

High Performance Imaging Using Large Camera Arrays

Bennett Wilburn^{1*} Neel Joshi^{2†} Vaibhav Vaish² Eino-Ville Talvala¹ Emilio Antunez¹
Adam Barth² Andrew Adams² Mark Horowitz¹ Marc Levoy²

¹Electrical Engineering Department ²Computer Science Department
Stanford University Stanford University

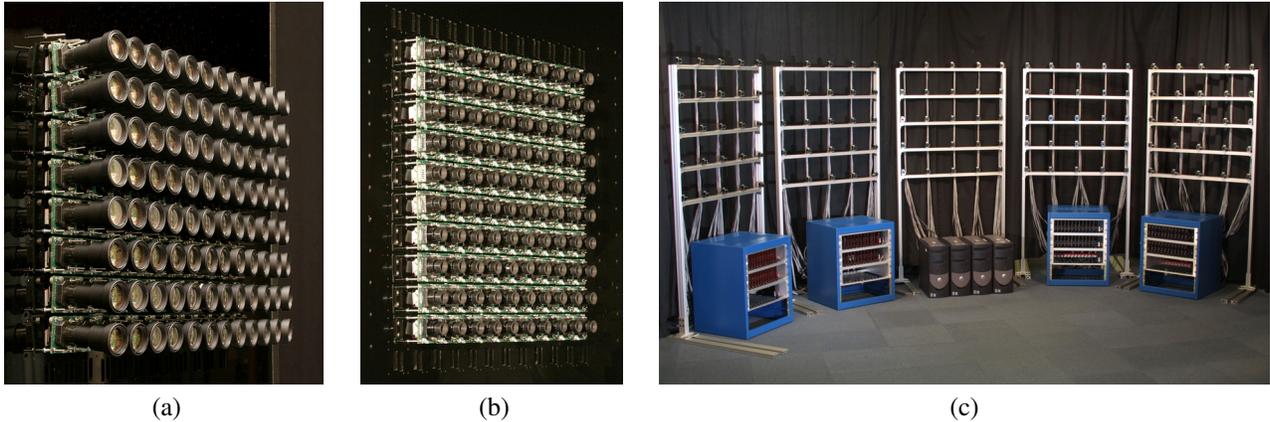


Figure 1: Different configurations of our camera array. (a) Tightly packed cameras with telephoto lenses and splayed fields of view. This arrangement is used for high-resolution imaging (section 4.1). (b) Tightly packed cameras with wide-angle lenses, which are aimed to share the same field of view. We use this arrangement for high-speed video capture (section 4.2) and for hybrid aperture imaging (section 6.2). (c) Cameras in a widely spaced configuration. Also visible are cabinets with processing boards for each camera and the four host PCs needed to run the system.

Abstract

The advent of inexpensive digital image sensors and the ability to create photographs that combine information from a number of sensed images are changing the way we think about photography. In this paper, we describe a unique array of 100 custom video cameras that we have built, and we summarize our experiences using this array in a range of imaging applications. Our goal was to explore the capabilities of a system that would be inexpensive to produce in the future. With this in mind, we used simple cameras, lenses, and mountings, and we assumed that processing large numbers of images would eventually be easy and cheap. The applications we have explored include approximating a conventional single center of projection video camera with high performance along one or more axes, such as resolution, dynamic range, frame rate, and/or large aperture, and using multiple cameras to approximate a video camera with a large synthetic aperture. This permits us to capture a video light field, to which we can apply spatiotemporal view interpolation algorithms in order to digitally simulate time dilation and camera motion. It also permits us to create video sequences using custom non-uniform synthetic apertures.

*email:wilburn@graphics.stanford.edu

†Neel Joshi is now at the University of California, San Diego.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—imaging geometry, sampling; C.3 [Computer Systems Organization]: Special Purpose and Application-Based Systems—real-time and embedded systems

Keywords: camera arrays, spatiotemporal sampling, synthetic aperture

1 Introduction

One of the economic tenets of the semiconductor industry is products that sell in large volumes are cheap, while products that sell in lower volumes are more expensive, almost independent of the complexity of the part. For computers, this relationship has changed the way people think about building high-end systems; rather than building a custom high-end processor, it is more cost effective to use a large number of commodity processors.

We are now seeing similar trends in digital imaging. As the popularity of digital cameras grows, the performance of low-end imagers continues to improve, while the cost of the high-end cameras remains relatively constant. In addition, researchers have shown that multiple images of a static scene can be used to expand the performance envelope of these cameras. Examples include creating images with increased resolution [Szeliski 1994] or dynamic range [S.Mann and R.W.Picard 1994; Debevec and Malik 1997]. In other work, Schechner and Nayar used spatially varying filters on a rotating camera to create high-resolution panoramas that also had high dynamic range or high spectral resolution [Schechner and Nayar 2001]. Another use for multiple views is view interpolation to create the illusion of a smoothly moving virtual camera in a static

or dynamic scene [Levoy and Hanrahan 1996; Gortler et al. 1996; Rander et al. 1997; Matusik et al. 2000].

Most of these efforts employ a single moving high-quality camera viewing a static scene. To achieve similar results on dynamic scenes, multiple cameras are required. This motivated us in 1999 to think about designing a flexible array containing a large number of inexpensive video imagers. The multiple camera array that resulted consists of 100 video cameras, each connected to its own processing board. The processing boards are capable of local image computation, as well as MPEG2 compression.

In section 2, we review prior work in building multiple video camera systems. While these systems are generally directed at specific applications, they provide valuable insights into the requirements for a flexible capture system. Section 3 gives an overview of our multiple camera array and explains in a little more depth the features we added to make it a general purpose research tool.

The rest of this paper focuses on our recent results using the camera array in different imaging applications. We start by exploring ways of using multiple cameras to create an aggregate virtual camera whose performance exceeds the capability of an individual camera. Since these applications intend to approximate a camera with a single center of projection, they generally use densely packed cameras. In particular, section 4 explores the creation of a very high-resolution video camera in which the cameras are adjusted to have modestly overlapping fields of view. We then aim the cameras inward until their fields of view overlap completely, and we use our system's fine timing control to provide a virtual video camera with a very high frame-rate. In both of these applications, the large number of cameras provide some opportunity that would not be present in a single camera system. For the virtual high-resolution imager, one can perform exposure metering individually on each camera, which for scenes with spatially varying brightness allows us to form a mosaic with high dynamic range. For the virtual high-speed imager, one can integrate each frame for longer than one over the frame-rate, thereby capturing more light per unit time than is possible using a single high-speed camera.

Sections 5 and 6 consider applications in which the cameras are spread out, thereby creating a multi-perspective video camera. One important application for this kind of data is view interpolation, whose goal is to move the virtual observer smoothly among the captured viewpoints. For video lightfields, the problem becomes one of spatiotemporal interpolation. Section 5 shows that the optimal sampling pattern to solve this problem uses cameras with staggered, not coincident, trigger times. It also describes a spatiotemporal interpolation method that uses a novel optical flow variant to smoothly interpolate data from the array in both time and virtual camera position.

In section 6 we consider combining the images from multiple viewpoints to create synthetic aperture image sequences. If we align, shift, and average all the camera images, then we approximate a camera with a very large aperture. By changing the amount of the shift, we can focus this synthetic camera at different depths. Using the processing power on each camera board, we can focus the synthetic aperture camera in real time, i.e. during video capture. Alternatively, we can shape the aperture to match particular characteristics of the scene. For example, we freeze a high-speed fan embedded in a natural scene by shaping the aperture in both time and space.

2 Early Camera Arrays

The earliest systems for capturing scenes from multiple perspectives used a single translating camera [Levoy and Hanrahan 1996] and were limited to static scenes. Dayton Taylor extended this idea to a dynamic scene by using a linear array of still cameras [Taylor 1996]. By triggering the cameras simultaneously and hopping from one camera image to the next, he created the illusion of virtual camera movement through a "frozen" dynamic scene. Manex Entertainment used more widely spaced cameras and added an adjustable trigger delay between cameras to capture images corresponding to a virtual high-speed camera flying around their scenes. Both of these systems used still cameras, so they were limited to capturing one specific virtual camera trajectory through space and time that was fixed by the camera arrangement.

