

Volumetric 3D Displays and Application Infrastructure



With vendors lowering the barrier to adoption by providing compatibility with new and legacy applications, volumetric displays are poised to assume a commanding role in fields as diverse as medical imaging, mechanical computer-aided design, and military visualization.

Gregg E.
Favalora
Actuality Systems

Volumetric displays produce volume-filling three-dimensional imagery: Each volume element or *voxel* in a 3D scene emits visible light from the region in which it appears. Given their ability to project volume-filling autostereoscopic imagery, these displays are being adopted in fields as diverse as medical imaging, mechanical computer-aided design, and military visualization. Here, I use the term autostereoscopic to describe a display property that lets observers experience a 3D effect without requiring any additional eyewear.

Barry Blundell and Adam Schwarz, developers of the cathode ray sphere volumetric display, explain that “[a] volumetric display device permits the generation, absorption, or scattering of visible radiation from a set of localized and specified regions within a physical volume.”¹

The following volumetric-display taxonomy uses a vocabulary supported by previous researchers:¹

- *Swept volume.* Researchers proposed as early as 1912 that a volume-filling image can be produced by reflecting or transmitting light from a rotating or oscillating 2D surface within the desired 3D volume. It is possible to use a zero or 1D surface instead of a 2D surface, such as dragging an ink-filled stylus through a gel substrate. Normally, though, the surface is a plane or helix that translates or rotates through the volume. As the projection surface sweeps through the volume, it reflects or emits light

synchronized to its location. If the volume is refreshed frequently enough—say, 20 volumes per second—the observer will perceive a 3D image.

- *Static volume.* Volumetric displays can also generate 3D imagery by coaxing a volume into emitting light in which the bulk properties remain static. For example, glasses doped with rare-earth ions can emit visible points of light when excited by dual intersecting infrared laser beams.²
- *Holograms and holographic stereograms.* The 3D display community has yet to agree on whether or not holograms create volumetric images. Zebra Imaging has created holograms like the one shown in Figure 1 for applications such as automotive ergonomic studies; they generate full-color, full-parallax imagery that the user essentially perceives to be volume-filling.
- *Highly multiview 3D displays.* An emerging class of 3D displays physically reconstructs 3D lightfields by projecting from 30 to 200 views of a 3D scene with various trajectories from or through an image surface. In these cases, the imagery satisfies a so-called *super multiview condition* in which the observer’s eye automatically focuses on each voxel as if it were projected from that region of space.³ When placed in the middle of the image space, a diffuser, such as a business card, appears to be slicing into an object. These characteristics may



Figure 1. Volumetric holograms. Full-color, full-size holographic images such as this car interior suggest that computer-generated holograms should be considered volumetric displays. Even though commercially viable, interactive-rate holographic computer displays have yet to be developed, these static holograms made from CAD data sets are available today. Photo courtesy of Zebra Imaging.

compel developers to admit highly multiview 3D displays to the volumetric display family.

The literature includes a concise and fully illustrated history of volumetric displays⁴ and a thoughtful overview of autostereoscopic displays in general.⁵ In addition, the “Other Volumetric Displays” sidebar describes the development of other methods for generating volume-filling 3D imagery.

Several types of volumetric displays are already commercially available, including Actuality Systems’ Perspecta Spatial 3D System and LightSpace Tech-

Other Volumetric Displays

Many alternative methods for generating volume-filling 3D imagery exist, including the following canonical systems.

Stationary projectors

Several 3D displays, such as the Perspecta Display, require an embedded stationary projector to relay crisp imagery onto a rotating screen. Ordinarily, a fixed projector illuminating a rotating screen would induce considerable blur. Figure A shows one of the earliest solutions to this problem, described by Max Hirsch in an ingenious patent filed in 1958: Several mirrors are mounted to the same platform as the rotating screen, thereby ensuring that the focus does not change with respect to the screen’s angle.

Electron guns

Figure B shows a volumetric 3D display architecture that




Figure A. Patent illustration showing one of the earliest solutions to the problem of maintaining image focus while projecting onto a rotating screen. A series of mirrors (53, 54, 55) rotates along with the projection screen (56), keeping their relative positions constant regardless of the housing’s angle. (US patent 2,967,905)

employs electron guns that illuminate a rotating phosphor screen. Richard Ketchpel introduced one of the first designs for this technique in 1960. Barry Blundell and Adam Schwarz improved on this technique in several versions of their cathode ray sphere, which uses custom electronics and optimized vector-scanning software to generate animated volumetric imagery.

