Published in SPIE Proceeding #1732 "Holographics International '92" (SPIE, Bellingham, WA, July 1992) paper #1732-48.

New Approaches To Holographic Video

Mark Lucente
Pierre St. Hilaire
Stephen A. Benton
Derrick L. Arias and John A. Watlington

MIT Media Laboratory, Spatial Imaging Group
20 Ames St., Cambridge, MA 02139 USA

ABSTRACT

Recent advances in both the computation and display of holographic images have enabled several firsts. Interactive display of images is now possible using the bipolar intensity computation method and a fast look-up table approach to fringe pattern generation. Full-color images have been generated by computing and displaying three color component images (red, green, and blue). Using parallelism to scale up the first generation system, images as large as 80 mm in all three dimensions have been displayed. The combination of multi-channel acousto-optic modulators and fast horizontal scanning continue to provide the basis of an effective real-time holographic display system.

INTRODUCTION

The first real-time display of holographic images was reported in 1990 by the MIT Spatial Imaging Group[1,2]. This work tackled the two fundamental difficulties facing interactive holographic displays: to reduce the holographic fringe pattern to a size that can be handled by existing computers, and to modulate light using some high-bandwidth modulation device. In this paper, recent work at MIT answers several questions regarding the real-time display of holographic images. First, computation methods that allow interactive computation are presented. Second, the addition of green and blue components to the previously red image allows for full-color real-time holographic imaging. Finally, work in progress to scale up the display system to a larger image and viewing zone is examined.

Background

Typically, holograms made using coherent visible light contain fringe patterns with spatial frequencies as high as 1000 cycles/mm or more. For typical holograms sizes, the information content of the holographic fringe pattern is over 10^{10} bytes, which is far too large to be handled using conventional computers and electronics. By eliminating the vertical parallax of the image to create a horizontal-parallax-only (HPO) hologram, the information content is reduced by a factor of about 500 or more[3]. Further reduction in the information content of the hologram was achieved by sacrificing the image size and range of viewing angle. The first CGH patterns displayed by the MIT system contained 2 MB of information, and consisted of 64 lines (*holo-lines*) of 32 KB each. The viewing angle range was 15 degrees, and the three-dimensional image occupied a volume of approximately 40 mm on each side. The task of computing these fringe patterns was assigned to a massively parallel supercomputer, with a framebuffer capable of serially reading out the 2-MB pattern line-by-line to produce an analog signal of bandwidth 110 MB/sec. Since no high-bandwidth dynamic spatial light modulator with a horizontal resolution of 32 Kpixels was commercially available, an acousto-optic modulator (AOM) was used to impress the CGH pattern onto a beam of light using what has been called a scophony geometry.

The MIT Approach: The Scophony Geometry

The scophony geometry exploits the fact that only a small part of a hologram needs to be illuminated at any one moment as long as two conditions are met: First, all sub-holograms comprising the overall hologram must be illuminated during the image latency time (integration time) of the human eye. Second, the instantaneously illuminated sub-hologram must be large enough so that the limit set by diffraction allows a resolution that is acceptable for the image. A typical AOM made of tellurium dioxide and operated in the slow shear mode dynamically modulates a beam of light with approximately 2000 collinear samples at a given moment. This pattern scrolls across the aperture of the modulator at the speed of sound (617 m/s). In the scophony approach, a horizontally scanning mirror is used to compensate for the rapid propagation of the modulating acoustic signal within the AOM, and to provide a means of spatial multiplexing. In this way, the image of the AOM is virtually scanned to produce a full 32-KB holo-line. A vertically scanning galvanometric mirror provides the vertical positioning of each holo-line. As shown in Figure 1, the MIT display employs a telescope lens pair to relay and demagnify the light diffracted by the AOM. The first lens ("Fourier transform lens") focuses the light modulated by the AOM to the horizontal and vertical scanning mirrors. The second lens provides a demagnification. Since a given type of AOM supports spatial frequencies in a range determined by its input RF bandwidth and the speed of sound in its acousto-optic material, the demagnification is used to scale the modulating pattern to match the spatial dimensions dictated in computing the CGH pattern. Essentially, the telescope effect acts to amplify angular propagation angles of the light diffracted by the AOM, at the expense of image width. The MIT display behaves much as a conventionally illuminated CGH, but one that has a small aperture rapidly moving across it. The light diffracted by the fringe pattern in each sub-hologram is imaged to the correct horizontal and vertical positions in order to build up the three-dimensional (3D) image brick-by-brick and at a rate faster than that perceivable by the human eye. Finally, typical diffraction-limited spot sizes in the 2000-sample modulation system is about 60 microns, which is smaller than can be perceived by the human visual system.

