The A-Desk: A Unified Workspace of the Future

A Surround Environment for Seamless Interaction with Physical and Digital Media

Andrei State
University of North Carolina at Chapel Hill
InnerOptic Technology, Inc.

Herman Towles
University of North Carolina at Chapel Hill
Applied Research Associates, Inc.

Tyler Johnson
University of North Carolina at Chapel Hill
Scalable Display Technologies, Inc.

Ryan Schubert
University of North Carolina at Chapel Hill
University of Central Florida

Brendan Walters
University of North Carolina at Chapel Hill
Pendo.io, Inc.

Greg Welch
University of North Carolina at Chapel Hill
University of Central Florida

Henry Fuchs
University of North Carolina at Chapel Hill

It is made from wood, projectors, and a computer running multiprojector blending software. Developed in 2008–2009 in response to an agency interest in future work environments for intelligence analysts, the immersive curved A-Desk supports a variety of tasks, such as traditional office or development work, design, and even immersive telepresence (Figure 1).

Digital Object Identifier 10.1109/MCG.2019.2951273
Date of current version 6 January 2020.

INTRODUCTION

The A-Desk ("Analyst Desk") is a user-enveloping display system consisting of a bowl-shaped, two-dimensionally curved display and work surface. The A-Desk concept and working prototypes were developed after consultations with the U.S. National Geospatial-Intelligence Agency (NGA) and with support from the U.S. Intelligence Advanced Research Projects Activity (IARPA) (Future Analyst Workspace (A-Desk), Principal Investigators Henry Fuchs and Greg Welch, U.S. Air Force Office of Scientific Research, Agency

Figure 1. Left: Conceptual drawing of A-Desk system in intelligence analysis scenario. Right: 160-degree prototype constructed at UNC in 2009, demonstrating a (simulated) medical remote collaboration scenario. Note the selectively illuminated physical document on the horizontal desk surface.

Figure 2. A-Desk prototypes running the IARPA demo application with freely movable (simulated) animated windows and OpenGL 3-D graphics in the background. Left: 90-degree prototype illuminated by blended imagery from two projectors mounted high above the unit (February 2009). Right: 160-degree system; the curved display is produced by blended imagery from seven projectors mounted above the display surface and visible at the top of the image (July 2009).

Number FA8750-08-2-0209, IARPA Analyst Workspace for Exploitation (A-SpaceX) Program, Program Manager Dr. Jeffrey G. Morrison).[1] The A-Desk aims to address these agencies' need for innovative display technologies for use by analysts whose complex daily tasks include visualizing intelligence and surveillance data from multiple sources comprising different modalities, developing and managing hypotheses about past and future activities, scenarios, or events, and communicating such information to, as well as discussing and processing it with, collaborators.

Our A-Desk design combines ideas already present in IARPA presentations associated with the "Analyst Workspace for Exploitation" (A-SpaceX) Program (managed by Dr. Jeffrey G. Morrison) as well as in our own, independently developed concepts, including some aspects of our 1998 "Office of the Future" and related work.[2,3] We developed two modular A-Desk hardware prototypes (Figure 2) and a number of software prototypes implementing an IARPA-specific demo application and several other demonstrations. The demonstrated system showcases a variety of features, including the seamless transition between the curved vertical display and the horizontal work surface, automatic image-based calibration, gesture recognition, and hand tracking. The A-Desk design also includes the capability to recognize and digitize physical materials such as documents, which we had demonstrated in related preceding work.[4]

Multiprojector Display

Before the development of large flat-panel displays, developers used projectors to create large virtual environments. In 1992, the CAVE[6] was showcased at Siggraph. This 4-wall (3 side walls and floor) immersive theatre was implemented using four rear-mounted CRT projectors with synchronized visuals rendered by multiple rendering pipelines. Geometric continuity at the display wall boundaries was controlled largely by mechanical manipulation of each projector's pose. Position tracking of the user's eyes within the rear-projection display space enabled perspectively correct rendering of visuals, even across room corners.

With the introduction of cheaper commodity LCOS/DLP projectors and user-programmable graphics hardware in the late 1990s, many researchers began developing higher-resolution visualization walls by simply scaling up the number of projectors. These rear-projected solutions most often required custom room designs and used a complex manual calibration process of adjusting projector pose and/or clipping (electronically or physically) overlapping images to achieve seamless geometric tiling of the display array.

Meanwhile, a research team at the University of North Carolina developed a vision of converting any room into a visualization environment with front projection and cameras that automate the process of "calibrating" an array of casually placed projectors. The rendering task was to create geometrically seamless imagery on arbitrarily shaped surfaces illuminated by multiple overlapping projectors, which requires knowledge of each projector's lens characteristics (focal length and distortion) and spatial pose (position and orientation), of the three-dimensional (3-D) shape of the display surface, and of the 3-D position of the human observer, all in a common coordinate system. Creating photometrically seamless imagery when projector images overlap requires knowledge of which pixels of each projector should be attenuated such that the total illumination is consistent.[7]

The key to solve these issues was the realization that a projector can be used as the equivalent of a camera, and that it is possible to use a calibrated stereo camera pair and projected structured light from each projector to both calibrate each projector and reconstruct the shape of the entire display surface in a common 3-D coordinate system. This 3-D information is also sufficient to compute an attenuation or blending mask for each projector in order to achieve photometric uniformity.

Given this information, and with knowledge of the viewer's 3-D position in the display space, it is then possible to efficiently render pre-distorted and blended images (using GPU shader programs) for an array of casually positioned projectors, producing geometrically and photometrically seamless imagery on the walls of any room. The OpenGL-based software framework developed at UNC for this purpose, and which encapsulates both calibration and rendering, is referred to as Wide Area Visuals (WAV).[8,9]

While our A-Desk prototype was previously mentioned briefly, in connection with a telepresence application,[5] this article contains the first complete description of the system, its background, implementation rationale, and technical details. Even though the work described here was performed in the 2008–2009 time frame, we believe that the seamless combination of real and virtual surfaces in an everyday work environment is still interesting today. Furthermore, the overall concept presents an interesting alternative to the currently popular head-worn immersive displays, for scenarios where a fixed display is appropriate, with benefits such as a wide field of view, natural interaction, and comfortable multi-user collaboration.