High-speed emitters

Figure C shows an emissive swept-screen display that uses a surface covered with high-speed emitters such as LEDs. Edwin Berlin invented one of the earliest examples of this display type, described in his patent filed in 1977.

Varifocal mirror

The Genisco/BBN SpaceGraph, a once commercially available 3D display, provided the focus for pioneering research on real-world applications. This display used a varifocal mirror to




Figure B. Patent for a “Three-Dimensional Display Cathode Ray Tube,” in which a cathode ray gun illuminates a swiftly rotating phosphor plate. (US patent 3,140,415)

nologies' DepthCube. Several vendors are developing software that lowers the barrier to adoption by providing compatibility with new and legacy applications. Our firm is developing a display-independent visualization environment with the aim of accelerating the adoption of 3D displays overall.

PROPERTIES AND SPECIFICATIONS

Volumetric 3D displays have many useful properties, including a wide field of view. In addition, these displays avoid accommodation and vergence mismatches, support multiple simultaneous ob-

servers, provide high-resolution imagery, and excel at several generic visualization tasks. A collection of advantages, parameters, and tradeoffs applies to many volumetric displays, including the following.

Motion parallax

Most volumetric displays have the extraordinary property of full parallax imagery, allowing one or more observers to inspect a 3D scene from a variety of horizontal and vertical viewpoints. Unrestricted lookaround is a key benefit of volumetric displays, as opposed to contemporary multiview

relay image components to various 3D depth planes and was used in applications such as medical imaging. Figure D shows an illustration of this technique from Alan Traub's 1966 patent application. Lawrence Sher described variations on this concept in his 1978 US patent 4,130,832.

Scanning displays

Beginning in about 1976, several organizations developed 3D displays that project a group of scanned zero-dimensional (point) images, rather than 2D images, onto a rotating screen.¹ These vector displays produced imagery with diameters greater than one meter. For example, the Space and Naval Warfare Systems Center developed several vector-scanned 3D displays using three lasers that illuminated a 36-inch-wide by 18-inch-high rotating helix.²

However, a variety of laser-scanning mechanisms generated the images, which limited the ultimate image complexity to the

bandwidth of the light modulation and scanning subsystems. The second-generation SPAWAR system, a technical tour de force when it was introduced in 1994, generated 40,000 voxels per each of three color channels at 20 Hz.² Contrast this with today's 3D displays, such as Perspecta, which repurpose off-the-shelf microdisplays to generate volumetric images composed of approximately 100 million voxels per color channel.

Other, more exotic scanned-beam 3D displays exist. So-called solid state or crossed-beam volumetric displays generate 3D imagery upon the interaction of dual light beams that intersect in the appropriate solid or gas.^{3,4} However, to date no commercially available systems use this methodology.

References

1. R. Hartwig, *Vorrichtung zur Dreidimensionalen Abbildung in Einem Zylindersymmetrischen Abbildungsraum*, German patent DE2622802C2, filed 1976, issued 1983.
2. P. Soltan et al., "Laser-Based 3-D Volumetric Display System (Second Generation)," *Naval Engineers J.*, vol. 107, no. 3, May 1995, pp. 233-243.
3. J.D. Lewis, C.M. Verber, and R.B. McGhee, "A True Three-Dimensional Display," *IEEE Trans. Electron Devices*, vol. 18, 1971, pp. 724-732.
4. E.A. Downing et al., "A Three-Color, Solid-State, Three-Dimensional Display," *Science*, vol. 273, 1996, pp. 1185-1189.




Figure C. Display that generates volumetric imagery by illuminating LEDs mounted on a rotating panel. (US patent 4,160,973)





Figure D. A system that uses a variable-focus mirror to re-image segments of a 3D scene to different focal depths, thereby creating virtual imagery. (US patent 3,493,290)



(a)



(b)

Figure 2. Perspecta Spatial 3D Display. (a) The system generates 10-inch-diameter volume-filling imagery with a full 360-degree field of view. (b) To provide the volumetric imagery, the display projects a series of 198 2D patterns, called slices, onto an optimized diffuser rotating at or above 900 rpm. The display sweeps the entire volume twice for every complete screen revolution, resulting in a visual refresh rate of 30 Hz.

lenticular or parallax barrier displays, which force observers to keep their heads uncomfortably within several centimeters of a designated viewing zone. Also, lenticular and parallax barrier displays permit only horizontal lookaround and are thus classified as *horizontal-parallax-only* displays.