Figure 1: Diagrammatic view of the 6-MB MIT display.

FAST COMPUTATION: INTERACTIVE DISPLAY OF HOLOGRAMS

The first step in the real-time display of holographic images is the time-consuming computation of the CGH fringe pattern. Initially even a Connection Machine Model 2 supercomputer (CM2) required over twenty seconds to compute a very simple image. Generally, the object to be imaged is modeled as a 3D collection of points scattering light toward the plane of the hologram. (This model is only an approximation to the generally continuous surfaces found in nature.) CGH generation requires a great deal of computation due to the many-to-many nature of physical light propagation. Light of a specified wavelength is propagated from every point on the object ($\sim 10^4$) to every sample ($\sim 10^7$) of the CGH using complex numerical simulation. The elimination of vertical parallax made each holo-line an independent representation of a single horizontal slice of the image[3], though still the multiplicity of calculations remained. The solution was to employ a look-up table approach to compute the CGH.

Bipolar Intensity

The first step toward the development of the look-up table approach was to introduce the simplifying concept of *bipolar intensity*, first reported in Reference 3. The electric field distribution used to represent a propagating wavefront of light is a complex entity, i.e., it consists of both real and imaginary components, or alternatively, both magnitude and phase. When a reference beam and an object wavefront interfere, the intensity (square of the magnitude) of the total field is recorded in the holographic medium. The total field contains three different types of fringe components representing the three possible combinations of object and reference light. The fringes representing interference between each point of object light and the reference beam are necessary and sufficient for the real-time reconstruction of a 3D image. The fringes that occur when light from one object point interferes with light from another object point (*object self-interference*) result in an unwanted noise or image “halo”. Finally, the reference beam alone contributes a spatially nearly invariant intensity that is physically a necessary reality but offers only a useless waste of dynamic range in a CGH. By not including this unwanted reference bias nor the object self-interference in the computational model, a simple algebraic manipulation demonstrates that the CGH can be computed through the use of only the real component of the field distributions[3]. This real component has the curiously non-physical property that it possesses negative as well as positive intensity values, thus the name “bipolar intensity”.

The bipolar intensity method offers several advantages to the computation of any type of CGH[3]. Object self-interference noise is eliminated. Computation speed is approximately doubled. More importantly, it is now possible to linearly sum the fringe contributions representing light from each object point; in other words, fringes are “incoherently” summed in that they do not interfere with each other. Finally, efficient use of the dynamic range of the display system is automatic, without regard to reference beam ratio. (When optically recording a hologram, the reference beam ratio – the ratio between reference and overall object intensities – is the experimental parameter adjusted in order to make use of the dynamic range of the recording medium.) The intensity of the reference beam is no longer relevant to bipolar intensity calculations, in which the final CGH is generally scaled and normalized before transmission to the display system.

Look-up Table Approach

The linear summation made possible by bipolar intensity allows the use of a look-up table approach to rapid computation[3]. When this computing system is initialized, a large array of elemental fringe patterns is precomputed – using the bipolar intensity method – and stored in a look-up table. Each point on the object to be imaged is described by its intended brightness and by its horizontal (x), vertical (y) and depth (z) locations. For a given image point, the appropriate elemental fringe is indexed from the look-up table using x and z, and is then scaled depending on the intended brightness. This scaled fringe is then accumulated into the appropriate holo-line. Therefore, computation has been reduced to only a single multiplication and addition per image point per hologram sample. This approach has also been implemented on several kinds of serial processing computers; in this case, it is more efficient to store the elemental fringes indexed by z-position, and then to perform the necessary shift in x during the accumulation routine. The following is a table comparing the computation times required to compute an image with 10,000 image voxels. Note that the CM2 now performs this task is under one second by using the look-up table approach, as compared with two and a half hours on a Sun 4 workstation performing the unnecessary full complex method. Note also that the speedup achieved with the look-up table approach by a factor of 25 on the CM2, and a factor of 42 on the serial machine, demonstrating that the look-up table approach is especially well suited for use on simpler computation systems.

