Computer-Generated Holography as a Generic Display Technology

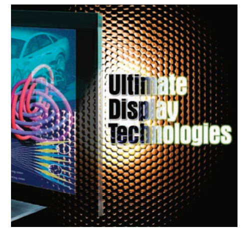

Computer-generated holography is a powerful technology suitable for a wide range of display types, including 2D, stereoscopic, autostereoscopic, volumetric, and true 3D imaging. Although CGH-based display systems are currently too expensive for many applications, they will become a viable alternative in the near future.

Chris Slinger
Colin Cameron
Maurice Stanley
QinetiQ

Invented in 1947 by Dennis Gabor, holography—from the Greek *holos,* for whole—is a 3D display technique that involves using interference and diffraction to record and reconstruct optical wavefronts. Holography's unique ability to generate accurately both the amplitude and phase of light waves enables applications beyond those limited by the light manipulation capabilities of lens- or mirror-based systems.

Computer-generated holography is an emerging technology, made possible by increasingly powerful computers, that avoids the interferometric recording step in conventional hologram formation. Instead, as Figure 1 shows, a computer calculates a holographic fringe pattern that it then uses to set the optical properties of a spatial light modulator, such as a liquid crystal microdisplay. The SLM then diffracts the readout light wave, in a manner similar to a standard hologram, to yield the desired optical wavefront.

Compared to conventional holographic approaches, CGH

- does not rely on the availability of specialized holographic recording materials;
- can synthesize optical wavefronts without having to record a physical manifestation of them—for example, it can generate 3D images of nonexistent objects; and
- offers unprecedented wavefront control by making it easy to store, manipulate, transmit, and replicate holographic data.

Although CGH-based display systems can be built today, their high cost makes them impractical for many applications. However, as compute power and optical hardware costs decrease, CGH displays will become a viable alternative in the near future.

ADVANTAGES

CGH provides flexible control of light, making it suitable for a wide range of display types, including 2D, stereoscopic, autostereoscopic, volumetric, and true 3D imaging. CGH-based display technology can produce systems with unique characteristics impossible to achieve with conventional approaches.

Optical efficiency

Using a pure-phase analog SLM, it is possible to diffract nearly all light falling onto the SLM into the desired image. Conventional displays are often highly inefficient because they use some form of absorption or attenuation to vary pixels' gray levels. In addition, all the light falling onto a computer-generated hologram can be diffracted to any pixel of the image, resulting in very high dynamic ranges.

Published by the IEEE Computer Society

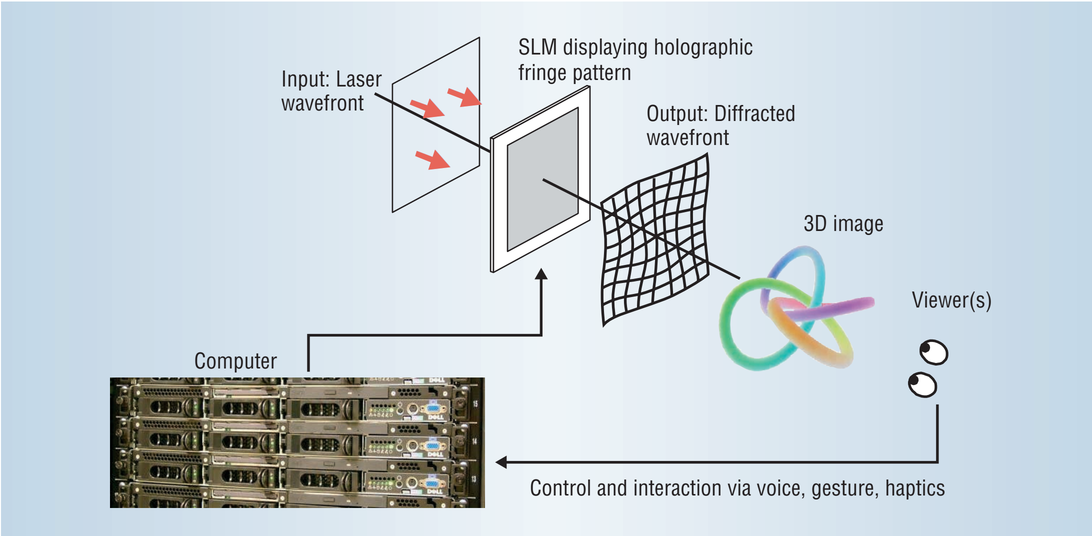

Figure 1. Computer-generated holography. A computer calculates a holographic fringe pattern for display by the spatial light modulator (SLM), which diffracts laser light to yield an interactive, true 3D image.