Vergence and accommodation

Volumetric displays project 3D imagery with consistent *vergence* and *accommodation* cues. Vergence refers to the eyes' tendency to rotate so that their optical axes converge at the region they are gazing upon. Accommodation is a synonym for focus.

Because volumetric displays create imagery that truly occupies a spatial region, a viewer's eyes comfortably swivel to fix on the same point they are focusing on. This is not the case for many other types of 3D displays, such as typical stereoscopic goggles or lenticular and parallax-barrier displays.

Size

Display image size varies from less than one inch to approximately three feet across. Some researchers prefer to describe the image's spatial extent in terms of its volume.

Resolution

The resolution and frame rate of its projection engine or equivalent govern a volumetric display's

resolution. A multiplanar volumetric display typically uses one engine composed of three projector components, each with a resolution of $1,024 \times 768$ pixels. Such an engine might have a frame rate of 5 to 10 kHz. Volumetric displays generally have resolution advantages over flat-panel autostereoscopic displays such as those using lenticular and parallax barrier screens, which trade off display resolution for the number of views they present to the observer.

Volume refresh rate

As in computer graphics, a 2D display's visual refresh rate occurs independently of the rendering engine's performance. The same holds true for volumetric displays: The display might update the volume buffer's contents more frequently than the optical projection machinery can relay the images for observation.

Optical bandwidth

The optical bandwidth of the embedded image projector determines any display's visual performance. For example, several 3D displays use the Texas Instruments Digital Light Processing technology as an enabling component because it generates a resolution imagery of $1,024 \times 768$ pixels at speeds above 8,000 single-bit-depth frames per second. In terms of raw optical bandwidth, this corresponds to 6.3 Gpixels per second.

Integrated display systems

The elements of a typical volumetric 3D visualization system are analogous to the parts of a contemporary graphics workstation. The display—a complex integration of optical, mechanical, electronic, and algorithmic subsystems—is driven by a host PC that directs commands to it from the application software. To provide compatibility with legacy applications, most systems introduce a library that passes the graphics API calls to the display processor while permitting the traditional video card to use the original library. Extending the analogy, our firm created a platform that provides a display-agnostic executive control layer to the visualization system.

PERSPECTA SPATIAL 3D SYSTEM

Our team developed a 100-million-voxel swept-screen volumetric display.⁶ Pictured in Figure 2, Perspecta is a turnkey system that incorporates a high-resolution 3D display and a display-agnostic operating environment.⁷ The system uses the OpenGL API to interface to legacy 3D software packages, providing compatibility to a wide range of existing 3D applications. Perspecta's diverse applications range from medical visualization for interventional planning to battlefield simulation and molecular modeling.

Optical architecture

Perspecta generates 3D imagery by projecting a sequence of 2D patterns, or slices, onto a swiftly rotating omnidirectional diffuser screen. As Figure 2 shows, Perspecta's heart is a high-speed 2D projection engine that relays a series of approximately 6,000 images per second onto the rotating screen. A series of relay mirrors also rotate about the screen's axis, thereby guiding the slices to the screen with focus and keystoneing that is invariant with instantaneous screen angle.

Perspecta's projection engine is based on the Texas Instruments Digital Light Processing technology, which uses a microelectromechanical-systems-based reflective microarray. Perspecta has DLP-based microdisplay chips for red, green, and blue color components. The system uses a region of 768×768 pixels for each microdisplay, updating the image every 168 μ s. Each voxel in the display space has, at highest spatial resolution, 3-bit color. However, the observer gets the benefit of full color through the system's use of proprietary volumetric dithering, quasi-half-toning algorithms.