Computation method	CM2	Serial
Full (complex) intensity	21.80 sec	9434 sec
Bipolar intensity	11.35 sec	4862 sec
Look-up table	0.84 sec	221 sec

In actual operation, additional procedures that are independent of object complexity must be performed, including normalizing the computed holographic pattern and moving it into the framebuffer. Therefore, a generally fixed overhead time (less than 0.1 second on the CM2) will be added to the above times.

Interactive Holographic Shading

Computation speeds obtained using a look-up table allow interactive display and manipulation of holographic images in real time. A 2-MB CGH representing a 10,000-point object can be computed in under one second on a CM2 endowed with 16384 processors working in parallel. The most recently developed software reads in a polygon-list representation of an object, and holographically renders the CGH. The scaling, rotation and positioning of the object are determined by the user's manipulation of dials and buttons, and realistic shaded images are generated interactively in under one second per update. In this interactive holographic shading computation ("holo-shading"), the CM2 front end (a Sun 4 workstation) first numerically transforms each polygon in 3D, as specified by its relative location in the object and the scale and the position of the object as specified by the user. The front end computes the direction in which the polygon faces. This *normal vector* is used to perform a *back-facing cull*, in which a particular polygon is abandoned if it is facing away from the user and is therefore invisible due to the virtual opacity of the object. The direction and intensity of virtual illumination sources are combined with the polygon normal vector to determine the brightness of the polygon. An ambient component is added, and the square root of this total brightness is used as the amplitude of the light scattered from the polygon. The final step in the holo-shading of a polygon is to populate the polygon with points at a resolution that is high enough so that the viewer sees a continuous surface. As the front end generates each point, its 3D location is used by the CM2 to perform a table look-up, and the calculated light amplitude is used to scale the precomputed fringe before accumulation into the overall CGH pattern.

The CM2 is free to perform the actual CGH computations since the polygon manipulation is performed by the front end computer of the CM2. This efficient partitioning of computational tasks allows for interactive holo-shading at rates that are as fast as non-shading computation. Any object can be selected, manipulated, and shaded interactively. An example will help to illuminate the computational feat at hand. The polygonal description of a VW "Bug" automobile body contains a list of 1078 polygons, spanning a total of 1231 vertices in 3D. After a typical scaling and rotation, about half of the polygons are back-face culled, leaving 510 polygons to populate with points. Typically these 3D image points or *voxels* number 5473. The image is 64 lines, typically discretized into 250 horizontal locations and 50 depth locations (for a total of 800,000 possible image voxels). In this example, the 5473 voxels are generated and used to compute a 2-megabyte hologram in 0.46 seconds. This represents a speed of 84 microseconds per voxel, the same as reported with the non-shading approaches to real-time hologram computation.

Smooth shading has not yet been implemented. The polygons are filled with a constant point brightness, giving shaded surfaces a slightly blocky appearance. Instead of a single amplitude calculation for the entire polygon, the addition of smooth ("Gouraud" or "Phong") shading can be facilitated by interpolating brightness values across the polygon. Virtual specular reflections and other directionally dependent elements of realism are as yet absent due to their computational time requirements.