Dynamic range in conventional displays is limited as each of the display's pixels is unambiguously mapped to a corresponding pixel in the image.

Ease of tiling

CGH pixels need not be contiguous. For example, high-pixel-count CGH-based displays can be made by assembling multiple microdisplay SLMs into a 2D array. Gaps between the SLMs, necessitated by interface and power connections, are not visible in the image projected from such computer-generated holograms. In conventional displays, pixel-level continuity and accuracy are required to prevent the eye from seeing tiling artifacts.

Tolerance of pixel defects

CGH delocalizes information associated with any point in a 2D or 3D image and spreads it over the hologram surface. Consequently, "dead" pixels or even multiple rows or columns of defective pixels do not cause noticeable degradation in the image. In contrast, even small numbers of dead pixels are unacceptable in conventional displays, a major factor in production costs.

Wide color gamut

CGH generally requires replay using lasers, which, though expensive now, are becoming ubiquitous and decreasing in price. A major advantage of lasers is that their monochromaticity, or spectral purity, yields a very wide color gamut, approaching 90 percent of the color range the human eye can perceive. Most existing display systems deliver only about 40 percent of the color gamut or less, as they use relatively broadband color filters.

Full depth cues

Able to generate a wide range of optical wavefronts, CGH is the only technology capable of synthesizing a true 3D image—that is, one with all the human visual system's depth cues. Stereoscopic, autostereoscopic, and volumetric displays are deficient in some depth cues, or worse, produce conflicting cues. This can result in the viewer experiencing discomfort, nausea, and other negative effects, especially after long-term use.

Conflicting depth cues also inhibit performance of various tasks in non-CGH-based displays. In a stereoscopic (glasses-based) display system, convergence and accommodation (focusing) depth cues conflict.[1] Viewer-position-independent obscuration is impossible in volumetric displays. Even advanced multiview autostereoscopic systems, such as integral imaging systems, generate optical wavefront approximations and resolutions significantly below the eye's visual acuity capabilities.

High system volume and image resolution

The lower resolutions of most display technologies, which use geometric optics for ray-based approximation, result from their need to minimize diffraction effects. In contrast, holographic systems directly exploit diffractive effects and therefore can be more compact as well as deliver images with higher information content. Appropriate CGH algorithms can adjust image resolution dynamically, allowing users to tailor system computing resources for specific activities.

Artifact-free binary modulation

Displays, particularly microdisplays, are often binary in nature. Examples include Texas Instruments' Digital Micromirror Device, various ferroelectric liquid crystal microdisplays, and Silicon Light Machines' Grating Light Valve technology. To generate images perceived as having good gray scale, these devices run at high frame rates and use temporal multiplexing. In some applications, this can result in image degradation and artifacts. CGH, however, can produce gray-scale images from binary fringe patterns, thus no multiplexing is necessary,

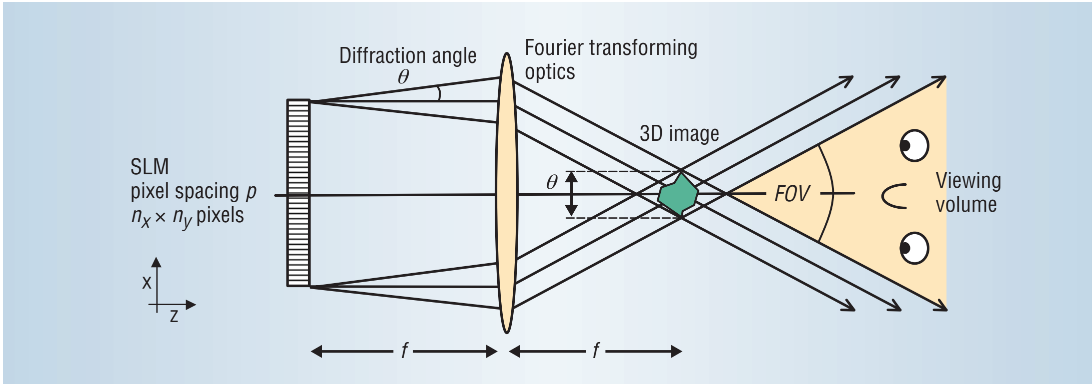

Figure 2. Relationship between image size (I) and field of view (FOV) in true 3D floating-image generation. I × FOV is proportional to the number of pixels across the SLM, and is independent of the focal length (f) of any optics. The same relationship holds for other replay geometries.

hardware frame rates can be reduced, and temporal dither artifacts are absent.