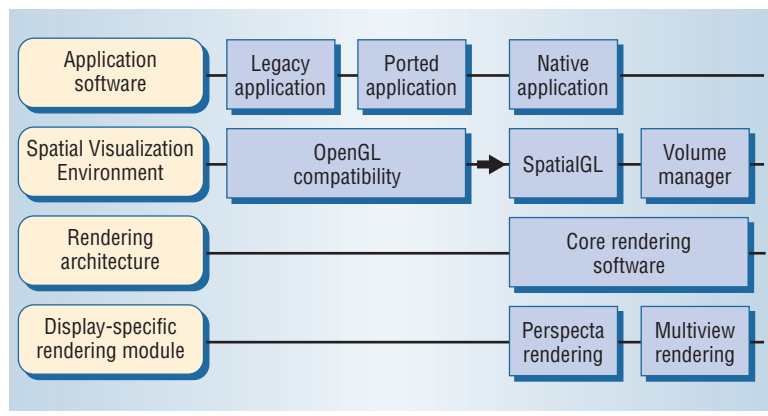


Figure 3. 3D rendering OS and hardware. The high-level software architecture consists of key components such as a spatial visualization environment, OpenGL compatibility layer, and core rendering software and electronics.

3D rendering OS and hardware

The Perspecta software platform employs a technology that unifies true 3D rendering under a single display-agnostic interface that lets the same application run in displays as different as a spinning-screen volumetric display and a 40-view quasi-holographic display. The Spatial Visualization Environment, which plays a key role in providing this capability, uses a new API—SpatialGL—to provide display-independent connectivity between applications and displays. Figure 3 shows the SVE's high-level software architecture, which consists of the following components:


Spatial Visualization Environment. The SVE provides a uniform abstraction of graphics with 3D displays. The SpatialGL API forms the basis for representing graphical assets and operations. A volume manager provides a virtualization layer to allocate and share rendering and display resources. It also provides mechanisms for interactive user feedback, such as supporting a 3D pointer.

OpenGL compatibility layer. To lower the barrier to adoption, the SVE also provides a compatibility layer to applications written using the OpenGL API. This layer works with legacy applications that were written without special consideration for 3D displays, such as molecular modeling software, VRML model viewers, and others.

Core rendering software. The core rendering software implements a Spatial 3D Server that provides the execution environment to run rendering operations for a particular 3D display. One major and original aspect of core rendering software, the rendering framework, lets new display types with peculiar geometries take full advantage of the existing graphics accelerators. For example, we have demonstrated the same application on two very different 3D displays: the Perspecta multiplanar display and an experimental, 40-view quasi-holographic display.

Core rendering electronics. Finally, a core rendering electronics subsystem delivers considerable

Figure 4. DepthCube volumetric display. (a) The display generates imagery by projecting a sequence of 2D images into a multiplanar optical element. (b) The display's subsystems. A computer relays 3D data to a dual-volume buffer. A digital projector relays the sequence in the buffer into a multiplanar optical element so that each relayed image corresponds to the single liquid crystal panel. Photograph and diagram courtesy of LightSpace Technologies.



optical and computational bandwidth to a variety of 3D displays. This subsystem accepts data from single or dual Gigabit Ethernet ports, executes the Spatial 3D Server environment using Linux running on a 2-GHz AMD Athlon 64 chipset, and performs proprietary high-speed geometric and volumetric rendering operations using an onboard Nvidia GeForce Ultra GPU. The subsystem reads computations over the GPU bus, making them available to three high-speed DLP projector drivers at the system output. A voxel engine, implemented in FPGAs, performs voxel routing and memory management tasks.

Display independence

These technologies embody a 3D visualization platform that incorporates a volumetric display, custom 64-bit graphics engine, a display-agnostic Spatial Visualization Environment, compatibility with legacy applications through the OpenGL API, and SpatialGL, a new API that exploits the unique functionality of a variety of 2D and 3D displays. The system's display independence provides a unique and valuable characteristic that lets it operate with a wide variety of hardware and software.

DEPTHCUBE

Figure 4 shows LightSpace Technologies' DepthCube, a volumetric display with a translational rather than rotational multiplanar architecture.

Optical architecture

The multiplanar display's projection volume contains no macroscopic moving parts. Instead, it relies on persistence of vision to generate 3D scenes.

Rendering software acquires and processes data from native and legacy applications. Once the system has sampled a desired 3D scene into a linear multiplanar representation, a high-speed digital projector reconstructs the scene by illuminating the successive planar regions of a *multiplanar optical element*. This 3D MOE volume consists of a stack of liquid crystal panels oriented parallel to the front

face of the 3D display. Each panel can switch between a transparent and light-scattering state. Generally, at any instant, every panel is transparent except for one: An MOE driver controls the MOE stack to synchronize the high-speed projector's 2D slice projection to the observer's appropriate depth plane. The DepthCube projects a series of 20 frames onto the MOE in synchrony with the time-varying location of the single light-scattering panel.