For capturing a more general data set, researchers turned to arrays of video cameras. Like still cameras, video cameras must be synchronized, but they also present a new challenge: enormous data rates. The pioneering multiple video camera array design is the Virtualized RealityTM project [Rander et al. 1997]. Their goal was to capture many views of a scene for video view interpolation. The first version of their system records video using VCRs, giving them practically unlimited recording durations but low quality. Their second version uses 49 video cameras capturing to PC main memory. This system has better quality (VGA resolution at 30 frames per second), but is limited to nine-second capture durations. Every third camera captures color video. To handle the bandwidth of the video cameras, they require one PC for every three cameras.

While the Virtualized RealityTM project uses relatively high quality cameras, two other groups experimented with large arrays of inexpensive cameras. Yang et al.'s Distributed Light Field Camera renders live dynamic light fields from an 8x8 array of commodity webcams [Yang et al. 2002]. Zhang and Chen's Self-Reconfigurable Camera Array uses 48 commodity Ethernet cameras with electronic horizontal translation and pan controls to improve view interpolation results [Zhang and Chen 2004a; Zhang and Chen 2004b]. Although the design of these systems make them much cheaper than Virtualized RealityTM in terms of per camera costs, significant compromises were made to use these commodity cameras. First, neither of the arrays could be synchronized, causing artifacts in the view reconstructions. Furthermore, since they were looking at single applications, neither system addressed the bandwidth challenges of building a general purpose large camera array. Yang et al. chose to implement a "finite-view" system, meaning each camera transmits only enough data to reconstruct a small number of light field views per frame time. Zhang and Chen's cameras use JPEG compression, but their choice of Ethernet and a single computer to run the array limits them to a resolution of 320x240 pixels at 15-20 frames per second.

Results from these efforts helped guide our system design. Since our goal was to create a general purpose system, we wanted tight control over both the timing of cameras and their positions. We also needed to be able to record the data from all the cameras, but with far fewer PCs than the Virtualized RealityTM system. The system that we designed to address these goals is described next.

3 The Multiple Camera Array

While we had wanted to use "off-the-shelf" technology to build our camera array, it became clear early on that none of the commercial video cameras would have both the timing and positioning flexibility that our system required. As a result, we decided to build

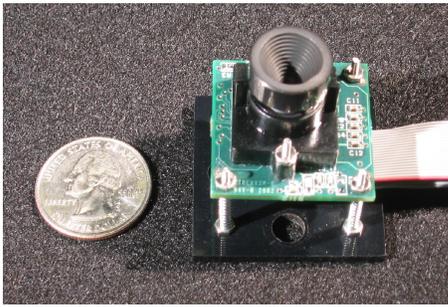


Figure 2: Our camera tiles contain an Omnivision 8610 image sensor, passive electronics, and a lens mount. The ribbon cables carry video data, synchronization signals, control signals, and power between the tile and the processing board. To keep costs low, we use fixed-focus, fixed-aperture lenses.

a custom imaging array, but one in which we leveraged existing standards as much as possible to minimize the amount of custom hardware that the system required for operation. A description of a preliminary version of this system was published in Wilburn et al. [2002].

3.1 Hardware Components

Our system consists of three main subsystems: cameras, local processing boards, and host PCs. The cameras are mounted on small printed circuit boards to give us maximum flexibility in their arrangement. Each camera tile is connected to a local processing board through a 2m long ribbon cable. These processing boards configure each of the cameras and can locally process the image data before sending it out to the host computer in either its raw form or as an MPEG2 video stream. A set of 4 PCs hosts the system, either storing the collected data to disk, or processing it for real time display.

Camera Tiles. One of the most critical decisions for the array was the choice of image sensors and their optical systems. While we thought it was reasonable to assume that computation would continue to get cheaper, we found it more difficult to make that same argument for high-quality lenses. Thus, we choose to use inexpensive lenses and optics as well as inexpensive sensors. In particular, we chose CMOS image sensors with Bayer Mosaic color filter arrays [Bayer 1976]. Although they have more image noise than CCD imagers, CMOS sensors provide a digital interface rather than an analog one, and they offer convenient digital control over gains, offsets, and exposure times. This makes system integration easier.

Figure 2 shows one of our camera tiles. For indoor applications, one typically wants a large working volume and a large depth of field. For these reasons, we use Sunex DSL841B lenses with a 6.1mm focal length, an F/# of 2.6, and relatively wide diagonal field of view of 57° . For applications that require a narrow field of view (usually outdoors), we use Marshall Electronics V-4350-2.5 lenses with a 50mm fixed focal length, an F/# of 2.5, and a diagonal field of view of 6° . Both sets of optics include an IR filter.

The camera tiles measure 30mm on a side and mount to supports using three spring-loaded screws. These screws not only hold the cameras in place but also let us change their orientations roughly 20° in any direction. For tightly packed camera arrangements, we mount the tiles directly to sheets of acrylic. For more widely spaced arrangements, we have designed plastic adapters that connect the tiles to 80/20 (an industrial framing system) components.

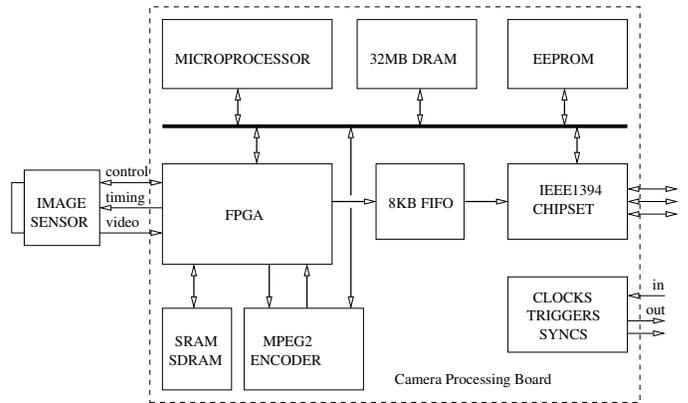


Figure 3: Camera processing board block diagram

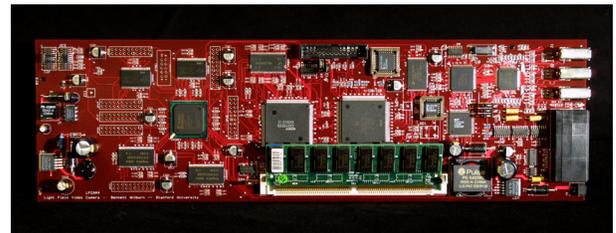


Figure 4: Camera processing board

Local Processing Boards. Figure 3 shows a block diagram of a complete camera system, and figure 4 shows the processing board for one camera. The processing board has five major subsystems: a micro-controller and its memory, an MPEG2 compressor, an IEEE1394 interface, a clock interface, and an FPGA which acts as master data router and programmable image computation unit. By choosing established standards, most of these subsystems could be implemented with existing off the shelf chip sets.

We chose the IEEE1394 High Performance Serial Bus [Anderson 1999] (also known as FireWire® and i-Link®) as our interface between the processing boards and the PCs. It guarantees a default bandwidth of 40MB/s for isochronous transfers, i.e. data that is sent at a constant rate. This is perfect for streaming video, and indeed many digital video cameras connect to PCs via IEEE1394. It is also well suited for a modular, scalable design because it allows up to 63 devices on each bus and supports plug and play. Another benefit of IEEE1394 is the cables between devices can be up to 4.5m long, and an entire bus can span over 250m. Thus, cameras based on such a system could be spaced very widely apart, possibly spanning the side of a building.

Even with this high-speed interface, an array of 100 video cameras (640x480 pixel, 30fps, one byte per pixel, Bayer Mosaic) would require roughly 25 physical buses to transfer the roughly 1GB/sec of raw data, and a comparable number of PCs to receive it. Rather than limiting the image size or frame rate, we decided to compress the video using MPEG2 before sending it to the host. The default 4Mb/s bitstream produced by our SONY encoders translates into a compression ratio of 17.5:1 for 640x480, 30fps video. To ensure that compression does not introduce artifacts into our applications, we designed the cameras to simultaneously store up to 20 frames of raw video to local memory while streaming compressed video. This lets us compare MPEG2 compressed video with raw video as an offline sanity check.