Aberration, distortion, and conformal correction

It is sometimes possible to precompensate the CGH pattern to prevent known optical aberrations or imperfections from degrading CGH-generated images. In the case of projecting a 2D image onto an arbitrary (but known or determinable) 3D surface, a CGH-based projector system can likewise produce a compensated, sharply focused image.

CHALLENGES

Despite the unique capabilities of CGH-based display technology, it remains prohibitively expensive for most commercial applications, largely because of the enormous pixel counts required to achieve sufficient image resolution and image size (I). In the case of direct-view CGH-based displays, generating images with an adequate field of view (FOV) also adversely impacts pixel-count requirements.

Image width and field of view

For direct-view CGH-based displays, the product of I and FOV is the primary influence on CGH pixel counts.

To understand why, consider Figure 2, which shows a typical Fourier transform configuration used to generate usable $I \times FOV$ irrespective of CGH pixel spacing. The angle θ over which light can be diffracted from a hologram fringe pattern is governed by the grating equation

$$\Lambda \sin\theta = \lambda,$$

where λ is the replay light wavelength and Λ is the spatial frequency of the fringe diffracting the light. The largest value of θ is given when $\Lambda = 2p$, where p is the pixel spacing on the SLM displaying the CGH pattern. Regardless of the optical elements' apertures and powers, $I \times FOV$ is proportional to the number of pixels along the appropriate CGH axis. Even for a relatively undemanding workstation application with, say, $I = 0.5$ m and $FOV = 60^\circ$, a full-parallax computer-generated hologram requires approximately 10^{12} pixels.

Using techniques such as restricting the hologram to horizontal parallax only—likely to be acceptable in many, but not all, applications—can reduce system pixel-count requirements to about 10^{10}. However, this is still three orders of magnitude higher than the pixel counts of other high-end display systems. The inability to calculate hologram fringe patterns or update an SLM system with this many pixels, at interactive rates, remain major bottlenecks to widespread CGH adoption.

Image resolution

For both 2D and 3D displays, the number of resolvable image points N_I in any plane of the image perpendicular to the optical axis is directly related to the computer-generated hologram's pixel count N_{CGH}. In fact, $N_{CGH} \geq N_I$, with typically $N_{CGH} = 2N_I$. For many display applications, meeting this constraint is less onerous than a large $I \times FOV$. However, for 2D projection systems, in which FOV is not a factor due to the scattering of light over a wide range of angles by the projection screen surface, image resolution is the primary influence on system complexity.

Image quality

The image from a computer-generated hologram cannot contain more information than the hologram itself. Ingemar Cox and Colin Sheppard[2] have defined and quantified the channel capacity C of an optical system. In the case of a CGH-based image generation system,

$$C \propto n_x \log[Q]\, n_y \log[Q] \log(1+S/N),$$

where Q is the number of quantization levels, $n_{x,y}$ are the resolvable pixels across the hologram and/or image, and S/N is the signal-to-noise ratio. Thus, binary or other highly quantized hologram fringe patterns, while maintaining sufficient $I \times FOV$, can compromise image quality by decreasing the amount of information passing through the system. System designers must therefore pay careful attention to image-resolution, size, field-of-view, and signal-to-noise tradeoffs.

Laser speckle

Unlike conventional volume holograms, computer-generated holograms must be illuminated (replayed) with narrowband and spatially coherent light—generally from red, green, and blue lasers—to produce color images. In addition to the high costs associated with this still-maturing technology, particularly green lasers, speckle can adversely affect image quality. Designers can overcome this problem through various speckle-reduction techniques, such as controlling laser system coherence.[3] In addition, CGH can generate an image of a surface that appears diffuse without having to rapidly vary the optical wavefront's phase structure, which will reduce perceived speckle.

ALGORITHMS FOR CGH DESIGN

Due to the daunting pixel counts that CGH-based display systems require, special attention must be paid to the efficiency of the algorithms used to calculate the holographic fringe patterns.[4] These fringe patterns, represented as a 2D array of pixel values, are usually analog in form. Some CGH modulators, such as QinetiQ's Active Tiling system, require converting the analog pixel array distribution into a binary distribution, with as little degradation in the replayed 3D image as possible. Minimizing any additional computational load of such a binarization step is particularly important for highly interactive systems.[5]

For all algorithms, the information for 3D image generation is typically exported from a commercial 3D computer-aided design tool as a triangulated mesh with material properties. The material properties can include diffuse, specular, and shininess components.