BACKGROUND AND MOTIVATION

One important motivation related to the A-Desk display was the need for displays with very high overall resolution (for the 2008 time frame) so that analysts might reliably identify certain small objects in satellite imagery. Normally, the analysts were using specialized desktop monitors; however, the monitors were dedicated to the sole task of high-resolution image analysis and required regular manual calibration by a technician to ensure the maintenance of effective resolution. Another motivation related to the integrated physical-virtual nature of the A-Desk was that all of the other material the analysts relied upon, including computer-based and printed reports, news stories, and other analysts' reports, was physically separated from the high-resolution imagery, complicating the correlation and integration of information. A third motivation related to collaboration was the difficulty analysts were having in sharing findings, ideas, and investigation state to other analysts, e.g., for

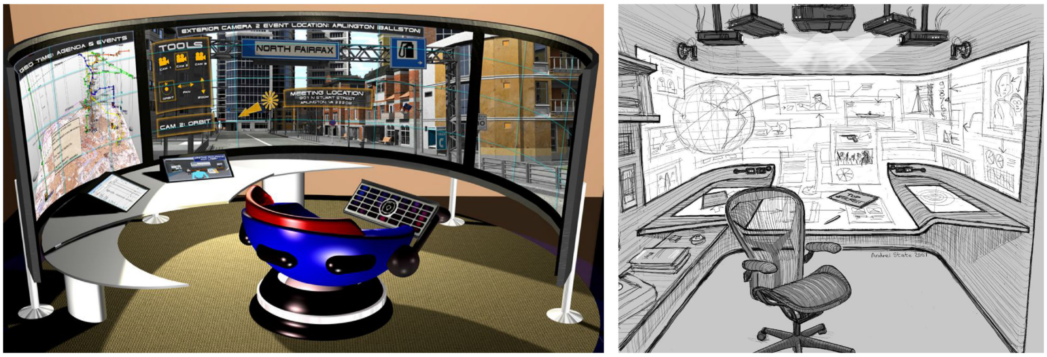

Figure 3. Left: IARPA surveillance concept illustration[10] (courtesy of Jeffrey G. Morrison; image created by Loma Media). Right: Independent UNC concept (2007) of projection-based office cubicle with colocated, blended surfaces for imaging/scanning, illustration (drawing/digitizer) and display. Multitouch interaction concepts such as pinch-zoom were not yet common at the time.

handoffs at shift changes. This is of course based on the information available to us at the time we began the development of the A-Desk concept in early 2008. The (highly secretive) intelligence community may have addressed and/or solved these issues in other ways over the past decade.

Previous research related to the IARPA Intelligence Analysis display challenge includes the already mentioned Office of the Future Project,[2] which described an office environment paradigm characterized by ubiquitous projective display on all available surfaces regardless of their shapes. The system proposed the use of computer vision, image processing, and traditional graphics to continually extract information, such as surface shape and characteristics for walls, furniture, arbitrary objects, and even people. That information can be used to project imagery on those surfaces such that it appears perspectively correct from any desired points of view—normally, from the locations of various head-tracked users in the environment. The paper's proof-of-concept prototype inspired much follow-on work, including in our lab. In some ways, the A-Desk is a more near-term realization of that "Office" vision, targeting a specific application domain, yet attractive for general use. The A-Desk display depends on seamlessly blended multiprojector imagery (see sidebar), a widely used technique at the time, today mostly displaced by large very high resolution panels and video hardware, and interfacing standards that can drive many millions of pixels from each individual output.

Beyond scientific research, the similarity of the display shown here—in particular when showcased with the "hand tracking" prototype described in the Applications section—with the imaginative, clairvoyance-based precrime analysis system in the 2001 sci-fi motion picture "Minority Report" is evident, and we acknowledge the influence of that film's production design and technical advisors.

EARLY CONCEPTS AND SIMULATION

When the Intelligence Analysis Display challenge was issued, IARPA had already produced some illustrated concepts for the possible workstation they were envisioning (Figure 3, left). At the same time, UNC scientists had been independently investigating "civilian" office scenarios, such as the office cubicle shown in Figure 3, right, and had demonstrated integrated processing of physical documents in such environments.[4] Our eventual A-Desk design (Figure 2) incorporated notions from all three of these, anticipating the deployment of technical methods first introduced in the Office of the Future concept, but improved, adapted, and streamlined as needed.

Our first task was to determine whether it would be possible to smoothly illuminate the large-field-of-view surround shape we were envisioning with blended and therefore overlapping imagery emitted by multiple projectors, and whether the required projectors could be practically mounted and operated such that they would not interfere with routine A-Desk