This process repeats 50 times per second in non-interlaced mode. The result is a $1,024 \times 768 \times 20$ -voxel, 15-bit-color image with full parallax and a view zone approximately equal to the extent provided by the portal that the system's front face affords. As in the Perspecta display, a core component of the DepthCube is a DLP-based digital projector subsystem, chosen for its high frame rate, contrast, and brightness.

Rendering

Rendering to the DepthCube 3D display relies upon the host PC's graphical processing unit, which underscores how researchers are repurposing GPUs for fields other than 2D display. The display's Cartesian geometry maps easily into the GPU, which has an inherently Cartesian coordinate system. Further, the presence of a z -buffer within the GPU provides the depth of every visible pixel in the color buffer so that DepthCube updates can occur at approximately real-time rates of less than 15 Hz for images containing single 3D surfaces. LightSpace Technologies uses the GPU's facilities for 3D texture maps, vertex and fragment shaders, and multiple rendering passes.

At this point, the system must read the rendered contents of the graphics card frame buffer back into main memory so that it can transfer them to the DepthCube. Transfer occurs through two mechanisms: the DepthCube API and the DepthCube GLInterceptor. The former provides direct control over the rendering and transfer processes, while the latter provides access to supported legacy applications that use the OpenGL API.

DepthCube API

The DepthCube contains dual 45-Mbyte multiplanar frame buffers that store the color and brightness of every voxel within the display, enabling the system to write the back buffer while the front buffer displays. The API can use the multiplanar frame buffer to display multiple 3D images, each at physically different depths, in both surface and volumetric data formats.

The API supports several addressing modes, such as plane and block modes. Given the contents of two adjacent depth planes, the plane mode summons antialiasing algorithms that improve the system's interplane perceptual color depth. Block modes offer the programmer direct control over individual voxels per plane. Additional functionality includes selectively updating small regions of the volume and overlaying multiple image types by relying on manual buffer swapping.

GLInterceptor

The GLInterceptor mechanism supports certain applications that were not written for the DepthCube but which utilize the OpenGL API. In the Microsoft Windows implementation, the GLInterceptor consists of two software components: the `openGL32.dll` dynamically linked library placed in the application directory, and a control panel executable accessed through the system tray. OpenGL applications load the GLInterceptor DLL, which then loads the original `openGL32.dll` from a system directory. The interceptor extracts and computes data relevant to the DepthCube, such as near and far clipping planes and the color of each pixel. In directly accessing the z-buffer, the current methodology obtains only those pixels closest to the viewer.

FUTURE ADVANCES

In the future, we predict the addition of several features and capabilities to volumetric displays.

Variable-opacity voxels

Today's commercially available volumetric displays generate translucent imagery in which every voxel functions as an omnidirectional emitter. Thus, these displays cannot depict regions with programmable opacity or effects based on the position of multiple observers. This limitation derives predominantly from the choice of omnidirectional diffuser screen surfaces and is not, as traditionally reported, an artifact of a rotating or translating screen.

SeeLINDER display

Generally, parallax-barrier displays would not be




Figure 5. Volumetric model-building application. The application can recognize a wide array of camera-tracked gestural-interaction techniques. Using a ray cursor, the viewer can drag and drop regions of the model by pointing at the display. Photo courtesy Tovi Grossman, University of Toronto Dynamic Graphics Project.

considered classically volumetric, but one group of displays in this class do wrap into a cylinder and create horizontal-parallax-only imagery. Examples include the SeeLINDER and earlier monochrome versions.³ These models use a series of bright LEDs mounted along the circumference of a rotating cylinder. A cylindrical parallax barrier rotates opposite the LEDs, while control electronics direct the LEDs to light in such a way that lines drawn from the potential positions of the observers' eyes intersect in the region of their quasivolumetric reconstruction. By definition, this view-based display can render voxels with programmable opacity.

Human-computer interfaces

Volumetric display users indicate that they want the display to function as a direct interaction device rather than having to shift focus between the host PC's 2D display and the volumetric display. Volumetric human-computer interfaces enable applications designed specifically for volumetric displays. A recent and groundbreaking body of work⁸ in this area used an early Perspecta Spatial 3D Display to develop methods of direct interaction with 3D imagery through visually tracked finger gestures that control a 3D geometric model-building application, as shown in Figure 5.