Object surface intensities are based on user-defined material and light properties. The algorithms populate the mesh with points at a density high enough that a human observer cannot see them. The display system then holographically reconstructs these points. For simplicity we focus on full-parallax, true 3D image generation, although all approaches can be modified to produce horizontal-parallax-only holograms, with significant savings in computational resources.

Ping-pong algorithm

One of the best-known and simplest Fourier-based CGH algorithms is the ping-pong technique devised by Yoshiki Ichioka, Masaharu Izumi, and Tatsuro Suzuki.[6] Building upon the crude first Born approximation, or superposition, approach, this technique introduces an obscuration operator that enables the reconstructed 3D images to exhibit hidden-line-removal effects.

The ping-pong algorithm first slices the object into planes perpendicular to the design plane. It then propagates light through the image's first plane (plane 1), toward the design plane, until it meets the next plane (plane 2). Next, the algorithm applies an occlusion operator that allows the selective attenuation and modulation of light as it passes through plane 2. It then adds light from plane 2, and propagation proceeds to plane 3. The process repeats through the other planes until the light reaches the design plane.

The ping-pong algorithm is easy to code using standard numerical subroutines. However, the basic approach can only generate images of self-luminous objects, so implementing full rendering effects and lighting models requires enhancements. It is also computationally inefficient. For these reasons, the ping-pong approach is not viable for a practical interactive system.

Special attention must be paid to the efficiency of the algorithms used to calculate the holographic fringe patterns.

Interference-based algorithms

These algorithms closely simulate light propagation in a conventional interferometric hologram recording. They essentially implement a 3D scalar diffraction integral capable of generating very high quality images including lighting effects and surface-reflection properties. Interference-based algorithms can reproduce all human visual depth cues and currently serve as the benchmark for image quality.

QinetiQ has developed this approach in its coherent-ray-trace algorithm. This algorithm uses Fourier transform geometry with an off-axis object to incorporate advanced rendering effects such as shadowing and Phong shading of curved surfaces. The core calculation consists of a double summation of all visible object points for each CGH pixel. This many-to-many mapping between CGH pixels and object points causes high computational loads.

Diffraction-specific algorithms

Mark Lucente pioneered diffraction-specific CGH computation while at MIT,[7] with more recent variants of the algorithm featuring Fourier-based geometries.[5] Importantly, these algorithms allow tradeoffs in image quality with computational speed. Diffraction-specific algorithms consider only the replay of a computer-generated hologram rather than its interferometric formation as in coherent ray tracing. By controlling the final holo-

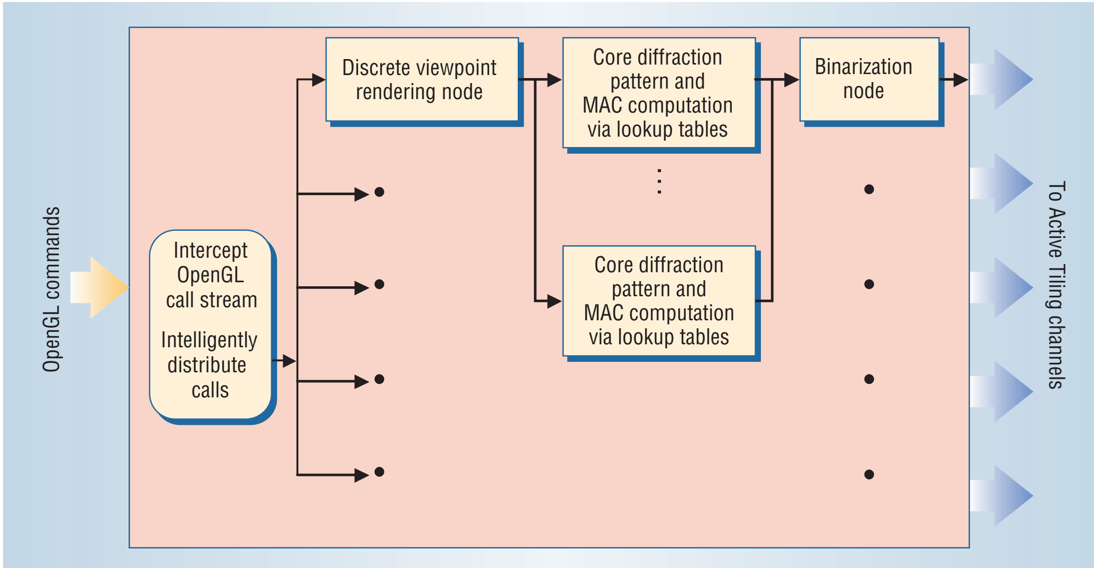

Figure 3. CGH pattern calculation pipeline for an Open GL capture, diffraction-specific algorithm system. The display system binarizes the hologram in the final computational phase using a partitioned approach and then passes the results to the optical subsystem over the parallel input architecture.

gram's information content, which often exceeds what the human visual system can interpret, these algorithms make it possible to achieve coarser resolutions, particularly during interactive image manipulation, that satisfy many applications.