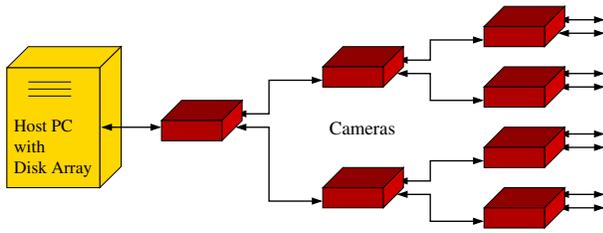


Figure 5: Camera array architecture

An embedded microprocessor manages the components in the camera and communicates with the host PCs over IEEE1394. The FPGA is used to route the image data to the correct destination, usually either the IEEE1394 chipset or the MPEG2 compression chip. It can also be configured to operate directly on the image data using its local DRAM for storing temporaries and constants and the SRAM as a frame buffer. Code in a small boot ROM configures the IEEE1394 interface so that host PCs can download a more sophisticated executable and configuration code to the board.

3.2 System Architecture

Figure 5 shows the high-level architecture of our system. Each of our cameras is a separate IEEE1394 device with three ports. The cameras are connected in a tree, with one port connecting to a parent and one or two ports leading to child nodes. The parent port of the root node is connected to the host computer, which has two striped IDE hard drives to capture the image data. For large arrays, we must use multiple PC's and IEEE1394 buses. Theoretically, the 40MB/s streaming bandwidth of IEEE1394 should accommodate 62 compressed video streams, but implementation details (bus arbitration and our inability to get cycle-accurate control over the bus) limits us to 30 cameras per bus. We run a networked camera control application that lets us drive the operation of the entire array from one PC.

The timing requirements for the array were stricter than could be achieved using IEEE1394 communication, especially with multiple PCs. To achieve the desired timing tolerance, we route a common clock and trigger signals to the entire array using an extra set of CAT5 cables. These cables roughly match the IEEE1394 topology, except they form a single tree even if multiple IEEE1394 buses are used. A single "master" root board in the array generates its own 27MHz clock and sends it to two children via CAT5 cables, which then buffer the clock and send it to two more children, and so on. The master also generates a trigger which is buffered and repeated to all other boards. This trigger is used to synchronize the cameras and provides a timing signal with no more than 200ns of skew between any two processing boards. To put this in perspective, 200ns is one thousandth of our minimum integration time of 205 μ s.

Most systems would use the trigger to synchronize all of the cameras. In fact the early prototype of our system [Wilburn et al. 2002] used it for this purpose as well. The final system provides an arbitrary, constant temporal phase shift for each camera. Because the timing signals for the image sensors are generated by the FPGAs, this was done by adding programmable timer reset values to the FPGA code. Thus, using just one trigger signal, we can reset all of the cameras to arbitrary phase offsets.

3.3 Results

Our multiple camera array captures VGA video at 30 frames per second (fps) from 100 cameras to four PCs. The default MPEG bit rate is 4Mb/s, but we are free to alter the bit rate or even stream I-frame only video. At 4Mb/s, we can capture sequences up to two and a half minutes long before we reach the 2GB file size limit of our operating system. We have not yet needed to extend this limit.

4 Improved Imaging Performance

By combining data from an array of cameras, we can create an aggregate virtual camera with greatly improved performance. Although one could design optical systems that ensure a common center of projection for all of the cameras, these systems become costly and complex as the number of cameras grows. Instead, we pack the cameras as closely as possible to approximate a single center of projection and compensate for parallax in software. Here, we discuss two high-performance applications: high-resolution, high-dynamic range video capture; and high-speed video capture.

4.1 High-Dynamic Range and High-Resolution Video

If we tightly pack our cameras and aim them with abutting or partially overlapping fields of view, we create a high-resolution video camera. Using this configuration and existing techniques from the image mosaicing literature, we can register and blend the images to create a single image of high resolution. One advantage of using many cameras for this task is that we can meter them individually. This allows us to capture scenes with a greater dynamic range than our cameras can record individually, provided that the dynamic range in each camera's narrow field of view is small enough. For scenes in which even the local dynamic range exceeds our sensors' capabilities, we can trade resolution for dynamic range by increasing the overlap of the cameras' fields of view, so that each viewing ray is observed by multiple cameras with different exposure settings.

To demonstrate this idea, we arranged our cameras in a dense 12x8 array with approximately 50% overlapping fields of view, shown in figure 1(a). Each camera has a telephoto lens with a roughly six degree diagonal field of view. With 50% overlap between adjacent cameras, most points in the scene are observed by four cameras, and the entire array has a total field of view of 30 degrees horizontally and 15 degrees vertically.

Color Calibration. Because the inexpensive sensors in our array have varying color responses, we must color match them to prevent artifacts in the image mosaic. Color calibration is important in any application involving multiple cameras, but it is critical in this application, since different parts of the image are recorded by different cameras. We must also determine the response curves of our cameras if we wish to create high dynamic range images. With gamma correction turned off in the cameras, the response curves of our sensors are reasonably linear except at the low and high ends of their output range. We have devised an automatic color matching routine that forces this linear response to be identical for all of the cameras and color channels by iteratively adjusting the offsets and gains for each color channel in every camera. Our goal is to ensure uniformity, not absolute accuracy—our final mosaics can be converted to another color space with one last transformation.

Each iteration of our calibration routine takes images of a white target under several different exposure levels. The target is placed close enough to the array to fill the field of view of all cameras. The exposure setting is the actual duration for which the sensor integrates light and is very accurate. The routine calculates the slopes and offsets of the sensor responses, then computes new settings to match a target response. We choose a line mapping the minimum response to 20 and the maximum to 220, safely inside the linear range of our sensors. Doing this for each channel using images of a white target also white balances our sensors. The entire process takes less than one minute.

Assembling HDR Image Mosaics. We use Autostitch [Brown and Lowe 2003]) to create our image mosaics. Autostitch uses a scale-invariant feature detector to detect corresponding features in overlapping images, bundle adjustment to estimate globally optimal homographies to align all of the images, and a multi-band blending algorithm to combine the registered images into a single mosaic. The cameras need not be precisely aimed, because Autostitch finds appropriate homographies to perform seamless image stitching. Given the 34mm separation of our cameras and our scene, roughly 120m away, we can tolerate +/- 20m of depth variation with less than 0.5 pixels of disparity in the mosaiced image.

For our application, we have modified Autostitch in two ways. First, we use our response curves and the cameras' exposure durations to transform pixel values from the cameras into a floating point, relative irradiance value before blending. Thus, the output of the blending is a floating point image. Our second modification is replacing the weights for the multi-band blend with a confidence measure that is high for pixel values in the middle of the sensor response and low for saturated or underexposed pixels, as well as being low for pixels at the edges of each camera.

Results. Figure 6 shows a comparison of 3800 x 2000 pixel mosaics captured with uniform and individually selected camera exposure times. The uniform exposure loses details in the brightly lit hills and dark foreground trees. The individually metered cameras capture a wider range of intensities, but they still have saturated and under-exposed pixels where their dynamic range is exceeded. An even better picture can be acquired by taking advantage of the cameras' overlapping fields of view to image each point with different exposure durations. Figure 7 (a) shows a mosaic captured using cameras with one of four exposure times (0.20ms, 0.62ms, 1.4ms, and 3.07ms). The increased local dynamic range can be seen in the covered walkway in the inset (c).

To evaluate the overall image quality, we took a picture using a 3504 x 2336 pixel Canon 20D configured with nearly the same field of view and compared it to one frame of our high-resolution video (figure 7(b)). The results are encouraging. While the insets show that the Canon image is superior, the effective resolution difference is modest. Plotting pixel intensities across edges in the two images showed that the Canon's resolution is roughly 1.5 times better. Since we could easily add cameras, or reduce overlap to increase resolution, this degraded resolution is not a serious limitation. In fact, resolution chart measurements with our cameras indicate that their effective resolution is about 400 pixels horizontally, not 640, so the resolution of the mosaic is not much worse than what we see from a single camera.