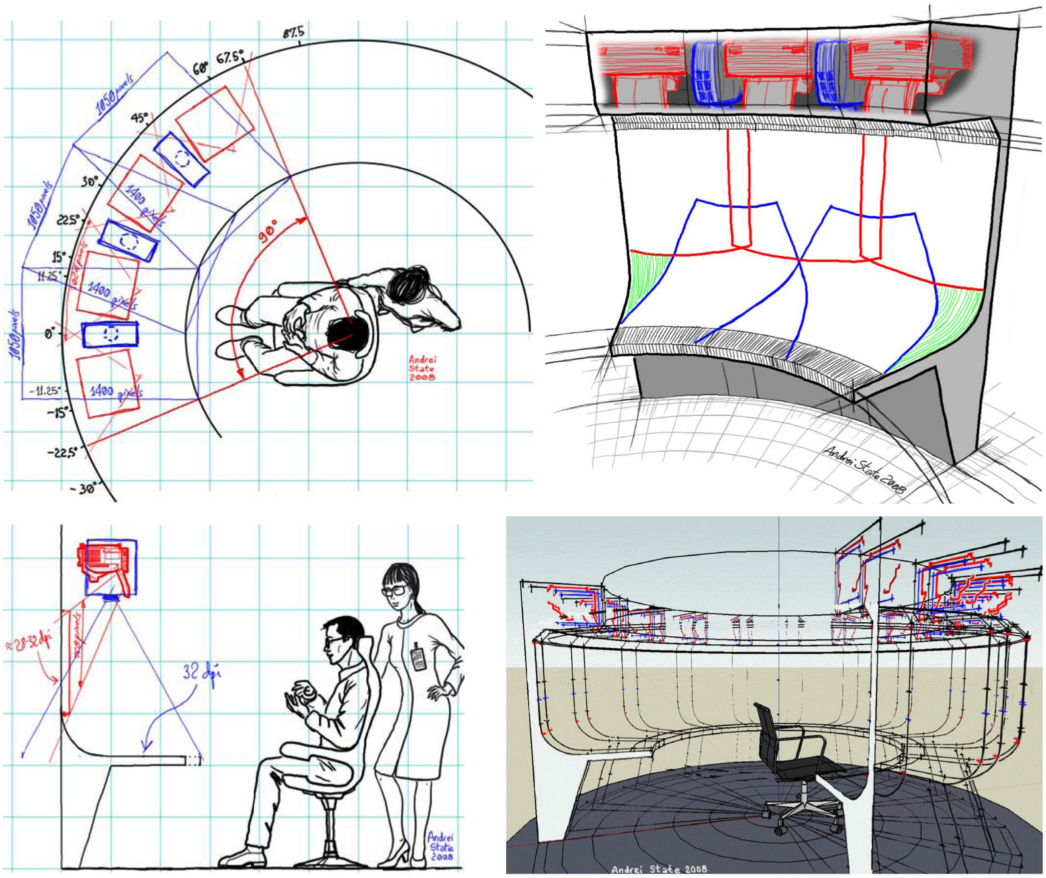

Figure 4. Left column and top row: Conceptual illustrations for work/display surface and determination of compact projector layout, with approximate pixel density calculation. Bottom right: 3-D concept (in Sketchup) of nearly 270-degree prototype with large diameter, suitable for multiuser work (cf. Figure 1, left), created by rotational development of the cross section at bottom left. (February 2008)

use. We considered issues such as users' freedom of motion, minimizing shadows cast by users, heat dissipation, and effective resolution on the display surfaces. Drawing from previous experience with camera and projector placement designs,[11] we conducted basic geometric preplanning and developed early concepts via hand-drawn sketches and simple calculations, which then evolved into basic engineering diagrams (Figure 4, top and left) and Sketchup-based 3-D models (Figure 4, bottom right—such a model was also used as the basis for the illustration in Figure 1, left).

PROTOTYPE CONSTRUCTION

The geometry developed in the early planning phase (Figure 4) served as the basis of full-scale mock-ups (Figure 5, left), which preceded the actual CAD design and construction of the prototype (Figure 5, center and right).

The operational prototype surround workspace (Figure 5, right, and Figure 6) is a hemicylinder with a diameter of 108 inches. The desktop surface is 30 inches deep with a central 48-inch-diameter user opening. The projection surfaces (desktop and walls) are covered with a diffuse, white laminate surface.

The display imagery is produced by seven projectors mounted in the cylindrical *bridge* located at the top of the surround surface. Three Projection Design F20 projectors are positioned 39 degrees apart and are aimed to project onto the desktop and onto part of the cylindrical walls (yellowish shapes in Figure 6 left, numbered as 2/3/5 in Figure 6, right). Four very-short-throw NEC WT610E projectors are interleaved between the F20 units and are aimed to illuminate

Figure 5. Left: Foam-and-cardboard mock-up for usability evaluation (May 2008). The bulky top-mounted volumes correspond to the projector mounting areas (cf. Figure. 4). Center and right: Wooden construction of the prototype's vertical cylindrical frame (October 2008).

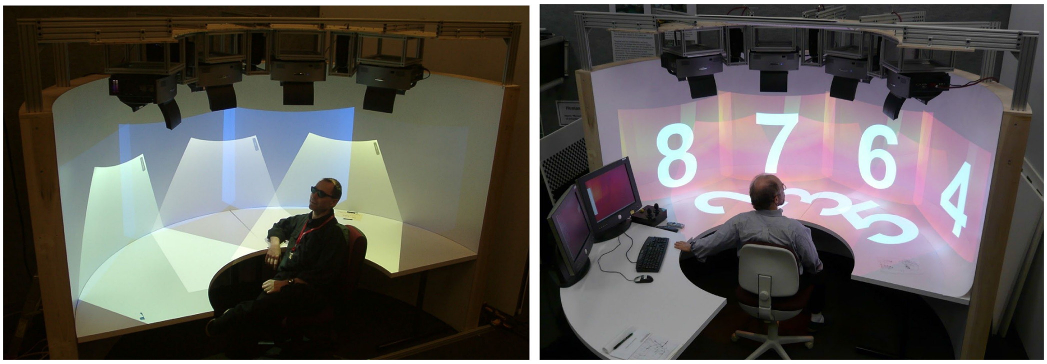

Figure 6. Initial display tests of the 160-degree prototype surround workspace showing areas illuminated by the seven projectors. At left, the curved (toroidal) transition surface between the vertical cylinder and the horizontal surface has not been attached yet (early 2009).