Diffraction-specific algorithms accomplish this by spatially and spectrally quantizing the hologram as *hogels* and *hogel vectors*, respectively, that they can manipulate to vary the computational load. The algorithms can precompute much of the CGH calculation, storing the results as diffraction tables accessible via lookup during the interactive fringe-pattern calculation. The computational bottleneck in the latter procedure is *hogel vector decoding*. This step, in many forms of the algorithm, is a simple multiplication and accumulation calculation (MAC), open to optimization on numerous architectures.

COMPUTATIONAL LOADS

To better understand CGH computational load requirements, consider the following example. A horizontal-parallax-only computer-generated hologram with spatial resolution comparable to high-definition television and a comfortable viewing zone requires approximately 400,000 horizontal and 1,024 vertical pixels. A 4,096-pixel hogel can achieve spatial resolution exceeding common HDTV specifications, thus generating the hologram requires 100 hogels. Given that, on average, each hogel must generate half of the possible image points, the total estimated computational load is 100 (hogels) × 1,024 (image points) × 4,096 (diffraction-table lookup entries) × 1,024 rows = 424 giga MACs.

Spatial or temporal multiplexing can generate color, and requires three times as many MACs. Converting a MAC to floating-point operations, the total computational load for a single frame is 2.6 teraflops, which equates to the sustained throughput of the world's 130th fastest computer (www.top500.org/lists/2005/06). In addition, this figure is for 1 frame per second; most likely a system will need to meet at least 15 fps or faster to guarantee a perceptually smooth image update during interactive operation. However, in certain circumstances, optimizations can significantly reduce the computational load.

Hardware architectures

The computational load's large size and method of calculation—simple MACs from a lookup table—largely dictate the appropriate architecture. The entire computation is naturally parallel, with each hogel and all rows within each hogel independent. Candidate architectures include supercomputers (massively parallel, nonuniform memory architecture systems and high-performance clusters), special-purpose CGH compute engines such as MIT's Cheops,[8] symmetric multiprocessor computers, clusters of commodity PCs and graphics processing units,[9] digital signal processors (DSPs), field-programmable gate arrays (FPGAs), and application-specific integrated circuits (ASICs).

However, DSPs and FPGAs do not provide the kind of close, fast access to large memories that lookup table calculations offer. In addition, ASICs are not commercially attractive given the low initial volume of likely electroholography products. Thus, the most likely hardware architecture is based on high-performance computing systems.

In recent years, we have compared the performance and cost benefits of different architectures including the Cray T3E, the SGI Origin, the SGI Onyx Reality Engine, and Intel IA-32 processor clusters with various interconnects. Our results indicate that IA-32 systems with a Myrinet interconnect offer the most cost-effective and flexible high-performance cluster.

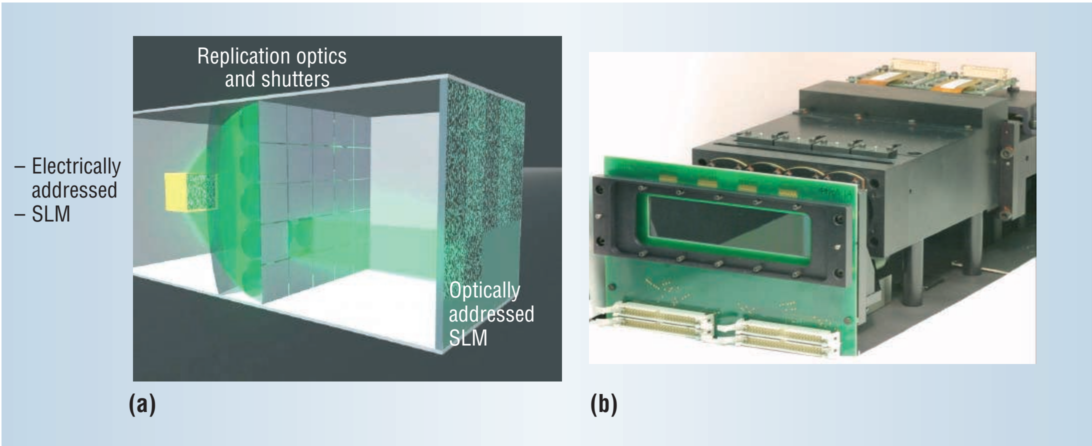

Figure 4. QinetiQ's Active Tiling system. (a) Schematic view of a single channel, which can generate 25 million pixels. (b) Photo of 1 × 4 channel Active Tiling unit capable of delivering full-color, temporally multiplexed computer-generated holograms with pixel counts up to 10^8.