What is more surprising is that the contrast of our image mosaic is noticeably worse than the D20. This is due to light leakage and aberrations in the lenses. Overall, these results show that it is possible to use large numbers of inexpensive cameras to build a virtual camera of both high dynamic range and high resolution. In this example we use large overlaps so four cameras view each pixel. Our array can easily be configured to reduce the overlap and create

larger mosaics. For example, reducing the camera overlap to 10% would yield very large mosaics (roughly 6900 x 3500 pixels) using the same number of cameras. (Remember that these are video cameras; we know of no non-classified video camera of comparable resolution.) This flexibility raises the question of how to optimally allocate camera views for imaging. This answer in turn depends on the dynamic range of the scene and the algorithm used for adaptively setting the exposure times. We are starting to look at adaptive metering algorithms for camera arrays to address this issue.

4.2 High-Speed Video

The previous application takes advantage of our flexible mounting system and exposure control to increase the resolution and dynamic range of video capture. The timing precision of our array offers another opportunity for creating a high-performance aggregate camera: high-speed video capture. We have previously described a method for configuring the array as a single, virtual, high-speed video camera by evenly staggering the camera trigger times across the 30Hz frame time [Wilburn et al. 2004]. Using 52 tightly packed cameras oriented with wholly overlapping fields of view, we simulated a 1560 frame per second (fps) video camera.

One benefit of using a camera array for this application is that frame rate scales linearly with the number of cameras. Also, compressing the video in parallel at each camera reduces the instantaneous data rate and permits us to stream continuously to disk for several minutes. By contrast, typical commercial high-speed cameras are limited to capture durations that fit in local memory, often as low as a few seconds, and require some means to synchronize the capture with the high-speed event. Finally, unlike a single camera, the exposure time for each frame can be greater than the inverse of the high-speed frame rate. In other words, we can overlap frame times among the cameras. This allows us to collect more light and reduce noise in our images at the cost of increased motion blur. By temporally deconvolving the captured video, we can recover some of the lost temporal resolution [Wilburn et al. 2004; Shechtman et al. 2002].

As with the image mosaics before, we must account for the slight parallax between views from different cameras. We assume a relatively shallow or distant scene and use planar homographies to align the images from all cameras to the desired object plane. This leads to artifacts for objects not at the assumed scene depth. In the next section, we extend this high-speed method to the case of more widely spaced cameras, and in section 5.2 we describe a technique for interpolating between the views produced by the cameras. As we will see, this technique can also be used to correct the misalignments in our high-speed video.

5 Spatiotemporal Sampling

We now turn to a different regime for the array: cameras spaced to sample a very wide spatial aperture. Data captured from such arrangements can be used for synthetic aperture photography, view interpolation, and analysis of scene structure and motion. We treat synthetic aperture photography in section 6. For the other two applications, a major challenge is establishing correspondences between points in different views. Generally speaking, algorithms for computing correspondences perform better when the motion between views is minimized. In this section, we show how to reduce image motion between views of dynamic scenes by staggering camera trigger times. Section 5.2 describes a new view interpolation algorithm based on optical flow.



(a)

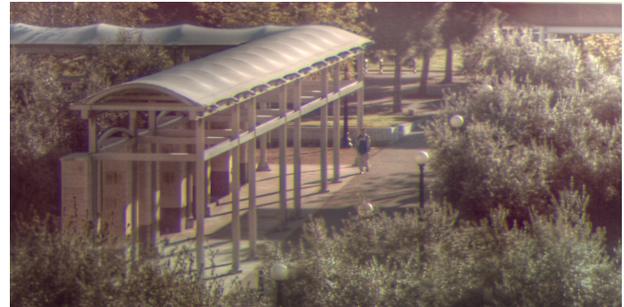


(b)

Figure 6: High Dynamic Range Panoramic Video. By metering cameras individually, we can increase the total dynamic range of the panoramic video. (a) In this image, all cameras are set to the same exposure. Notice the saturated areas in sunlight and dark regions in shade. (b) For this mosaic, each camera's exposure was set such that the average pixel value is in the middle of the sensor range, and the resulting high dynamic range image was tone mapped for display (and printing). More details are revealed, including the radar dish and hills on the horizon and dark areas in the foreground trees. The roof of the covered walkway, however, was outside the range of the cameras that viewed it. The gray color is due to tone mapping—we do not actually know how bright the roof should be. The sky in the top left of the panorama was also overexposed.



(a)



(b)



(c)



(d)

Figure 7: Comparison with a Canon 20D. (a) Setting the exposure times so each pixel is viewed by four cameras with varying exposure durations (0.20ms, 0.62ms, 1.4ms, and 3.07ms). This scheme increases the local dynamic range of the mosaic relative to figure 6(a) or 6(b). The inset (b) shows that we now have valid data for the covered walkway. The color variations along the borders of the panorama in (a) result from viewing those portions of the scene with fewer than four different exposures. This leads to artifacts in areas where we have no valid data. (c) An image of the same scene taken with a Canon 20D, which has a resolution of 3504 x 2336 pixels. (d) is the inset of the covered walkway from the Canon, for comparison. Our panorama has as much (or more) dynamic range as the Canon image. However, the Canon images are sharper and have more contrast than the panorama. The latter is due to stray light and aberrations in our relatively low-quality lenses.

5.1 Planar Camera Arrays

To reason quantitatively about view sampling in space and time, we will consider a planar camera array whose images are all aligned to a common fronto-parallel reference plane. This arrangement is used for light field rendering as well as many of the applications in this paper. More complicated surfaces can be tessellated to form triangles of cameras for which this analysis also applies. Given this framework, we ask two questions. First, what is the maximum possible image motion between two views from different positions and times? Second, how should we trigger fixed frame rate cameras to minimize image motion?

Figure 8 shows how motion in the images on the reference plane is related to the scene geometry and velocities. We assume the scene has near and far depth limits with signed distances Δz_{near} and Δz_{far} from the reference plane, and the reference plane is optimally placed at a depth Z_0 as described by Chai et al. [2000]. For a camera spacing of Δx , the parallax Δp in the aligned images for a point P at a distance Δz_p from the reference plane is $\Delta p = \Delta x \cdot \Delta z_p / (\Delta z_p + Z_0)$. If we define the “relative depth” d of the point to be $\Delta z_p / (\Delta z_p + Z_0)$, this simplifies to $\Delta p = \Delta x \cdot d$.

The worst-case parallax occurs at the near and far depth planes. The worst case temporal motion will occur if P is moving at the maximum velocity in the scene, v , on the near-depth plane, such that the vector $P_t P_{t+\Delta t}$ is orthogonal to the projection ray from C_0 at time $t + 1$. If we assume a narrow field of view for our lenses, we can approximate this with a vector parallel to the focal plane, shown as $v\Delta t$. If P has velocity v , the maximum temporal motion of its image in C_0 is $v\Delta t Z_0 / (Z_0 + \Delta z_{near})$. Equating this motion to the maximum parallax for P in a neighboring camera yields

$$\Delta t = \frac{\Delta x \Delta z_{near}}{v Z_0} \quad (1)$$

This is the time step for which maximum image motion between views at the same camera equals the maximum parallax between neighboring views. If we represent a view by two spatial (x, y) coordinates and one time coordinate t , measuring time in increments of the time step Δt and space in units of camera spacings provides a normalized set of axes to relate space-time views. Because motion due to parallax and temporal motion are not orthogonal, the true distance measure is the Euclidean spatial distance plus the temporal distance. Minimizing this distance measure between views minimizes the maximum image motion.

This metric gives us a method to optimize our distribution of samples in space and time. Figure 9 plots the (x, t) coordinates of captured views for a linear camera array with different values of Δx and Δt . Since the object motion is often not known a priori, we want a sampling that works for a wide variety of motion vectors. In scenes with little motion (figure 9(a)) the temporal pattern makes little difference, since the main image motion is from parallax. When object motion causes large image changes (figure 9(b)), synchronized time samples are one of the worst sampling patterns, since it creates dense rows of samples with large blank areas. In this case, the best timing for the cameras is one where the available time resolution increases with increasing parallax distance from the main sample. As shown in figure 9(b), across an array of N cameras, every one of the frame-time/ N possible starting times is used. Note that using this offset timing pattern does not hurt if scene velocities are small, because the changes in time make little difference in the images that are formed.