Synthetic Phase-Front Modulation

Currently, work involving synthetically generated elemental fringes produce shaded images with improved realism. Essentially, the object light is no longer modeled as sections of spherical wavefronts. In fact, a significant amount of object or reference beam phase-front variation is allowed when computing the elemental fringes. Recalling that the scophony approach is tantamount to the rapid scanning of a 2 mm aperture in front of a fully illuminated CGH pattern, the maximum phase coherence required in the CGH is 2 mm. Furthermore, strictly phase-continuous cosinusoidal fringes are not necessary since the CGH can be thought of as being illuminated by incoherent nearly monochromatic light. (The display can employ narrow-band incoherent light, though laser light is used primarily for the convenience of a bright narrow-bandwidth source[4]. For the current system, a linewidth of 1 nm blurs each image point by an imperceivable 40 microns.) The phase-front modulated fringes contain a stocastically varying phase component and thus do not diffract light to a diffraction-limited spot. They have been designed, for example, to diffract light to a small volume matched to the 3D resolution (width, height and depth) required in a particular image. The presence of random phase shifts in the object wave at regular intervals suffice to enlarge the spot size. These enlarged image elements fill in a larger portion of the image without requiring additional computation time. For polygon-shading, this means a move away from the simple point-based computation methods and toward larger more varied elemental image representations[5]. Current work involves the investigation of how to use other synthetic phase-front modulation schemes to enhance the realism of the final image. Also, the flexibility gained by allowing periodic phase shifts in the fringe pattern can also be exploited in order to compress CGH information for the purpose of rapid communication, e.g., from the CM2 to a framebuffer system. This flexibility also leads to some computational simplifications that are now being investigated. In particular, the generation of a stereogram-type CGH[3] is greatly improved by utilizing synthetic elemental fringes. Finally, this new approach may help to circumvent the limitations posed by the quantization noise and limited dynamic range of the display system electronics.

Dynamic Range and Quantization

Many questions remain regarding the rapid computation of holograms for real-time display. As images become more complex, i.e., consist of more and more points, the total brightness of the image decreases. In the MIT display system, the limited brightness becomes evident in complex shaded images. Given the range of output signal, the framebuffer is capable of producing some minimum and maximum signal, and is quantized into $2^8 = 256$ evenly spaced levels. A fringe pattern that spans all 256 levels (8 bits) of the framebuffer RAM memory produces a signal that diffracts the maximum amount of light into the image. Since the image is comprised of many superimposed cosinusoidal fringes of different (smaller) amplitudes, the total amount of light is essentially divided among the points comprising the image.

Consider the following single holo-line example: A fringe that ranges from 0 to 1 (one bit) after being quantized by the framebuffer diffracts light with brightness of 1 (arbitrary units). Adding the overlapping fringe contributions from 128 such image points fills the entire extent of the framebuffer, and gives a total image brightness of 128. However, if the image is a single point represented by a fringe ranging from 0 to 255, it has a brightness of $255^2 = 65025$. This example illustrates that a simple one-point image is over 500 times brighter than a complex image populated with 128 points. It is important to note that this problem is a reality of any signal processing system with a limited absolute range, particularly synthetic holography. It is here manifest through the first facile generation of complex fully-shaded images. In practice, noise from the electronics and the display system further limit the dynamic range available to the image.

FULL-COLOR HOLOGRAPHIC DISPLAY

The first successes in real-time display holography naturally produced monochromatic images. The MIT system employed the 632.8-nm red wavelength of a HeNe laser both for convenience and because it is at the longer end of the visible spectrum and therefore requires proportionally lower spatial frequencies. In 1991, the MIT system was transformed into a full-color imaging system through the addition of green and blue components[6]. The realization of full color in the MIT display was straight-forward, though the dependence of the acousto-optic effect on wavelength led to two new concerns: alignment of the three color components, and the unequal viewing zone sizes.

Computationally, three primary color CGHs were generated from the object description using a color computer-graphics lighting models. Each of these 2-MB CGHs were stored in one of the three framebuffer channels, and read out in parallel as three analog RF signals. A three-channel AOM was substituted for the single-channel AOM, with one of the RF signals feeding each. The red laser light was aligned to pass through one modulation window in the AOM. A green beam (532 nm from a diode-pumped frequency-doubled Nd:YAG laser) and a blue beam (442 nm from a HeCd metal vapor laser) were aligned to pass through their respective modulating signals. The holographic pattern diffracted the beams horizontally into three different ranges of angles. Fortunately, careful analysis of the Bragg effect in this particular device showed that the diffraction angles were approximately linearly related to wavelength and were therefore compensated for using a single grating. This grating had a spatial frequency equal to that of the modulating signal at the center (75 MHz) of its RF bandwidth (50 MHz to 100 MHz). This approximation was fine-tuned by adjusting the reference beam angles used to compute the three CGH color components. The three modulated beams passed through the same scanning system used in the monochromatic display.