Parallel inputs

Many SLM systems have parallel inputs, which eases the CGH computational challenge. For example, Figure 3 shows the flow of information from the application or data source, with each box representing one or more processing nodes. In the simplest case, each row of nodes could be associated with a single hogel, requiring perhaps 100 rows to meet the desired image size and field of view. Alternatively, the architecture could deploy varying numbers of processing nodes at each stage to improve the update rate or it could use one row to compute five hogels to reduce the computational cost.

The post-MAC result is a gray-scale CGH fringe pattern with 32-bit floating-point precision. The display system binarizes the hologram in the final computational phase using a partitioned approach[5] and then passes the results to the optical subsystem over the parallel input architecture.

Optimizations

Various optimizations—including pixel prioritization, dynamic level of detail, and wireframe mode for interaction[5]—can either accelerate CGH computation by increasing the frame rate or enable computation at the same frame rate with fewer resources, thereby reducing cost. Depending on scene geometry content, these techniques can reduce the computational load from 16 to 200 times. This translates to 162 gigaflops per frame at the most conservative level of savings. Given that a 3-GHz Intel Pentium 4 has a theoretical peak computational load of 12 Gflops per single-precision floating point, sustained throughput could be around 25 percent or 3 Gflops. Therefore, the computational load requires 54 processors per frame, or 810 for 15 fps.

QinetiQ researchers have also focused on achieving high, sustained throughputs on IA-32 processors using SSE2 vector capabilities to increase efficiency. If compute performance continues to adhere to Moore's law, the current requirement for 810 IA-32 processors—or 405 dual-processor nodes—will decrease to a modest cluster of fewer than 100 nodes by 2006.

Finally, for electroholography to be commercially successful, CGH systems must exploit a range of standard and customized visualization software packages. Ideally, users should be able to run their existing applications to drive the display without modification. A particularly powerful approach is to capture the graphics calls that applications such as OpenGL and DirectX produce and redirect them to the cluster-based CGH pattern calculation engine.

OPTICAL HARDWARE CONSIDERATIONS

Researchers have explored a number of SLM technologies, most notably the holovideo system[10] developed by the late Steve Benton and colleagues at MIT. Loosely based on the Scophony TV system of the 1930s, this 18-channel, parallel acousto-optic modulator and mechanical scanner can synthesize a 36-million-pixel computer-generated hologram and produce a horizontal-parallax-only image of $150 \times 75 \times 160$ mm^3 (W × H × D). However, it is not possible to scale the current system to deliver the 10^{10} pixel counts that a practical holographic workstation requires.

Active Tiling

A more recent development, and probably the leading scalable solution, is QinetiQ's Active Tiling system.[11] This system exploits the existing and projected strengths of both electrically addressed SLM (EASLM) and optically addressed SLM (OASLM) in an optimal way, trading the former's high temporal bandwidth for the latter's large spatial bandwidth.

As Figure 4a shows, an Active Tiling channel optically replicates or "tiles" output from a fast EASLM over an area of the OASLM. One configuration opens shutters in turn to build up the pattern on the write side of the OASLM, synchronizing the shutters to the EASLM's appropriate frames. This enables updating the entire OASLM at video rates.

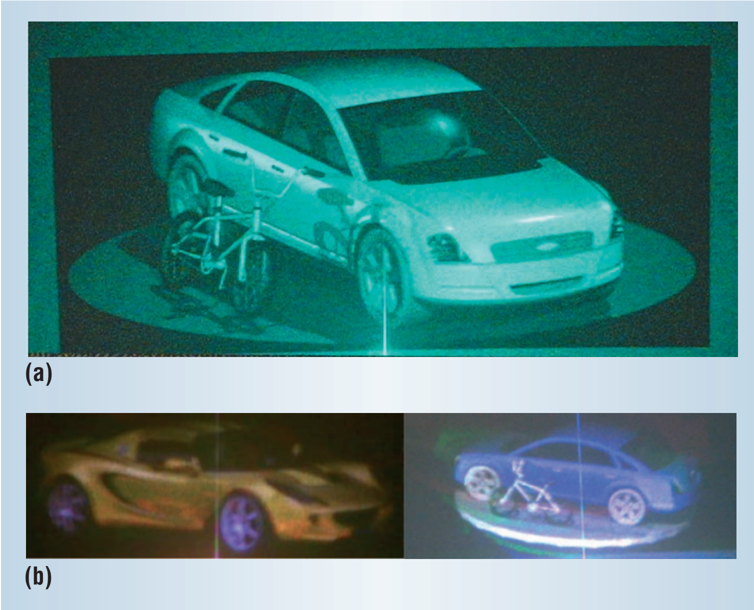

Figure 5. Computer-generated holograms. (a) Monochromatic and (b) color, full-parallax, 3D image produced by the 10^8 pixel Active Tiling system.