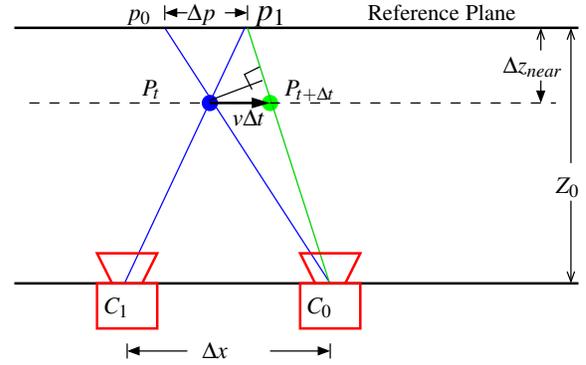


Figure 8: The temporal and spatial view axes are related by image motion. For a given scene configuration, we can determine a time step Δt for which the maximum image motion between temporal samples is equal to the maximum parallax between spatially neighboring views. If we measure time in increments of Δt and space in increments of the camera spacing, then the Manhattan distance between view coordinates corresponds to the maximum possible image motion between views.

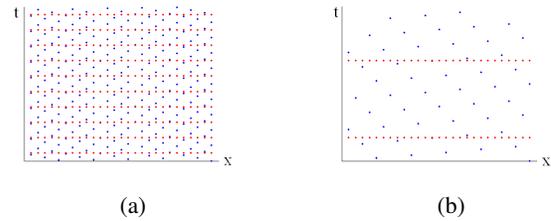


Figure 9: Plots showing (x, t) view coordinates for different amounts of object motion, and different sampling patterns. Both figures show a uniform time sampling in red and an optimal distribution of samples in blue. (a) For scenes with large camera spacings or very slow motion, time shifting of the cameras makes little difference. (b) For scenes with small camera spacings or high velocities, uniform sampling creates dense rows of samples and leaves most of the area unsampled. An optimized sample pattern starts each camera at $Q * i \bmod N$, where i is the index of the camera, N is the number of cameras, and Q is chosen to be roughly $1/3$ and also relatively prime with N .

5.2 Multibaseline Spatiotemporal Optical Flow

Better spatiotemporal sampling would improve even the simplest view interpolation algorithms like blending, but the sampling densities required for ghost-free images using blending are prohibitive. Instead, we created a novel optical flow variant for generating new views from a planar video camera array. Our modified spatiotemporal optical flow algorithm has two novel features. Optical flow typically computes flow between two images by iteratively warping one towards the other. Our first modification is to solve for a flow field at the (x, y, t) location of our desired virtual view. We were inspired to compute flow for the pixels in the new image by Kang et al. [2003]. They noted that for a video sequence, computing flow at a frame halfway between two images in a video sequence handles degenerate flow cases better and avoids the hole-filling problems of forward-warping when creating new views. We extend the method to compute flow at a desired view in our normalized (x, y, t) view space. We modified the robust optical flow estimator of Black and

Anandan [1993] using code available on the author’s web site. We iteratively warp the nearest four captured images toward the virtual view and minimize the weighted sum of pairwise robust data and smoothness error terms.

Motion cannot be modelled consistently for four images at different spacetime locations using just horizontal and vertical image flow. The second component of our algorithm is separately accounting for parallax and temporal motion. The standard intensity constancy equation for optical flow is:

$$I(i, j, t) = I(i + u\Delta t, j + v\Delta t, t + \Delta t) \quad (2)$$

Here, (i, j, t) represent the pixel image coordinates and time, and u and v are the horizontal and vertical motion at an image point. Our modified intensity constancy equation represents constancy between the virtual view and a nearby captured image at some offset $(\Delta x, \Delta y, \Delta t)$ in the space of source images:

$$I_{\text{virtual}}(i, j, x, y, t) = I_{\text{source}}(i + u\Delta t + d\Delta x, j + v\Delta t + d\Delta y, t + \Delta t) \quad (3)$$

The flow components are separated into parallax motion, determined by a point’s relative depth d and the spatial distance between views, and temporal motion, the product of the time between views and the projection (u, v) of the temporal motion onto the image plane.

For each virtual view, we choose input views for the flow algorithm by computing a three-dimensional Delaunay triangulation of the camera sampling points and selecting the views from the tetrahedron which encloses the desired (x, y, t) view. These images are progressively warped toward the common virtual view at each iteration of the algorithm. We cannot test the intensity constancy equation for each warped image against a virtual view. Instead, we minimize the error between the four warped images themselves using the sum of the pairwise robust intensity constancy error estimators. This produces a single flow map, which can be used to warp the four source images to the virtual view. We currently do not reason about occlusions and simply blend the warped images using their barycentric weights in the tetrahedron.

Results. For our experiments, we configured the cameras in a 12-by-8 array with a three inch camera spacing. We determined experimentally that nine staggers across the 30Hz frame time would be sufficient for our scene, so we created a 3x3 grid of triggers that is locally uniform and replicated it across the array. Because our application compares neighboring images, locally uniform sampling is sufficient. We calibrated our cameras to determine their relative displacements in the camera plane using the plane plus parallax framework described by Vaish et al. [2004].

Figure 10 shows the results of improved spatiotemporal sampling and our view interpolation algorithm. For reference, we show a cross-dissolve between two subsequent frames from one camera to illustrate the temporal motion between frames. Cross-dissolves, or blending, are the simplest interpolation method for arrays of cameras synchronized to trigger simultaneously. Staggering the camera trigger times to sample more uniformly in space-time improves even this simple interpolation method. Figure 10(b) shows a weighted blend of four views from the same array with staggered trigger times. The ghosting is greatly reduced. Finally, the image on the right shows the results of our multibaseline spatiotemporal optical flow algorithm. Because the computed flow is consistent for the four views, when the source images are warped and blended, the ball appears sharp.

Discussion. We used improved sampling to create a relatively simple interpolation method that uses optical flow to account for both parallax motion and true object motion in the scene. This method

allows us to estimate any camera image that is inside the time and spatial extent of the original camera area. If we hold the virtual viewpoint steady and synthesize new views at each trigger time, we produce a registered high-speed video. We are free, however, to alter the virtual view position and time arbitrarily (within the span of the array), enabling both time dilation and virtual camera motion.

While our spatiotemporal works well in practice, it does occasionally suffer from the usual artifacts of optical flow, such as large dominant motions masking the motion of smaller regions and problems when the image motion is too large. Thus as camera spacings increase, more sophisticated methods will be required to interpolate new views. Many methods developed to work with synchronized cameras should benefit from using cameras with more optimal sample timing. For example, segmentation-based stereo methods have recently been proven very useful for spatial view interpolation [Zitnick et al. 2004] and analysis of structure and motion in dynamic scenes [Tao et al. 2001; Zhang and Kambhampettu 2001]. Because these methods match small image regions across views, one would expect them to benefit from reduced image motion between nearby space-time views.

The high-resolution video capture application divided the total mosaic resolution by four to increase the dynamic range. By contrast, staggered camera triggers increase temporal sampling resolution with essentially no cost. Thus, we believe that staggered timing for video cameras arrays is always beneficial. If scene velocities are small, the temporal offsets are inconsequential. If the velocities are large, staggered cameras can capture events that would otherwise go unnoticed, minimizing interpolation artifacts.

6 Synthetic Aperture Photography

Spatiotemporal view interpolation simulates a narrow moving aperture in space-time. If instead of interpolating views, we align the images taken across the aperture to a plane and average them together, we approximate a camera with a very large aperture. Shifting the aligned images varies the focal depth for the system [Levoy and Hanrahan 1996; Isaksen et al. 2000; Vaish et al. 2004]. Warping them in addition to shifting them permits the focal plane to be tilted [Vaish et al. 2005]. In these experiments, we accelerate the computation by having the FPGA in each camera align and shift the video before it is compressed and sent to the host PCs. This gives us a real-time (live) synthetic aperture videography system. Specifically, as the user interactively adjusts the object focal depth, the host PCs broadcast the required image shifts to the cameras. Currently, the processing power of our host PCs limits us to 15 video cameras per PC.

The aperture of a traditional camera is a cylinder in space and time. The height corresponds to the exposure time and the cross section is the shape of the lens aperture. Synthetic aperture photography increases the spatial extent of the aperture by sampling it with many cameras. We now consider two exotic aperture shapes made possible by our array. The first, matted synthetic aperture photography, tailors an aperture to only capture rays that see through a partial occluder. The second creates a hybrid space-time aperture that images with high depth of field and low motion blur in low-light conditions.