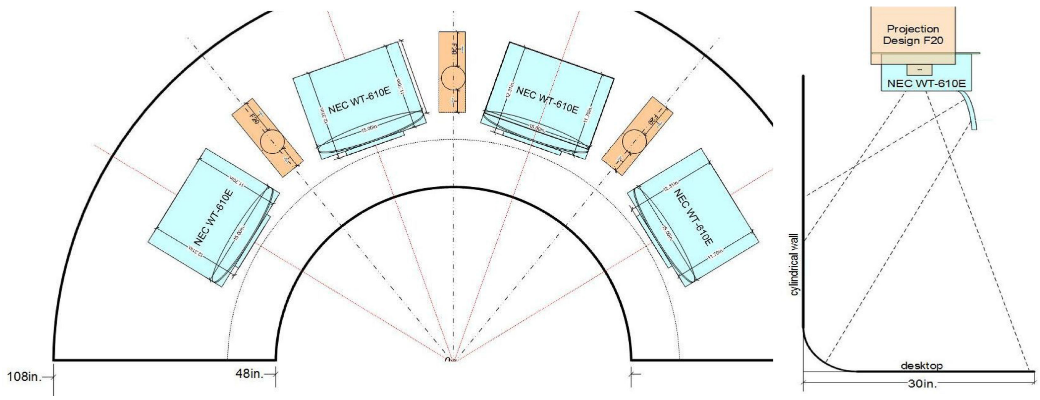

Figure 7. Left: Top view of the projector layout above the surround workspace. The four NEC projectors (numbered 4, 6, 7, 8 in Figure 6, right) and three Projection Design projectors (numbered 2, 3, 5 in Figure 6, right) are shown in blue and orange, respectively. Right: Side profile of the surround showing projector layout relative to cylindrical wall and desktop surface (cf. Figure 4, bottom left).

approximately 160 degrees of the cylindrical wall in front of the user (bluish shapes in Figure 6, left, numbered as 4/6/7/8 in Figure 6, right). Figure 7 (left) details the final positioning of the projectors mounted in the overhead bridge relative to the cylindrical wall.

It was deemed desirable early on, for both technical and aesthetic reasons, to flair the vertical cylindrical display surface into the horizontal desktop surface (Figure 3, right, and Figure 4, top right) such as to avoid a 90-degree first-order discontinuity in the display surface, and thus in the imagery.

To achieve this complex curved surface in the prototype, we simply positioned overlapping 8.5” × 11” sheets of heavy white paper to create a smooth curved surface between the two display surfaces, as is clearly visible in Figure 2, left.

In addition to the main 160-degree unit described above, we also built a smaller 90-degree module (Figure 2, left). It was used for early testing and then for parallel development of software applications. This smaller module was designed to be attached to the larger one if desired, creating a nearly 270-degree system as illustrated in Figure 1, left.

DISPLAY AND COMPUTATIONAL HARDWARE

As previously detailed, the workspace uses four NEC WT-610E and three Projection Design F20 projectors in the surround display (Figure 7). The NEC projectors feature an unusual design with aspheric mirrors and a folded optical path designed to produce a large 1024 × 768 resolution projection from an ultra-short throw distance, while the ProjectionDesign units have a conventional short-throw lens design and project at 1400 × 1050 resolution.

In addition to those seven projectors, the workspace is augmented with two flat panel displays and a keyboard positioned on a quarter-circle extension of the desktop surface. These components are used to manage the startup of the A-Desk’s system software, as well as to provide a conventional user interface. They are visible at the left edge of the image in Figure 2, right, and Figure 6, right.

The entire display configuration is hosted by a single-processor deskside PC with AMD 2.5-GHz quad-core CPU, equipped with four dual-output NVIDIA Quadro 580 GPUs. Seven of the eight available displayport outputs are connected to the top-mounted projectors, while the 8th output is interfaced with a Matrox DualHead2Go unit that digitally splits the 3200 × 1200 resolution input signal into two (1600 × 1200) channels to drive the two flat panel displays. The deskside PC is running Microsoft Windows XP.

To facilitate basic projector control such as ON/OFF, RS-232 serial ports on each projector are interfaced with a Digi multi-channel ethernet-to-serial server.

SOFTWARE FRAMEWORK

The UNC WAV software framework (see sidebar)[8,9] is a custom library connecting the application code and the OpenGL driver for the four Nvidia Quadro 580 GPUs driving the projectors; WAV offers the application a single OpenGL graphics context across all projectors, with no application-level awareness of geometric or photometric blending. WAV uses Nvidia’s “GPU Affinity” capability to optimize performance. The framework supports static camera-based calibration (see the Calibration section).

After each display frame is rendered by the application, it is subsequently warped and blended. This task proceeds by first rendering the virtual 3-D model of the display surface, texture-mapped with the application image. This texture-mapped 3-D display surface model is then rendered using the projection matrix of each projector to yield the prewarped image to be displayed by each projector. All projector’s prewarped images combine to form the complete application image on the physical screen surface. Before each prewarped image is finally sent to its projector for display, its precalculated static blend mask is applied, in order to achieve the uniform intensity of the displayed imagery across all projectors.

The WAV framework supports arbitrary observer viewpoints and is not even limited to continuous surfaces, or to surfaces that can be described with simple geometry as in this case (washer, cylinder, torus). WAV can project perspective imagery onto a geometrically “chaotic” environment and make the imagery appear consistent from a specific, optionally tracked and thus dynamic, viewpoint. This capability would be required if one wanted to add head-tracked stereoscopy to the A-Desk; we discuss this option under the Future Work section.