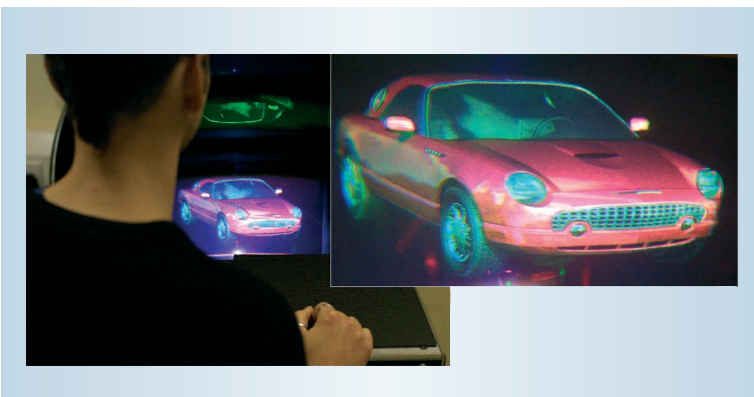

Figure 6. Replay of a spatial-multiplexed, 3 × 8 billion-pixel, full-parallax, full-color, 3D image.

Active Tiling channels are modular and can be stacked together in parallel to deliver the required pixel counts, as Figure 4b shows. This parallel approach also minimizes the required data feed rates for each channel and closely matches the capabilities of cluster-based rendering architectures.

A typical Active Tiling channel configuration consists of a binary, 1,024 × 1,024-pixel ferroelectric crystal on silicon EASLM operating at a 2.5-kHz frame rate; a binary-phase diffractive optical element and associated refractive optics performing the 5 × 5 replication; and an OASLM using an amorphous silicon photosensor, light-blocking layers, a dielectric mirror, and a ferroelectric liquid crystal output layer. The output for each channel is therefore 26 million pixels.

Current Active Tiling system attributes include a pixel areal density of more than 2.2×10^6 pixels cm^{-2}; a compact system volume exceeding 2.4×10^9 pixels m^{-3}; binary pixels, with a spacing of 6.6 μm; an updatable 1 × 4 channel arrangement (5,120 × 20,480 pixels = 104 megapixels) in both monochromatic and frame-sequential-color operation; and designs for scalable channel tilings in both height and width to deliver systems with a pixel count in excess of 10^9.

Replay optics

CGH configuration requires using various optical elements in conjunction with the SLM holding the holographic patterns.[5] For example, many SLM pixel spacings are relatively large compared to the wavelength of replay light. In such situations, an optical Fourier replay geometry can substantially reduce system volumes. Sophisticated multimirror configurations, folded in three dimensions, can further minimize overall system size. Using inexpensive carbon fiber or electroformed mirrors can also reduce system costs.

User interaction

An important element of CGH display systems in many applications is user interaction with the image. Researchers have developed intuitive interfaces using voice, gesture, and haptics. The synergistic effects of 3D sound can also be exploited to facilitate image interaction and modification.

EXPERIMENTAL RESULTS

Figures 5 and 6 illustrate state-of-the-art true 3D images generated at QinetiQ laboratories using what we believe to be the largest pixel-count CGH systems designed for this purpose. We calculated these computer-generated holograms using a 102-node PC Linux cluster of dual IA-32 Pentium III 1.26-GHz Tualatin-core 512-Kb cache CPUs, with 1 Gbyte of memory per node and two 36-Gbyte Ultra160 SCSI disk drives, along with a Myrinet interconnect and a 7.5-Tbyte storage area network.

Figure 5 shows monochromatic and color, full-parallax, 3D images produced by a computer-generated hologram having 10^8 pixels, while Figure 6 shows a replay of a spatial-multiplexed, 3 × 8 billion-pixel, full-parallax, full-color, 3D image. As these images illustrate, standard computer graphics techniques can be incorporated into CGH design, including environment mapping, radiosity modeling, and transparency effects. When viewed by eye, the full images look sharp and do not appear fuzzy—the eye naturally accommodates as the viewer focuses attention on different parts of the image, as with a real object.