6.1 Non-linear Synthetic Aperture Photography

The synthetic aperture camera effect permits one to see a subject hidden behind partial occluders by blurring the occluder across the

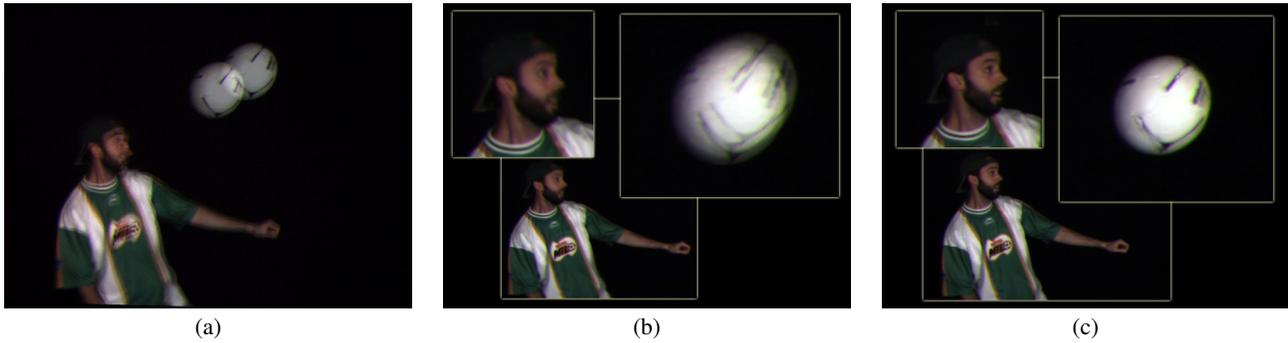


Figure 10: Better spatiotemporal sampling improves view interpolation. (a) A simple cross dissolve between two subsequent frames from one 30Hz camera. (b) Synchronizing the cameras with staggered trigger times increases our view sampling density in space and time. This view, created using a weighted average of four input views, shows much less ghosting. (c) Better spatiotemporal view sampling reduces image motion between views, making optical flow more robust. Here, we warp the four source images to the desired virtual view using our multibaseline spatiotemporal optical flow algorithm. The warped images are blended using the same weights as in the center image. No double images are present because parallax and motion for the ball are correctly recovered.

image. However, the occluder is not rendered invisible, and the synthetic aperture photograph attenuates the signal of interest, i.e. the subject. Suppose that N cameras view the scene with measurement noise ϵ . To create the synthetic aperture image, we align the views from all cameras to one plane and average them together. If only K cameras see through the occluder to any given point on the subject, then the signal in the synthetic aperture image is attenuated by a factor of K/N , while the measurement noise falls by $1/\sqrt{N}$. Thus, the SNR has fallen by at least K/\sqrt{N} relative to the SNR of a single image. Since the occluder does not completely average out, it will add an additional noise component.

If we knew, for each camera, which pixels saw through the partial occluder to the subject, we could average only the contributions from the unoccluded pixels. Averaging just the K unoccluded pixels would increase the SNR of a single image by \sqrt{K} and does not reduce the contrast of image by attenuating the signal. In practice, many pixels are mixture pixels, containing information from both the foreground and the background, so the SNR improvement will be smaller than \sqrt{K} .

To implement this, we create a binary matte image for each camera. The matte is one for pixels which are not blocked by the occluder and zero otherwise. Although binary mattes discard information, in order to use fractional (i.e. alpha) values, we must also recover the foreground color. The binary matte is a robust, conservative solution. To create the matted synthetic aperture image, we divide the sum of the aligned, matted input images by the sum of the aligned mattes at each pixel.

There are several ways one might imagine creating the occlusion mattes. One that we have implemented identifies all of the pixels that vary significantly over time in video from each camera. Barring motion of the occluder and interreflections between the occluder and the subject behind it, these pixels capture some time-varying portion of the subject and hence are not occluded. We identify these pixels by computing the variance of each pixel over each second of input video and thresholding.

Results. Figure 11 shows the results of our matted synthetic aperture method filming people through foliage. By shaping the aperture to reduce contributions from occluders, matted synthetic aperture produces a more accurate image of the hidden subjects. Mixture pixels prevent the occluder from being eliminated entirely, and spaces where no rays get through are left black. We compared the mattes we produced using the image variance in time with a

“ground truth” matte we constructed by imaging white and black backgrounds placed behind the occluder. We found little discernable difference in using the two mattes.

Discussion. As we have seen in this section, customizing the rays that contribute to a synthetic aperture image can lead to significant improvements. Computing mattes based on the temporal variance of each input pixel works well for static occluders. We are interested in extending our techniques to handle moving occluders using other matting techniques. Some possibilities include matting based on color thresholding for homogeneous occluders, shape from stereo or focus, and active range finding.

So far we have shown how to shape the aperture in space, but there is no reason we could not shape the aperture in both time and space. For example, if we could estimate the motion of partially occluded subjects, we could shape a space-time synthetic aperture that follows the object’s path. This aperture should generate an even better image, where information present in some views could be added to views where it is missing. This section showed one way to customize an aperture for a specific problem. In the next section, we extend this idea to shaping an aperture in both time and space.

6.2 Hybrid Aperture Photography

Traditional cameras have two means of collecting more light: increasing the exposure time and increasing the lens aperture diameter. Both have side effects. Increasing the exposure time increases motion blur for fast-moving objects, and increasing the aperture diameter results in a smaller depth of field. Thus, to photograph a fast-moving object embedded in a wide depth of field, stationary or slowly moving scene, one would prefer to use a small aperture diameter and short exposure times. If the scene is not brightly illuminated, this can result in dark, noisy images.

As noted earlier, our array is not limited to cylindrical space-time aperture functions. We can partition our array into subarrays, thereby simultaneously capturing images of a scene using multiple different apertures. By combining the images captured through these different apertures, we effectively create a “hybrid” aperture, allowing us to properly photograph these scenes. As an example of this idea, in figure 12, we consider the problem of photographing a spinning fan in the middle of a deep room. To create a hybrid aperture specialized for this scene, we simultaneously image the scene through the three following apertures:

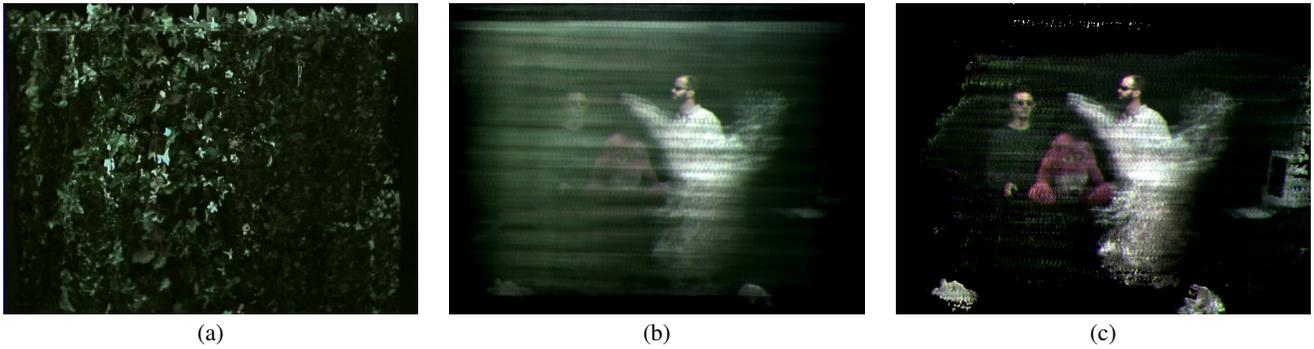


Figure 11: Matted synthetic aperture photography. (a) A sample image from one of 90 cameras used for this experiment. (b) The synthetic aperture image focused on the plane of the people, computed by aligning and averaging images from all 90 cameras as described in the text. (c) Suppressing contributions from static pixels in each camera yields a more vivid view of the scene behind the occluder. The person and stuffed toy are more clearly visible.

- A spatially narrow, temporally wide aperture, s . The entire scene will be in focus, but will have motion blur for fast-moving objects. The image through a single camera with a small aperture lens and long exposure implements this aperture.
- A spatially wide, temporally narrow aperture, t , focused on the subject. The narrow temporal aperture eliminates motion blur, but the large spatial aperture means everything not at the subject’s depth will be out of focus. We capture this using a synthetic aperture photograph taken by cameras configured with short exposure times.
- A spatially and temporally wide aperture, w . This image will have both defocus blur due to limited depth of field and motion blur for the subject. We acquire this using a synthetic aperture photograph taken by an interleaved array of cameras with long exposure times.