The color display performs very well. The images are crisp and exhibit good contrast. Color registration is excellent despite the linear approximation made to horizontally align the three beams. The remaining problem is that the shorter wavelengths of green and blue light will not span the entire range of angles spanned by the red light. (Given the highest and lowest spatial frequencies supported by the AOM, a beam of light diffracts through a range of angles that is approximately proportional to wavelength.) The effect is that at one extreme of the viewing zone, the viewer loses the blue and then the green components of the image before the red is finally vignetted. Since the color fringing happens only in one eye (the other remains closer to the center of the viewing zone), the image is not severely degraded; though the image may appear a bit odd at one extreme, the human visual system seems to perceive an image as possessing some color despite the lack of color components in one eye.

SCALING UP THE DISPLAY

The first increase in the size of the holographic image was the result of the use of parallelism[7]. Similar in approach to the addition of color, each of the three buffers of the framebuffer card contained every third line of a 6-MB CGH. The framebuffer read out the three parallel signals to a three-channel AOM. The image size was increased by a factor of three; in this case, the vertical resolution was tripled. Though the overall bandwidth of the display was increased by a factor of three, each of the signals still possessed the 110-MB/sec bandwidth of the original system. Therefore, the same RF processing electronics were utilized, though multiplied in threes. The optical system (Figure 1) was similar, with the substitution of the three-channel modulator for the single-channel AOM. In addition, a HOE beam aligner was placed after the AOM in order to move the 3 beams closer. (They emerged from the AOM separated by several millimeters.) Thus, a scale-up by a factor of three was achieved at less than three times the expense.

Using this first introduction of parallelism as an example, it is clear that the MIT scophony geometry can be scaled to a larger higher-bandwidth system by making use of parallelism. Multi-channel AOMs can be custom made to have 32 channels or more, limited by the delicate process of assembling the ultrasonic transducers onto the fragile acousto-optic crystal. The next step in the development of the MIT holographic video system was to employ an 18-channel AOM at the heart of a second generation system capable of displaying a 36-MB CGH. The new display system is designed to have 18 parallel channels of framebuffer intermediate storage and 18 parallel RF analog output signals and subsequent RF processing.

The first problem in this 36-MB system is to compute and rapidly read out the 18 channels in parallel. The 3-channel system was supported by a single high-end graphics framebuffer (which is designed nominally to store the red, green and blue components of a computer graphic image). Perhaps six framebuffer cards can be used in parallel, but it may be difficult or impossible to synchronize these 18 output signals. Another problem – the presence of blanking intervals – remained unsolved in the first system. A standard framebuffer is designed to read out lines of video, each with a length of up to 2 KB, with a blanking interval inserted after each line (to allow time for the horizontal retrace on a standard CRT display). Though this interval may be shortened by reprogramming the framebuffer microcontroller code, it generally cannot be reduced to zero. Therefore, each holo-line (of length 32 KB in 2-MB system) was broken in fifteen places by the blanking intervals that appeared as a dark "picket fence" across the image at the plane of the hologram (image of the AOM). This severely limited the usable depth of the imaging system, and required that images in the MIT system be at least 60 mm from the plane of the reconstructed hologram. Hence an additional requirement on any new framebuffer system is that it be capable of producing a continuous stream of data that is long enough to represent an entire holo-line. Since in the 18-channel system holo-line-length is 256 KB, the decision was to use six custom made framebuffers capable of storing a 36-MB CGH, and reading out 18 holo-lines at a time in continuous 256-KB lengths. The framebuffer system is called "Cheops" and was designed at the MIT Media Laboratory for research into scalable digital television, in which large reconfigurable framebuffers are used to perform real-time compression and decompression of video signals[8].