CALIBRATION

Instead of calibrating each projector individually, we used a single Point Grey research stereo camera pair on a tripod, imaging the entire display for calibration. The camera pair was calibrated using Zhang’s camera calibration algorithm[12] with the aid of a physical checkerboard pattern. A series of binary-coded structured light patterns were displayed by each projector in sequence and captured by this camera pair. Detection of the structured light pattern yields a set of image-to-image correspondences between the calibration cameras and each projector. For each such stereo correspondence, a 3-D point is reconstructed on the display surface, yielding a dense 3-D point cloud reconstruction of the display surface. This dense reconstruction was then registered to the known virtual 3-D model of the display surface (based on CAD and physical measurements) using the Iterative Closest Point algorithm.[13]

To obtain a full 3-D calibration for each projector, the set of original features detected in its structured light pattern are then associated with their 3-D reconstructed counterparts that have been registered to the display surface model in the previous step. Each projector's set of 3-D–2-D correspondences is then used to compute its full 4×3 projection matrix and associated Brown lens model distortion parameters, reconstructing the projector relative to the display surface. We then compute a blend map for each projector using previously described approaches.[2,7]

APPLICATIONS

Simulated Desktop. We developed an initial demonstration application (multiwindow simulation) in OpenGL, linked with the WAV framework. The demonstration displays movable, deformable, textured polygonal shapes that look like desktop windows playing videos and displaying scrolling text and moving 3-D graphics by means of simple texture-based animations. There is a dynamic 3-D background with abstract shapes, enabled by the WAV framework. The most highly developed part of the demo application is the fully functional calculation for deforming and moving windows across the 2-D-curved surface; the challenge was to make the windows maintain their apparent rectangular shape even as they were moving across areas of high local two-dimensional curvature, such as the toroidal transition between the flat disk and the cylindrical wall. (Note how both of the early concepts illustrated in Figure 3 avoid that difficulty through the absence of any toroidal display surface elements.)

Intelligence Analysis. IARPA's original applications included simultaneous visualization of multimodal factual and speculative material such as projections or hypotheses, historical and/or real-time data, possibly from mass surveillance (as we now know), database info (cartographics, infrastructure, etc.), highlighting various relationships between such data items. Intelligence applications might also include "running point" in real time during operations like in the "24" TV show or the "Bourne" movies (plausibility aside), where support teams are surrounded by a multitude of information on many displays, as well as other complex scenarios requiring large amounts of simultaneously displayed information. Due to the lack of actual scenario data, we demonstrated the feasibility of such applications only symbolically, by means of a simple 3-D visual simulation showing analysis items as clusters of spheres interconnected with sticks (Figure 2). The rationale was that if OpenGL can be shown to run in a unified fashion (single display context) across such a display, it is conceivable to deploy complex graphical applications of any kind onto such a system.

Classic office multi-document apps. In this rather more conventional area, we envision applications such as document preparation, website design, code development, and others, on many files at once, and adapted for comfortable two-user collaboration. We did not demonstrate these concepts beyond the simulated desktop OS described above.

Processing of physical documents. The A-Desk concept also incorporates streamlined automatic recognition, illumination, scanning, and capture (including optical character recognition as needed) of physical documents, such as books, journals, magazines, or post-it notes, placed on the work surface, a capability that enables seamless transitions from such materials to virtual content. This streamlined interaction paradigm for physical documents is a natural consequence of the unified work-display-lighting-capture surface. Figure 1, right illustrates the recognition and illumination concept, while previous-related work[4] demonstrated almost all components of this pipeline (Figures 8 and 9).

Panoramic imagery. Viewing of panoramic imagery is another natural application of the A-Desk system as it is well-suited to its geometric layout (Figure 10).

Figure 8. View of "multisurface, multiresolution" workspace[4] with high-resolution LCD display panel insert whose imagery is continuous with the one on the projectively illuminated surfaces. Note 3-D model displayed continuously (from the camera's point of view) across the sharp edges between the mutually perpendicular display surfaces.

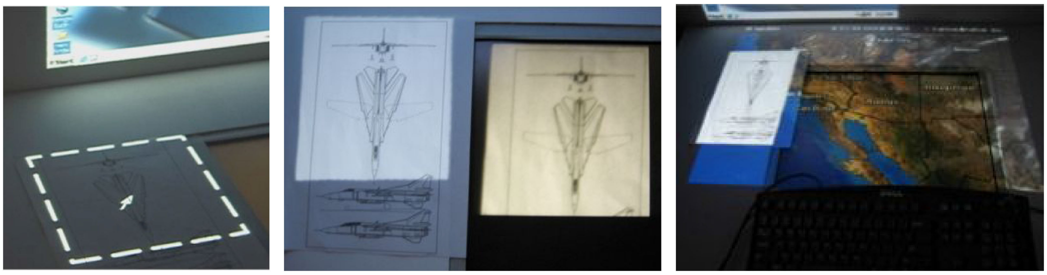

Figure 9. Another view of the "multisurface, multiresolution" workspace[4] experimental system demonstrating illumination and capture of physical document. In addition to lighting the physical object, one can capture it with the camera and then process it digitally.

Figure 10. Panoramic viewing. Note distortion at the bottom, which makes imagery look correct from the user's viewpoint—not the camera's (ca. November 2008).

Telepresence. The user-enveloping nature of the A-Desk makes it suitable for telepresence applications, such as remote control of robotic devices. This application was introduced in follow-up work on animatronic avatars.[5] Figure 11 demonstrates how such a system could operate; as opposed to the static imagery in Figure 10, the live panoramic imagery shown in this demonstration (Figure 11, left) is acquired in real time by the mobile surround-camera system in Figure 11, right.