Figure 12 shows the images I_s , I_t , and I_w captured through the apertures s , t , and w . Each of these apertures collects much more light than would be collected by a camera with a spatially and temporally small aperture. Observe that I_s has motion blur for the fast-moving subject (the fan), I_t has defocus blur for everything not at the depth of the subject, and I_w has both. Because I_w is focused at the subject, the motion blur of the subject is in focus and therefore identical to the motion blur in I_s . Similarly, because the two synthetic aperture photographs are focused at the same depth, the defocus blur for the rest of the scene is equivalent in both images. Therefore, we can compute our desired image from $I_s + I_t - I_w$, after normalizing each image for exposure, as shown in (d).

The synthetic aperture images show aliasing artifacts because we are point sampling the spatially wide apertures. In order to capture the two images simultaneously, we assigned half of the cameras in our array to one synthetic aperture and the remainder to the other, setting aside one camera for the spatially narrow aperture. We interleaved the two sets of synthetic aperture cameras in a checkerboard pattern on our planar array, but the slight displacements between views caused slight changes in the aliasing of the synthetic aperture images. The differences in the aliasing remain after subtracting I_w from I_t and cause artifacts.

Aliasing appears only in the defocused regions of I_t and I_w . In the final image, we wish the defocus blur to cancel. If we knew where the aliases appeared in I_t and I_w , we could matte out the defocused regions prior to composing the final image. We can construct such a matte from a depth map of the scene.

To reconstruct an alias-free synthetic aperture image, we first apply an appropriate reconstruction filter to the samples of w . This filter removes high-frequency components of the scene along with the aliases. We estimate the depth of features that survive this filtering by computing the variance across the synthetic aperture samples at each pixel. If we assume textured objects in the scene, variance will be high for objects not at the focal depth. We obtain a matte by thresholding this variance image. In practice, many objects do not have high frequency textures, but low frequency textures do not create aliases, so the technique is robust for our purposes.

Figure 12 (e) is the result of matting I_t and I_w before computing $I_s + I_t - I_w$. The aliasing artifacts are gone, and we have achieved both high depth of field for the scene and low motion blur for the fan. The last picture (f) is the image taken through an aperture of narrow spatial and temporal extent (i.e. one camera with a short exposure time). The motion of the fan is frozen and the statue is in focus, but the result is much noisier than the hybrid aperture image.

It is interesting to compare our approach to that of Stewart et al. [2003], which proposes a hybrid reconstruction filter for light field rendering in order to reduce “ghosting” artifacts. Their filter combines a wide spatial aperture to capture subject detail with a narrow spatial aperture to capture scene depth and view-dependent reflectance. Like them, we use a hybrid reconstruction filter, i.e. composed of several filters of different shape. Moreover, both hybrids include a diagonal filter in $uvst$ —equivalent to assuming objects are at a fixed depth. However, the two approaches differ in several ways. Because we consider dynamic scenes, our hybrid includes time, while theirs does not. As a result, we must consider signal-to-noise issues, which do not arise for the static scenes they consider. Secondly and more importantly, Stewart et al. apply both filters to the same light field. We instead sample the light field multiple times, with a different sampling strategy for each filter. Finally, their hybrid filter is linear, whereas ours is nonlinear due to the previously explained compositing step.

7 Discussion and Conclusion

We set out in 1999 to create a system that would allow us to experiment with the imaging capability of a large number of inexpensive cameras. The resulting camera array, while far from perfect, has accomplished this goal. Its key design features—small camera tiles with flexible mounting, accurate timing control of the imagers, and local processing and compression with each imager—have enabled

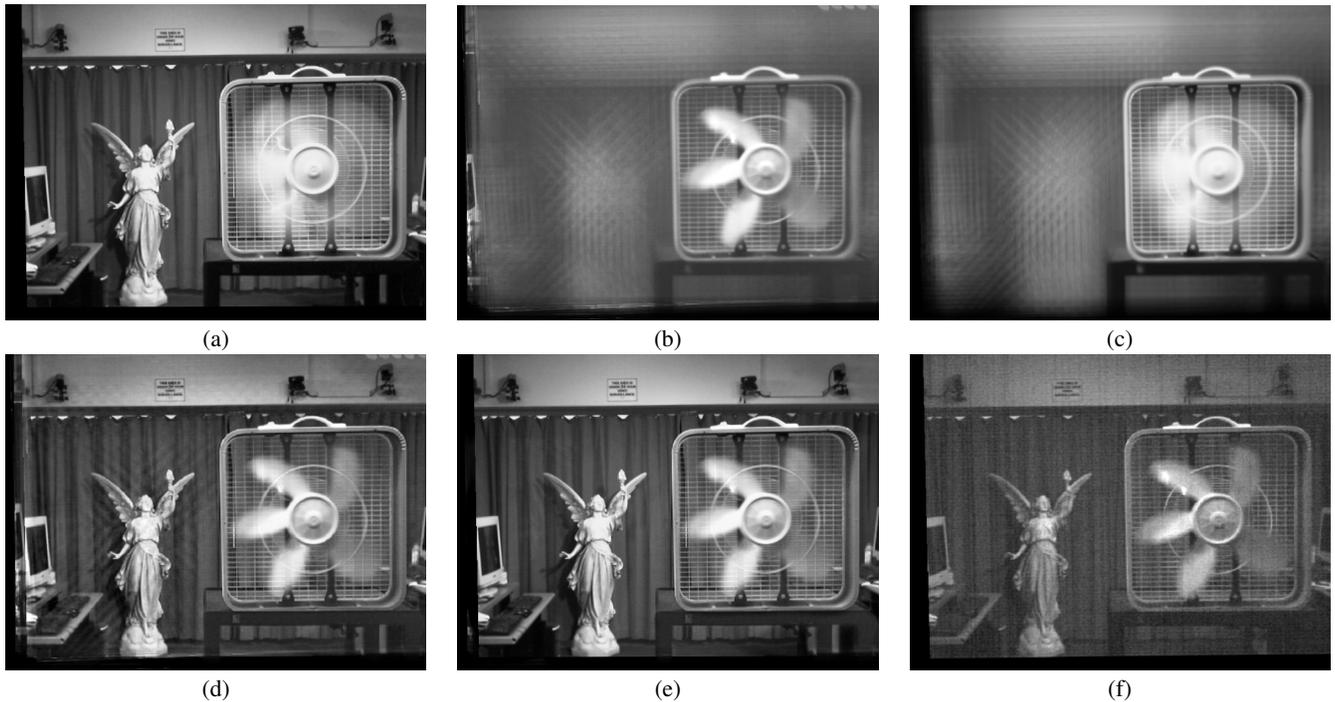


Figure 12: Hybrid synthetic aperture photography for combining high depth of field and low motion blur. (a-c) Images captured of a scene simultaneously through three different apertures: a single camera with a long exposure time (a), a large synthetic aperture with short exposure time (b), and a large synthetic aperture with a long exposure time. Computing $(a+b-c)$ yields image (d), which has aliasing artifacts because the synthetic apertures are sampled sparsely from slightly different locations. Masking pixels not in focus in the synthetic aperture images before computing the difference $(a + b - c)$ removes the aliasing (e). For comparison, image (f) shows the image taken with an aperture that is narrow in both space and time. The entire scene is in focus and the fan motion is frozen, but the image is much noisier.

a wide variety of imaging tasks. The high sampling density can be used to approximate cameras with extraordinary features even with the inexpensive imagers that we used. Abutting the views leads to high-resolution video mosaics, overlapping views can be used to raise the effective dynamic range or frame rate, and cameras can be allocated to accomplish all three simultaneously.

Although many of the techniques we have presented can be applied to high-quality cameras to extend their performance even further, we are particularly interested in exploring the limits of imaging with large arrays of cheap cameras. One open question is whether using many cameras and clever processing we can overcome the poorer imaging characteristics of inexpensive cameras and outperform a single high-quality camera. For example, the resolution of our high-resolution video capture system increases linearly with the number of cameras, but fabrication yields for high-resolution image sensors decrease exponentially with increasing pixel resolution, so the array approach seems superior. On the other hand, our system would make a poor camera for astronomy, which demands very low noise, because noise decreases only logarithmically with the number of cameras. These lines of reasoning indicate that high-quality cameras might be superior in general, but arrays can perform better in some cases.