Limitations of the Horizontal Scanner

The second obstacle in expanding the system involves the horizontal scanning mirror. The limitations of using a rotating polygonal mirror to perform the horizontal scanning became evident upon examination of the way in which design parameters scale with the width of the image and with the size of the viewing zone. The horizontal scanning mirror must satisfy several engineering parameters in the MIT scophony geometry. First, it must scan at such a rate as to compensate for the propagation of the ultrasonic modulating pattern within the AOM. Second, a complete scan must occur during the time τ required to feed a single holo-line from the framebuffer into the AOM. An 18-sided polygonal mirror rotating at a rate ω performed this function in the first generation (2-MB and 6-MB) MIT systems. To satisfy this second constraint,

$$\omega = \frac{2\pi}{N\tau} \tag{1}$$

where N is the number of facets. (For the purposes of this discussion, the simplest case – the confocal scanning case – is considered, though the general case exhibits the same limitations[4].) In the confocal case, the first lens of the system focuses to the scanner at focal length F_1; the Fourier transform of the AOM modulating pattern is imaged to the plane of the horizontal scanning mirror. For a signal with bandwidth $\Delta\nu$, the width Δx of the Fourier transform is

$$\Delta x = \frac{\lambda F_1 \Delta\nu}{s} \tag{2}$$

where λ is the wavelength, and s is the acoustic speed within the AOM material. To relate Δx to the width of the hologram w and the viewing zone angle θ, note that the minimum number of samples required to faithfully represent a line of this hologram is $2w \sin\theta/\lambda$ by the Nyquist sampling theorem. Given τ, the holo-line-time,

$$\Delta\nu = \frac{w \sin\theta}{\lambda\tau}. \tag{3}$$

Therefore, the minimum required scanner width is:

$$\Delta x = \frac{F_1 w \sin\theta}{\tau s}. \tag{4}$$

Thus, Δx scales linearly with hologram width and viewing zone size. When using a polygonal scanning mirror, the facet width and number of sides N determine the radius R. Combining Equation 1 with these geometric concerns[9],

$$R \geq \frac{2w\sin\theta}{\omega^2\tau^2}. \tag{5}$$

Choosing optimal parameters is a matter of engineering choices. A polygon with a larger number of facets can rotate more slowly (smaller ω) and allows the use of higher f/# optics but must be much bigger, heavier, and more expensive. A smaller number of facets necessitates an output lens with an impractical f/#. Inevitably, Equation 5 demonstrates that increasing the size of the image or viewing zone requires an equivalent increase in the size of the polygon. Numerical analysis demonstrated that using a polygon as the horizontal scanning element becomes impractical for images of more than a few centimeters in width[4]. As an example, an image with a horizontal size of 200 mm and a viewing zone of 20 degrees requires a polygon with a radius of 500 mm, assuming a 12-sided mirror.

Fast Galvanometric Horizontal Scanning

Clearly another solution is needed if we want larger images. One approach is to use a galvanometric scanner as our horizontal scanning element. The width of the horizontal scanning mirror needs to be at least Δx. Given the confocal MIT scophony geometry, this width is simply

$$\Delta x = \frac{w\sin\theta}{2\Omega} \tag{6}$$

where Ω is the mechanical excursion range of the galvanometric scanner[9]. The f/# of the output lens scales linearly with the scan angle θ, so in practice Ω is generally kept below 15 degrees (equivalent to an optical deflection range of 30 degrees.) Since we require a constant angular velocity over the entire deflection range, the galvanometric scanner is fed with a linear (triangle or sawtooth) waveform. The scanning frequency is generally limited by the inertia of the scanner assembly, with typical values ranging between 150 and 500 Hz.

At the relatively low rates achievable using a galvanometric horizontal scanner, the main limitation of this approach is its inability to complete a large number of horizontal scans within a given frame time (at least $1/30$ Hz to avoid flicker). As an example of this limitation, let us consider the specifications of the MIT galvo-based display. It currently has a horizontal viewing zone of $\theta = 26$ degrees and the hologram width is $w = 70$ mm. With a mechanical deflection range of $\Omega = 15$ degrees, the mirror must be at least 59 mm. Moving a mirror of this size is difficult for frequencies over 100 Hz since linear motion is required. By using a lightweight beryllium mirror and a high performance moving coil galvanometric scanner (Cambridge Technology Inc. model 6650) we were able to obtain an acceptably linear scan up to 150 Hz. Because the AOM-based display has no "memory" or latency, the display refresh rate must be high enough for the human visual system to see an image without flicker. For commercial devices, this refresh rate should be at least 60 Hz, but a 30 Hz refresh rate still results in acceptable images for our purposes. At this rate, each acoustic channel in the AOM completes only 150 Hz$/$30 Hz $=$ 5 horizontal display lines per frame. The solution to the design of a galvo-based display is to simultaneously modulate a large number of acoustic channels in parallel within a single crystal of acousto-optic material, a technology that is now well-documented and commercially available. It is also possible to double the vertical resolution by driving the scanner with a triangular waveform and by using two AOMs with their acoustic waves traveling in opposite directions. With this back-and-forth scanning scheme, only one AOM is driven at a given instant so that the linear velocity of its acoustic wave is tracked by the angular motion of the scanning mirror.