Hand tracking. To experiment with interaction concepts, camera-tracked finger-mounted attachments worn by the user can be used to manipulate a mouse pointer and move windows on the A-Desk display (Figure 12), or perform two-handed gestures (for example, to scale

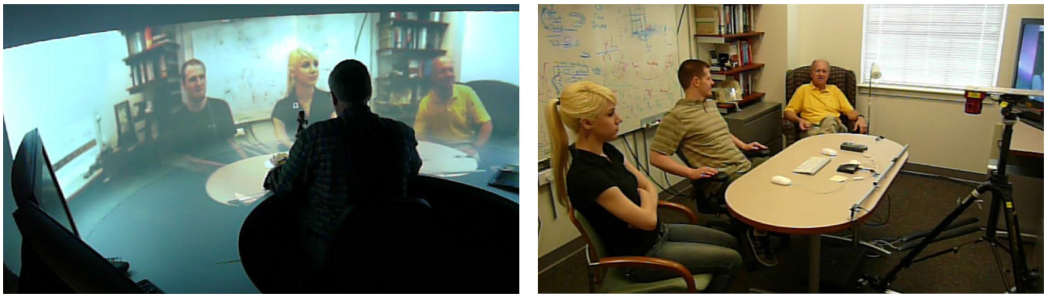

Figure 11. Telerobotic application.[5] Left: A-Desk display for panoramic viewing into a remote environment. Right: Remote "robot" placeholder (a tripod with a multicamera panoramic device, visible at right) acquires the panoramic imagery in real time (early 2009).

Figure 12. Three-dimensional gesture interaction: Moving a window (dark text window with bright dot at center) with a camera-tracked finger attachment (April 2009).

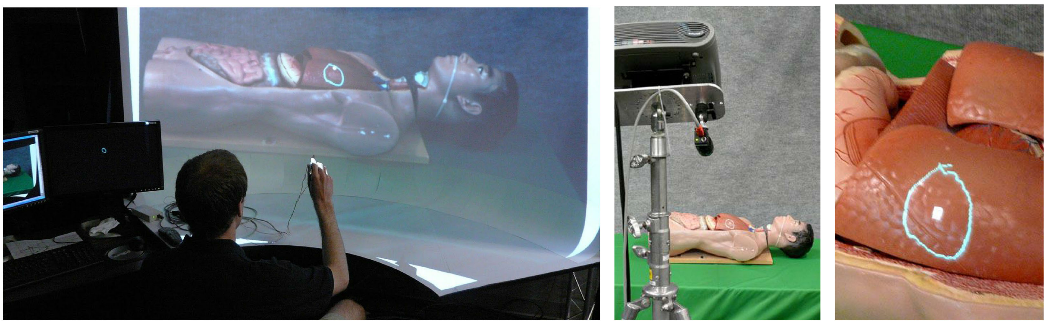

Figure 13. Three-dimensional tele-annotation prototype. The camera-based finger tracker (left) allows annotating the remotely located human model by means of a digital projector illuminating the model (center and right) (July 2009).

virtual imagery). Another application of hand tracking demonstrates freehand tele-annotation (Figure 13).

Customized OS window manager. We initiated an effort to port the Linux Compiz OpenGL-based window manager[15] to the A-Desk, aiming to demonstrate an actual functioning general-purpose computer running arbitrary existing applications that would be unaware of the unusual shape of the display surface. Thus, the applications would be able to run unmodified as all warping and blending would be done within the compositing window manager. In other words, the aim was to provide the complete functionality of a conventional window manager within the A-Desk environment—as opposed to the simulated desktop mentioned at the beginning of this section. Beyond the application to an immersive Linux workspace on arbitrary surfaces, this approach would also support virtual machine

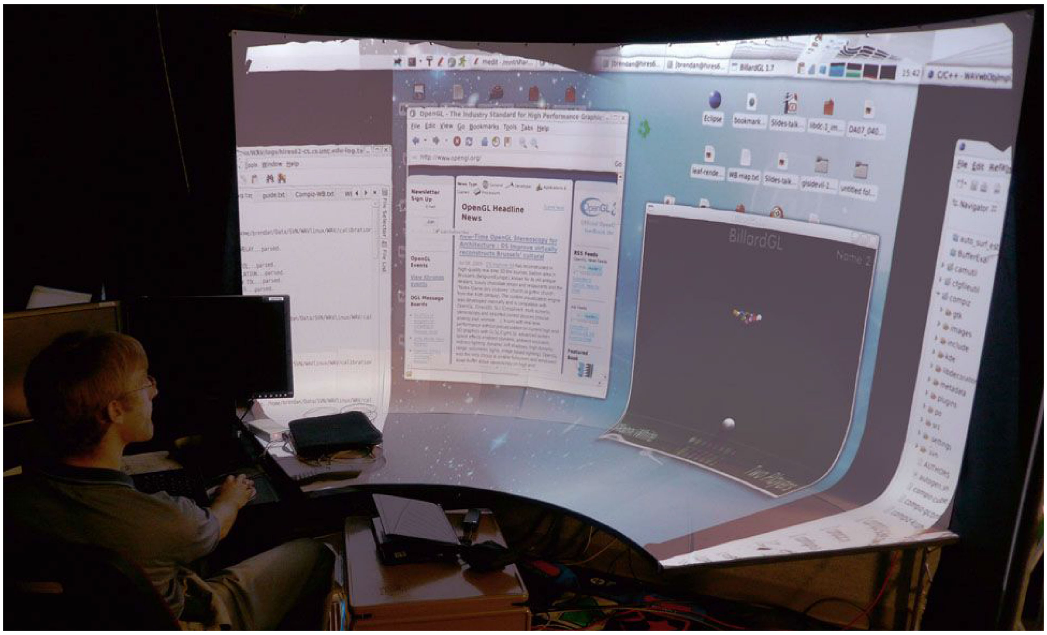

Figure 14. Partially implemented general-purpose Linux desktop environment using the OpenGL-based Compiz window manager (July 2009). Cf. Figure 2.

applications such as Oracle VirtualBox, which would then in turn support operating systems such as Microsoft Windows or Android with many of their native applications, all within the two-dimensionally curved A-Desk display surface.

The Compiz effort was not completed before the termination of the research project in late 2009 but had already shown promising preliminary results (Figure 14). The remaining problems included mapping and distribution of mouse motion and click events and some interfacing issues between Compiz and the WAV framework.