Aside from increasing imaging performance, our system can also be used to create images that could not have been captured using any normal camera. Some of these applications use cameras spread further apart, creating a wide synthetic aperture. A key issue with this wider baseline is how to allocate the cameras along the two spatial and one temporal dimensions. We show that for scenes with closely spaced cameras or fast motion, triggering all of the cameras at the same time is a poor sampling strategy. Instead, one can sample

the (x, y, t) view volume more uniformly by distributing the firing times of the cameras across the frame time. We take advantage of the resulting reduced image motion with an optical flow variant that explicitly accounts for parallax motion and object motion. This allows us to interpolate missing points in the spatiotemporal volume, creating virtual camera views from new positions in space and time. These techniques can be used to create Matrix-style “bullet time” effects in post-processing.

Based on our experiences with non-linear and hybrid synthetic apertures, we believe the most interesting applications of large camera array are those that do not try to approximate a conventional camera. In particular, we have shown that by shaping the synthetic aperture to avoid rays that do not hit the desired subject, or by creating non-cylindrical shapes in space-time, camera arrays allow one to create images that have not been possible before. We have explored only a fraction of the possible applications, and each one raises questions that suggest new opportunities.

Looking to the future, we would like to design a next-generation camera array. One straightforward improvement to our system would be adding more processing to the cameras. Our FPGAs are operating nearly at capacity doing relatively simple image processing tasks. In a future design, we would also not use image sensors with electronic rolling shutters. The rolling shutter is analogous to a mechanical slit shutter that scans across the image, causing rows at the bottom of the image to expose after rows at the top. This sampling pattern is inconvenient for many applications.

For real-time applications, a future system should support more flexible communication as well as increased processing power. Currently, all of the video from our cameras flows directly to the host

PCs. Live synthetic aperture video, which we demonstrate only for a modest number of cameras, would be easier if each camera could reduce the video it received, adding images from downstream cameras to its own warped input before transmitting it to upstream cameras. We could add this functionality to the current architecture by using multiple IEEE1394 interfaces in each camera to support point-to-point communication between devices, but other applications might have more complex communication needs. Thus, before designing a new architecture, we should investigate potential real-time array applications and how they would map to arrays of “smart” cameras.

Finally, many applications would benefit from incorporating active technologies into this system. For example, we envision using range sensors or projectors for active matting techniques in synthetic aperture photography. Calibration in very unstructured environments might be aided by lasers that could cast geometric calibration targets into our scenes. Projectors are particularly appealing because they seem poised to descend the same slope of falling cost that CMOS sensors are sliding down now. Many of the challenges working with large arrays of projectors are the same as those for cameras: bandwidth, control, and flexibility. As technologies like projectors and range sensors become more affordable and prevalent, we foresee creating large, hybrid camera arrays that not only passively observe, but also actively interact with their environments.

8 Acknowledgements

The authors would like to thank Harry Shum for very helpful suggestions early in the project. We are also grateful to Michal Smulski, Hsiao-Heng Kelin Lee, Monica Goyal, Katherine Chou, Guillaume Poncin, Georg Petschnigg, and Benjamin Jasen Levoy for their contributions to the camera array infrastructure. Construction of the camera array was funded by Intel, Sony, and Interval Research. The application work reported in this paper was supported by the NSF under contract IIS-0219856-001 and DARPA under contract NBCH-1030009.

References

- ANDERSON, D. 1999. *FireWire System Architecture, Second Edition*. Mindshare, Inc.
- BAYER, B., 1976. Color imaging array. U.S. Patent 3,971,065.
- BLACK, M., AND ANANDAN, P. 1993. A framework for the robust estimation of optical flow. In *Proc. ICCV 1993*, 231–236.
- BROWN, M., AND LOWE, D. 2003. Recognizing panoramas. In *Proc. ICCV*, 1218–1225.
- CHAI, J.-X., TONG, X., CHAN, S.-C., AND SHUM, H.-Y. 2000. Plenoptic sampling. In *Proc. SIGGRAPH 2000*, 307–318.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proc. SIGGRAPH 1997*, 369–378.
- GORTLER, S., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. 1996. The lumigraph. In *Proc. SIGGRAPH 1996*, 43–54.
- ISAKSEN, A., MCMILLAN, L., AND GORTLER, S. 2000. Dynamically reparametrized light fields. In *Proc. SIGGRAPH 2000*, 297–306.
- KANG, S., UYTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High dynamic range video. In *Proc. SIGGRAPH 2003*, 319–325.
- LEVOY, M., AND HANRAHAN, P. 1996. Light field rendering. In *Proc. SIGGRAPH 1996*, 31–42.
- MATUSIK, W., BUEHLER, C., RASKAR, R., GORTLER, S., AND MCMILLAN, L. 2000. Image-based visual hulls. In *Proc. SIGGRAPH 2000*, 369–374.
- RANDER, P., NARAYANAN, P., AND KANADE, T. 1997. Virtualized reality: Constructing time-varying virtual worlds from real events. In *Proceedings of IEEE Visualization*, 277–283.
- SCHECHNER, Y., AND NAYAR, S. 2001. Generalized mosaicing. In *Proc. ICCV 2001*, 17–24.
- SHECHTMAN, E., CASPI, Y., AND IRANI, M. 2002. Increasing space-time resolution in video sequences. In *Proc. ECCV 2002*, 753–768.
- S.MANN, AND R.W.PICARD. 1994. Being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures. Tech. Rep. 323, M.I.T. Media Lab Perceptual Computing Section, Boston, Massachusetts. Also appears, IS&T’s 48th annual conference, Cambridge, Massachusetts, May 1995.
- STEWART, J., YU, J., GORTLER, S., AND MCMILLAN, L. 2003. A new reconstruction filter for undersampled light fields. In *Eurographics Symposium on Rendering (EGSR)*, 150–156.
- SZELISKI, R. 1994. Image mosaicing for tele-reality applications. In *WACV 1994*, 44–53.
- TAO, H., SAWHNEY, H., AND KUMAR, R. 2001. A global matching framework for stereo computation. In *Proc. ICCV 2001*, 532–539.
- TAYLOR, D. 1996. Virtual camera movement: The way of the future? *American Cinematographer* 77, 9 (September), 93–100.
- VAISH, V., WILBURN, B., JOSHI, N., AND LEVOY, M. 2004. Using plane + parallax for calibrating dense camera arrays. In *Proc. CVPR 2004*, 2–9.
- VAISH, V., GARG, G., TALVALA, E., ANTUNEZ, E., WILBURN, B., HOROWITZ, M., AND LEVOY, M. 2005. Synthetic aperture focusing using a shear-warp factorization of the viewing transform. In *Proc. A3DISS 2005*.
- WILBURN, B., SMULSKI, M., LEE, H., AND HOROWITZ, M. 2002. The light field video camera. In *Media Processors 2002*, vol. 4674 of *Proc. SPIE*, 29–36.
- WILBURN, B., JOSHI, N., VAISH, V., LEVOY, M., AND HOROWITZ, M. 2004. High speed video using a dense array of cameras. In *Proc. CVPR 2004*, 294–301.
- YANG, J., EVERETT, M., BUEHLER, C., AND MCMILLAN, L. 2002. A real-time distributed light field camera. In *Eurographics Workshop on Rendering*, 1–10.
- ZHANG, C., AND CHEN, T. 2004. A self-reconfigurable camera array. In *Eurographics Symposium on Rendering*, 243–254.
- ZHANG, C., AND CHEN, T. 2004. View-dependent non-uniform sampling for image-based rendering. In *Proc. ICIP 2004*, 2471–2474.
- ZHANG, Y., AND KAMBHAMETTU, C. 2001. On 3d scene flow and structure estimation. In *Proc. CVPR 2001*, 778–785.
- ZITNICK, C., KANG, S., UYTENDAELE, M., WINDER, S., AND SZELISKI, R. 2004. High-quality video view interpolation using a layered representation. In *Proc. SIGGRAPH 2004*, 600–608.