The current MIT system consists of two AOMs having 18 acoustic channels each and has a 30 Hz refresh rate. Since the last of the five lines is blanked during the vertical retrace interval, the display exhibits a vertical resolution of $(5-1) \times 2 \times 18 = 144$ lines. Each of the AOMs transducers is supplied a signal from a channel of the Cheops framebuffer system. These 18 parallel signals are frequency-shifted by mixing them with a local oscillator; they are amplified and then demultiplexed to feed the two banks of transducers. This demultiplexing operation is more practical than using twice the framebuffer channels since at any time only one AOM needs to be active.

The CGH pattern being rapidly read out of the Cheops system is currently computed on the Cheops processor card using table look-up. This is done for convenience, since a fast I/O bus has not yet been installed in order to rapidly communicate CGH information from the CM2 into the 36 MB framebuffer memory. As an example of the versatility of the look-up table approach, the CGH is computed using only bit-shifting and integer addition and multiplication. Without floating-point arithmetic or even a cosine function, a CGH is computed without noticeable noise or other limiting visual effects. The image is bright and crisp, and computed in a matter of minutes on the single Intel 960 processor of the initial Cheops configuration. The signal-to-noise ratio is comparable to the best optically recorded silver halide hard-copy holograms. Images have been generated with a depth of 500 mm, although in practice the astigmatism due to the lack of vertical parallax limits the useful depth to 200 mm (as is the case for any rainbow or HPO hologram of comparable scale). Since Cheops generates a continuous stream of 256 KB of holo-line with no blanking intervals, the image can be placed at practically any depth, though the optical set-up favors the 500 mm immediately after the output lens. At the time of publication, images fill a volume of 80 mm (horizontal) x 70 mm (vertical) x 200 mm (depth) with a horizontal viewing zone angle of 24 degrees and a vertical resolution of 96 lines. As the last 6 framebuffer channels come on line, the vertical resolution will be increased to 144 lines. A high-speed I/O bus will provide a path for the 36-MB CGH to be loaded into the framebuffers. A compression scheme may be employed in order to take advantage of the processing power of the Cheops system.

CONCLUSION

Research into real-time display holography has produced a firm foundation for further investigation. For three years, the MIT system has proven the viability of real-time display of holographic images despite the obstacles of enormous information content and optical bandwidth. As described in this paper, many additional possibilities have been explored. Interactive computation, full-color images, and larger display systems are all possible. The use of parallelism has been at the heart of these advances, both in computation and in electronic and optical systems. Also at the heart of many advances has been the concept of a better-than-physics approach to holography. Optical holography can display 3D images that contain far more information than can be used by the human visual system, as evidenced by the extremely high image resolution and the continuous parallax generally exhibited by holograms. The reduction of image resolution, the elimination of vertical parallax, and the discretization of horizontal parallax (as in the case of stereograms) have made possible the rapid computation and display of synthetic holograms. For example, the introduction of bipolar intensity or "incoherent" fringe summation is an example of how a physically unrealizable process can be simulated numerically with benefits to computing and to image quality. Many such advances remain to be explored. Current research into optimized synthetic elemental fringes will likely lead to CGH data compression and faster computation.

With the initial questions now answered, a new set of questions are born. For example, is the MIT scophony approach the best way to display a CGH in real time? It is not obvious how such an approach can be used to create images possessing full parallax. And can the MIT display be supported by less exotic computer systems? Current work has demonstrated that a first generation system (6-MB) can be supported by a high-end computer-graphics workstation, with performance that is as good as the performance demonstrated by the CM2. As computers become cheaper and more powerful, holographic video research can move forward in greater leaps.