The versatility of computer-generated holography, combined with its unique ability to produce full-depth-cue 3D images at beyond eye resolution, floating in space, and with an extended color gamut has led some to label CGH the ultimate display technology. However, many CGH-based displays have an appetite for pixels that can far exceed other display types. Unique additional computational operations add to the cost of such systems, particularly high-frame-rate interactive systems. Thus, for many applications, lower-cost, simpler display technologies will be more appropriate.

Nevertheless, with compute power and display hardware continuing to decrease in price and other required technologies rapidly advancing, the question is not whether CGH systems will become a practical generic display technology but, rather, how soon. ■

Acknowledgments

The authors gratefully appreciate the work of all members of the holographic imaging project team at QinetiQ as well as our academic and commercial collaborators. Funding for this research comes from QinetiQ, the UK Ministry of Defence TG7, the Ford Motor Company, and Holographic Imaging LLC.

References

1. S.J.Watt et al., "Achieving Near-Correct Focus Cues in a 3D Display Using Multiple Image Planes," *Proc. SPIE,* vol. 5666, SPIE—Int'l Soc. for Optical Eng., 2005, pp. 393-401.
2. I.J. Cox and C.J.R. Sheppard, "Information Capacity and Resolution in an Optical System," *J. Optical Soc. America A,* vol. 3, no. 8, 1986, pp. 1152-1158.
3. J.I. Trisnadi, "Speckle Contrast Reduction in Laser Projection Displays," *Proc. SPIE,* vol. 4657, SPIE—Int'l Soc. for Optical Eng., 2002, pp. 131-137.
4. C.D. Cameron et al., "Computational Challenges of Emerging Novel True 3D Holographic Displays," *Proc. SPIE,* vol. 4109, SPIE—Int'l Soc. for Optical Eng., 2001, pp. 129-140.
5. C.W. Slinger et al., "Recent Developments in Computer-Generated Holography: Toward a Practical Electroholography System for Interactive 3D Visualization," *Proc SPIE.,* vol. 5290, SPIE—Int'l Soc. for Optical Eng., 2004, pp. 27-41.
6. Y. Ichioka, M. Izumi, and Y. Suzuki, "Scanning Halftone Plotter and Computer-Generated Continuous-Tone Hologram," *Applied Optics,* vol. 10, no. 2, 1971, pp. 403-411.
7. M. Lucente, "Diffraction-Specific Fringe Computation for Electro-Holography," doctoral dissertation, Dept. of Electrical Eng. and Computer Science, MIT, 1994.
8. V.M. Bove Jr. and J.A. Watlington, "Cheops: A Reconfigurable Data-Flow System for Video Processing," *IEEE Trans. Circuits and Systems for Video Technology,* vol. 5, no. 2, 1995, pp. 140-149.
9. V.M. Bove Jr. et al., "Real-Time Holographic Video Images with Commodity PC Hardware," *Proc. SPIE,* vol. 5664, SPIE—Int'l Soc. for Optical Eng., 2005, pp. 255-262.
10. P. St.-Hilaire, "Scalable Optical Architecture for Electronic Holography," *Optical Eng.,* vol. 34, no. 10, 1995, pp. 2900-2911.
11. M. Stanley et al., "3D Electronic Holography Display System Using a 100-Megapixel Spatial Light Modulator," *Proc. SPIE,* vol. 5249, SPIE—Int'l Soc. for Optical Eng., 2003, pp. 297-308.

Chris Slinger *is a QinetiQ Fellow and technical director of the Optronics Department at QinetiQ Malvern in Great Malvern, United Kingdom. His research interests include holography, computer-generated holography, high-performance displays, and novel imaging systems. Slinger received a D.Phil. in multiple-grating volume holography from the University of Oxford. He is an Institute of Physics (IOP) Fellow. Contact him at slinger@qinetiq.com.*

Colin Cameron *is business development manager for high-performance visualization in QinetiQ's Optronics Products Group. His research interests include technical computing, high-performance graphics, and advanced display technologies for high-performance visualization. Cameron received a BSc (Hons) in laser physics and optoelectronics from the University of Strathclyde, Glasgow. He is an IOP Fellow. Contact him at cdcameron@qinetiq.com.*

Maurice Stanley *is a QinetiQ Fellow and business manager for high-performance imaging in QinetiQ's Optronics Products Group. His research interests include spatial-light-modulator technology, high-performance displays, and 3D camera systems. Stanley received a PhD in electrical and electronic engineering from University College London. He is an IOP Fellow. Contact him at mstanley1@qinetiq.com.*