

A Synthetic-Aperture Camera Array

MIT9904-14

Progress Report: January 1, 2001 - June 30, 2001

Leonard McMillan and Julie Dorsey

Project Overview

We are exploring a new approach to computer graphics where images play a central role in the modeling and rendering processes. We have demonstrated computer-generated animations whose realism is comparable to a photograph, yet having the same flexibility of navigation and interaction as classical three-dimensional computer graphics models. These "image-based" representations are both easy to acquire and generate. They consist of a collection of images and their associated camera calibration information. Our rendering system proceeds by interpolating the set of rays needed to construct a virtual image. The interpolation process is controlled via a synthetic aperture camera model coupled with a dynamically variable user-specified focal plane. Our research has focused on end-to-end solutions for image-based rendering, including devices for acquisition, rendering algorithms, and three-dimensional display technologies.

In addition to developing new rendering algorithms and representations for image-based models, we plan, as part of this project, to construct a real-time system for acquiring dynamic image-based models. To date, nearly all research in image-based rendering has been limited to static scenes. In order to construct a system capable of rendering dynamic models it is necessary to build a two-dimensional camera array. The construction of such an apparatus would allow a new class of applications including three-dimensional teleconferencing and holographic television.

Progress Through June 2001

This spring we have constructed an initial prototype of our camera array system. Our current system is composed of two linear arrays of 8 cameras, each approximately half a meter in length. This configuration has allowed us to experiment with both larger 1-D arrays (16 by 1) and smaller 2-D arrays (8 by 2). Our prototype system connects to a host PC via an IEEE-1394 (Firewire) interface. Currently, each linear camera array is mounted onto a plexiglass backplane. We are using a single host computer for all of the cameras via multiple firewire input cards and hubs. We have demonstrated real-time streaming of the video from 16 cameras at 320x240 resolution at 30 fps using 4 firewire interface cards. The image resolution and frame rate can be adjusted to maximize particular features (i.e. higher resolution, but slower frame rate). We also have developed software to synchronize, preview, and dynamically switch between the multiple video streams. We have also completed both intrinsic and extrinsic calibration of the camera arrays. At this time we are modifying our image-based rendering algorithms to accommodate input from the real-time video streams of the camera array.

Research Plan for the Next Six Months

We have also recently we received the development kits for SoundVision's Clarity based digital cameras. Clarity is a microprocessor based on an ARM core developed by Soundvision, Inc. that is intended as a single-chip digital camera solution. Currently we are developing software to run image-based rendering algorithms on the Clarity development kit, and we are completing our camera pod design. Our goal is to develop one working camera pod prototype by this fall. Once the single camera module has been manufactured and tested, we will replicate the module to construct a second-generation camera array to replace our firewire design. The advantages of using the Clarity cameras over our existing prototype design is that we will be able to move significant parts of the image-based rendering algorithms on to the CPU controlling the camera. This will allow for better bandwidth utilization and support multiple simultaneous virtual streams from the array. We will continue to use firewire as the system interface for our second-generation system. The new system will transmit the interpolated data from virtual cameras over the firewire bus, whereas, our current system transfers all of the video data from each camera in the array to the host. The host is responsible to constructing virtual views. We are expecting a performance enhancement in the range 1-2 orders of magnitude for this architectural improvement.

In the next six months, we would like to extend our firewire prototype system to at least 32 cameras, and complete the fabrication of our programmable camera pod. We will also integrate our prototype array with our real-time image-based rendering software and demonstrate the first ever dynamic image-based renderings. We are hoping that our prototype system will provide sufficient results to warrant publication in the December time frame.