, increasing the capabilities of the renderer, and producing higher quality autostereoscopic displays.

Although we are able to capture images using our low-cost scanner camera, there are several problems with the image including an abnormal red hue, radial distortion effects, and parallax discrepancies. These problems are due to 1) limitations in the optics of in both the scanner and the attached plastic lenses and 2) lighting conditions not expected by the scanner manufacturer. Our strategy is to compensate and correct for these shortcomings with software. We will be developing post-processing software that automatically corrects all of these problems, thus providing a low-cost portable system for capturing reference images that can be used by our rendering system.

For the real-time camera, we plan on 1) doing the circuit board layout for the sensor pods and 2) fabricating and testing a small number of sensor pods. We also plan to build a test pod, which will simulate the interface between the sensor pod and the camera motherboard. This will allow us to test out our random-access imager design. Because the sensor pod control logic is implemented on a field programmable gate array, we will easily be able to make modifications to our design.

For rendering, we plan on 1) increasing the texture bandwidth of the real-time renderer so that we can render from larger datasets and 2) improving the accuracy of the renderings by using more accurate, hardware accelerated projective texture mapping. Furthermore, we will start design on an efficient renderer that will work with the scanner camera (the texture mapping method will not be sufficient for the real-time camera).

Although we have presented prototypes for the autostereoscopic display, we would like to improve on the quality and visual impact of the three-dimensional images. In addition, we will be exploring the possibilities of building a video display using our direct light field display technology.

Over the next six months, we hope to be much closer to our vision of a complete capture-to-display system. Given our successes during the past six months, we have every expectation to meet these goals.