These researchers developed an array of interaction methods. For example, a hand-pointing gesture can accomplish model selection and dragging. The user can add primitives to the scene by selecting among 3D models presented in a unique surface browser that covers the display's outer spherical surface.

A wealth of interaction modes are also available, including surface menus accessed by pressing menu items affixed to the display surface; a variety of postures and gestures such as point, pinch, curl, and trigger; model transformations such as scaling, rotation, and translation; and constrained transformations that restrict one object from entering the space of another.

The future of volumetric displays will see advances in system software, processing electronics, and the ability to depict regions of variable opacity and other viewer-position-dependent effects. Advances in these areas are gated by progress in high-speed MEMS microdisplays, GPU chipsets and manufacturer support of high-bandwidth readout modes directly onto the system bus, and the availability of point-source-like high-power illumination subsystems. Existing industry forces in the realms of high-performance computing and information display will likely provide ongoing advances in these areas. That said, the most urgent requirement for widespread adoption of volumetric 3D displays is the ability to integrate tightly into today's visualization software. Then, volumetric displays can be deployed across the fields where they are most critically needed: medical imaging, oil and gas exploration and production, and military command and control. ■

Acknowledgments

I thank Oliver Bimber, who provided the opportunity to share the viewpoints expressed here, and my colleagues at Actuality Systems, who created

the Perspecta Spatial 3D technologies, particularly SVE architects Joshua Napoli, Won-Suk Chun, and T.J. Purtell II. Thanks also to the following: the US government for supporting significant portions of this work under a NIST Advanced Technology Program contract; Michael Klug for providing the photograph of the Zebra Imaging hologram; Alan Sullivan and John Snuffer for providing the illustrations and descriptions of the LightSpace 3D display system; and Tovi Grossman of the University of Toronto for providing illustrations of his award-winning volumetric HCI work.

References

1. B. Blundell and A. Schwarz, *Volumetric Three-Dimensional Display Systems*, John Wiley & Sons, 2000.
2. E.A. Downing et al., "A Three-Color, Solid-State, Three-Dimensional Display," *Science*, vol. 273, 1996, pp. 1185-1189.
3. T. Honda et al., "Three-Dimensional Display Technologies Satisfying 'Super Multiview Condition,'" B. Javidi and F. Okano, eds., *Proc. Three-Dimensional Video and Display: Devices And Systems*, vol. CR76, SPIE Press, 2000, pp. 218-249.
4. K. Langhans, et al., "New Portable FELIX 3D Display," *Proc. SPIE*, vol. 3296, SPIE—Int'l Soc. for Optical Eng., 1998, pp. 204-216.
5. M. Halle, "Autostereoscopic Displays and Computer Graphics," *Computer Graphics*, ACM Siggraph, vol. 31, no. 2, 1997, pp. 58-62.
6. G.E. Favalora et al., "100 Million-Voxel Volumetric Display," *Cockpit Displays IX: Displays for Defense Applications*, *Proc. SPIE*, vol. 4712, SPIE—Int'l Soc. for Optical Eng., 2002, pp. 300-312.
7. W.-S. Chun et al., "Spatial 3D Infrastructure: Display-Independent Software Framework, High-Speed Rendering Electronics, and Several New Displays," *Proc. SPIE-IS&T Electronic Imaging*, vol. 5664, 2005, pp. 302-312.
8. T. Grossman, D. Wigdor, and R. Balakrishnan, "Multi-Finger Gestural Interaction with 3D Volumetric Displays," *Proc. ACM Symp. User Interface Software and Technology*, ACM Press, 2004, pp. 61-70.

Gregg E. Favalora is a founder of Actuality Systems Inc. His primary research interest is the development of volumetric, multiview, and holographic displays. Favalora received an SM in engineering sciences from Harvard University. He is a member of the IEEE, SPIE, and OSA. Contact him at favalora@actuality-systems.com.



**JOIN A
THINK
TANK**

Looking for a community targeted to your area of expertise? IEEE Computer Society Technical Committees explore a variety of computing niches and provide forums for dialogue among peers. These groups influence our standards development and offer leading conferences in their fields.

Join a community that targets your discipline.

In our Technical Committees, you're in good company.

www.computer.org/TCsignup/