The scophony system employs a modulator that is purely dynamic; a single sample of the CGH propagates to the end of the modulation aperture in only 19 microseconds. Because of this need for a continuous high-bandwidth signal, the scophony approach places a great deal of the burden on the computational and electronic systems. Current research into the possibility of alternative modulation schemes may make possible a display that exploits the latency of a modulator in order to relax the requirements on the electronics.

Finally, the variety of real-time holographic displays has enabled some preliminary investigations into the psychophysics involved in the viewing and perception of holographic images. Because the MIT system can display any arbitrary image, various tests of depth cues, optical illusions, and image realism have been performed interactively. Certain nonsensical images can be generated in order to study the response of the human visual system to unnatural, non-physical images. For example, a scene in which objects in the *background* of a scene occlude the objects in the *foreground* of the scene tests the importance of various depth cues. Real-time display holography is likely to leap forward as insight is gained in the science of 3D imaging.

ACKNOWLEDGMENTS

The authors wish to acknowledge the many researchers in the MIT Spatial Imaging Group past and present, especially John Underkoffler, Wendy Plesniak, Dr. Paul M. Hubel, Michael Halle, and Prof. Hiroshi Yoshikawa (Nihon University, Japan).

This research has been supported in part by: the Defense Advanced Research Projects Agency (DARPA) through the Rome Air Development Center (RADC) of the Air Force System Command and the Naval Ordnance station, Indian Head, Maryland, contract #N00174-91-C-0117; by International Business Machines, Honda R&D Co., Ltd., and NEC Corp.; and by the TV of Tomorrow consortium at the MIT Media Laboratory. "Connection Machine" is a trademark of the supercomputer's manufacturer, Thinking Machines, Inc., Cambridge, MA, USA.

REFERENCES

1. Pierre St.-Hilaire, Stephen A. Benton, Mark Lucente, Mary Lou Jepsen, Joel Kollin, Hiroshi Yoshikawa and John Underkoffler, "Electronic display system for computational holography", In S. A. Benton, editor, *Practical Holography IV*, vol. 1212-20, SPIE (Bellingham, WA), 1990.
2. Stephen A. Benton. "Experiments in holographic video imaging", In P. Greguss, editor, *Proc. of the SPIE Institute on Holography*, vol. IS#08, SPIE (Bellingham, WA), 1991.
3. Mark Lucente. "Optimization of hologram computation for real-time display", In S. A. Benton, editor, *Practical Holography VI*, vol. 1667-04, SPIE (Bellingham, WA), 1992. See also: Mark Lucente. "Interactive computation of holograms using a look-up table", *Journal of Electronic Imaging* (submitted for review).
4. Pierre St.-Hilaire, "Real-time holographic display: improvements using higher bandwidth electronics and a novel optical configuration", Master's thesis, Massachusetts Institute of Technology, June 1990.
5. John S. Underkoffler, "Toward accurate computation of optically reconstructed holograms", Master's thesis, Massachusetts Institute of Technology, June 1991.
6. Pierre St.-Hilaire, Stephen A. Benton, Mark Lucente and Paul Hubel, "Color images with the MIT holographic video display", In S. A. Benton, editor, *Practical Holography VI*, vol. 1667-33, SPIE (Bellingham, WA), 1992.
7. Pierre St.-Hilaire, Stephen A. Benton, Mark Lucente, Hiroshi Yoshikawa and John Underkoffler, "Real-time holographic display: Improvements using a multichannel acousto-optic modulator and holographic optical elements", In S. A. Benton, editor, *Practical Holography V*, vol. 1461-37, SPIE (Bellingham, WA), 1991.
8. V. Michael Bove, Jr. and John A. Watlington, "Cheops: a modular processor for scalable video coding", In *Proceedings of SPIE Visual Communications and Image Processing '91 Conference*, SPIE, Nov. 1991.
9. Pierre St.-Hilaire, Stephen A. Benton and Mark Lucente, "Synthetic aperture holography: a novel approach to three dimensional displays", *Journal of the Optical Society of America A* (to be published).