DISCUSSION, CONCLUSION, AND FUTURE WORK

The A-Desk system described here, which we designed, constructed, and briefly used more than ten years ago, was a "version-zero" prototype which we hoped would lead to further developments. Unfortunately, the program under which we implemented this initial design was quickly discontinued. We are particularly disappointed to have missed the chance to complete the Linux Compiz port, which would have elevated the existing prototype from a system made for proof-of-concept development and demonstrations to a practical, everyday work environment. Nevertheless, we managed to make sufficient progress in that direction to realize the importance of combining the general-purpose calibration-warp-blend-display framework (WAV) with a compatible general-purpose operating system and window manager (Linux Compiz) supporting the necessary geometric manipulations. We would like to offer that as the most important lesson learned from this development. In retrospect, we would have benefitted from starting that effort as early as possible (in simulation, before any hardware even became available), as it would also have helped with all our other software efforts.

Since then, projection and display technologies have evolved dramatically, as have computational and rendering power of CPUs and GPUs. Yet so far we have not witnessed the introduction of dual use of surfaces as both physical work areas and virtual display areas. Nor have we seen everyday use of unified display on surfaces of arbitrary shapes and relative orientations. Hence, we believe that the basic concept of the A-Desk remains interesting, in terms of both display geometry and of the capabilities we envisioned for it (and partially implemented). We think it is worthwhile to speculate about how such a system could be reimplemented with today's more advanced technology, as well as how it might compare to alternatives available today, such as commercial head-mounted displays (HMDs). Given that specific alternative, we will dedicate a considerable part of the discussion to the possible addition of stereoscopy to the A-Desk.

Modern reimplementation. First, let us consider a straightforward reimplementation of the same concept using state-of-the-art technology. We would replace all projectors with modern 3840×2160 units. For example, instead of the short-throw NEC WT-610E XGA projectors, we would use a device such as the Sony VPL-VZ1000ES short-throw "4k" projector.[16] Similarly, we would replace the top-down-projecting Projection Design F20 units with modern 4K projectors. We could still drive a large number of such projectors from a single PC. For example, a system based on an AMD Threadripper CPU with its large number of PCI-Express lanes could easily accommodate 4 high-end PCI-E video cards, thereby supporting 15 4k projectors and a conventional service display, which would yield a 270-degree A-Desk if using projection geometry analogous to the one shown in Figure 7. Such an updated system would have an average pixel density of roughly 100 dpi, three times higher than the original prototype and comparable to regular desktop monitors. Since modern projectors support High Dynamic Range (HDR), we could achieve better blending and higher contrast in the blended imagery, but might lose some or all of the HDR bits—depending on factors such as ambient lighting and display surface reflectance characteristics.

Stereoscopic projection. Some modern projectors also happen to support a severely crippled form of frame-sequential stereoscopic display (to be viewed with active shutter glasses): they can be driven with a side-by-side horizontally compressed, or with a top-bottom vertically compressed 1080p 60 Hz signal, whose bandwidth, unfortunately, is an astounding 4 times lower than that of the full-resolution signal, providing roughly only 1 megapixel per eye, which the projector then internally stretches to its native resolution (roughly 8 megapixels) and displays at twice the frame rate, in alternation with the other eye's signal, which is processed in the same way. This is in principle unnecessarily restrictive considering the available bandwidth and processing power within the projector (cf. superior passive stereoscopic panels made by LG, in production until 2016,[17] delivering 3840×1080 per eye at 60 Hz). It should be noted for completeness that many such projectors also support the "frame packing" format generated by 3-D Bluray players, which can deliver 1920×1080 pixels per eye but is restricted (again) to lower frame rates (24, 25, 30 Hz) unsuitable for computer work—not that it matters, because video card manufacturers did not find it necessary to enable their products' HDMI outputs to easily generate that kind of (still inferior) stereoscopic signal for OpenGL applications. It may be possible to generate such a signal by defining custom resolutions, or by inserting specialized devices such as the Lumagen Radiance converter[18] into the display paths, a rather significant complication that only adds yet another option to compromise between resolution and update rate, at up to a mere quarter of the full transmission and display bandwidth. Beyond that, some high-end "professional" video cards support frame-sequential stereoscopic signals such as 1920×1080 at 120 Hz (60 Hz per eye), a slightly more advantageous 50% of the 4k 60-Hz transmission/display bandwidth. Alas, procuring suitably compact 4k projectors that also accommodate that format appears difficult or impossible. We conclude that for our projection-based A-Desk design, it is not possible to obtain both very high monoscopic resolution and stereoscopic resolution that does not leave at least half of the available bandwidth unused.

Notwithstanding the industry's regrettable and generalized current neglect of non-head-worn stereoscopy, it is evident that the modernized A-Desk would be vastly superior to the original one and would provide a spectacular, unified surround display area of on the order of 100 million effective monoscopic pixels; that is more than an order of magnitude higher than the prototype we constructed. With the availability of projective stereoscopy (even crippled as it is), it also makes sense to add tracking of the user's eyes, which would be implemented either with modern camera-based methods (face and eye tracking, no additional headgear), or with precise optical methods, placing passive or active markers, perhaps aided by miniature inertial units, on the shutter glasses. Head-tracked stereo would extend the range of applications, for example into areas such as scientific and medical visualization, making the modernized A-Desk in many ways akin to the CAVE.[6]

Tiled LCD panels. High-resolution display panels have also made advances since 2009, and one could conceivably implement a variation on the original design using curved display monitors, if we accept these monitors' specific radius of curvature (usually around 1.8–2 m, matching the curvature of our larger original A-Desk concepts in Figure 1, left, and Figure 4) as the radius of the entire workspace. Then, such monitors could be used for some or all of the cylindrical portion of the system, at the expense of visible bezels. All other parts of the display surface would still have to be implemented with projectors, especially the non-cylindrical parts, i.e., the toroidal and washer-shaped areas; truly flexible displays have been demonstrated, but they are not yet widely available commercially, nor do they support 2-D curvature or arbitrary nonrectangular pixel arrays. Such a complex hybrid design might be "conceptually impure," but might have its advantages. The monitor-tiled cylindrical portions of the display surface could offer nonresampled native pixels at higher density than the rest of the workspace, HDR, and absence of blending artifacts, in exchange for visible inter-panel borders/bezels. Regarding stereoscopy, a small number of "3D Ready" LCD monitors still on the market are indeed suitable for frame-sequential stereo at rates up to 120 Hz (when driven by "professional" video cards); to avoid crosstalk, they must have fast Twisted Nematic (TN) panels, as opposed to other, more prevalent technologies such as In-plane Switching (IPS). Our extensive research has uncovered just one suitable curved monitor of this type, the ASUS ROG Swift PG27VQ.[19] Finally, a tiled stereoscopic display allows only small virtual protrusions out of the display surface, in order to avoid stereo conflicts with inter-monitor bezel areas; this is a considerable limitation of such a design.

Comparison with virtual reality headsets. Contemporary alternatives for surround imagery include HMDs, which typically offer angular resolutions well below 20 pixels/degree at fields of view on the order of 100 degrees,[20] resulting in angular pixel densities multiple times lower than the angular pixel density available to a user sitting in the center of the modernized A-Desk postulated above (but nearly approaching its angular pixel density in "crippled-stereo" mode). Furthermore, while some manufacturers offer operating-system-like GUI functionality in HMDs, such interactions are still tedious and awkward in virtual environments when compared to classic mouse-based GUIs, or to modern touch interfaces. Because of that, we expect that a modernized A-Desk with a completed Linux window manager port, and with hand and/or finger tracking as already shown, would provide a user experience that surpasses contemporary HMDs for desktop-oriented tasks. HMDs are also inferior in terms of collaborative work, except where multiple users have to navigate a shared virtual environment. The advantages of HMDs are the simultaneous presentation of left-right imagery (as opposed to alternating left-right frames), and superior stereoscopic depth perception, as HMDs are able to present extreme stereoscopic depth without excessive eye strain, thanks to their optics. In contrast, room-scale display surfaces have a more important accommodation-vergence conflict and do not support stereo virtual geometry that extends far in front of the display or deep beyond it; carefully watching "3-D" films on "3-D" HDTVs reveals that all geometry is flattened and contained within a fairly thin spatial layer extending from slightly in front to slightly behind the display panel. Stereoscopic content in motion pictures is created in this way to allow comfortable binocular fusion for almost everyone, by compressing depth such that the stereoscopic disparity on the display surface is small (the display surface is shared by both eyes, in contrast to HMDs). This is easily noticeable when "deep" space plays an important role, as in "Gravity" (2013), where the viewer fuses (triangulates) background stars only a few centimeters behind a normal-size HDTV panel. From these considerations, we can infer that HMDs are superior to an A-Desk system (and to all conventional stereo displays with a single display surface shared for both eyes) when it comes to stereoscopic visualization of deep geometry. However, non-head-worn stereoscopic displays can be considered more suitable for meeting-type telepresence applications such as our (nonstereoscopic) demonstration (Figure 11). They are also at the very least adequate for stereoscopic display of geometry whose depth extent is limited, such as volume renderings. But in all fairness, stereoscopy was not part of the original A-Desk requirements; rather, it is an added benefit of our hypothetical modernization, and as implemented in contemporary projectors, sadly, a modest one.

Augmented Reality. There are a number of commercially available optical-see-through AR devices. Current devices have limited fields of view (on the order of no more than 50 degrees diagonal), due to the complexity of optical combiner designs, and to considerations of weight and user comfort. Hence, it seems unsuitable to attempt to replace the A-Desk system with such a device. However, an AR headset might be used for the display of supplementary elements (such as stereoscopic geometry) together with a projector-based A-Desk. Such a hybrid system might be built with specially finished surfaces that appear quite dark when not projected upon, thus serving as high-contrast backgrounds for the additive compositing in AR headset combiners—consider that the demonstration footage we are usually shown for devices such as the Microsoft HoloLens is acquired by a built-in camera and is enhanced with nearly opaque augmentations, which do not accurately portray the user's view through the headset (it is treated as video see-through, not optical see-through).

Head-mounted projection. Head-mounted projection displays (HMPDs) use beam splitters to align the projector's optical centers with the user's eyes.[21] Because of that, the shape of the display surface does not need to be known to the system since the projected image always looks unchanged from the projector's (and therefore from the user's eye) point of view. This completely eliminates the complex calibration and simplifies the rendering pipeline as one can deploy an arbitrarily shaped display surface (even a randomly draped cloth), as long as the projector can achieve focus on it. However, these considerations apply per eye, and so one *must* use one projector per eye and show separate imagery to each eye, making stereoscopy "mandatory" for arbitrary display surfaces. Separation is achieved by projecting on surfaces with highly directional retroreflective properties; if necessary, polarization may be added to strengthen separation. We are aware of one commercial device using HMPD technology, the *Tilt Five*,[22] which also features built-in head tracking like many modern head-worn displays; its (projective) field of view is 110 degrees and its resolution is specified as 720p, unfortunately yielding an angular resolution that is even lower than the resolutions of current virtual reality headsets discussed above. However, if miniaturized projection at 4k or better resolutions became available, a more advanced HMPD could be a superior alternative to our modernized A-Desk. Advantages such as mobility and arbitrary display surfaces were already mentioned. Other differences include the fact that HMPD users must wear special gear (the current Tilt Five device appears fairly comfortable). Also, HMPD users can easily collaborate in 3-D if each user projects their own individual head-tracked stereoscopic imagery onto shared retroreflective surfaces.

Camera-based capture. The A-Desk concept relies on cameras in order to provide convenient, seamless imaging (scanning) of physical documents. Miniature cameras of extremely high imaging quality and resolution are now ubiquitous in mobile devices. They are also very inexpensive, and we can envision equipping a modernized A-Desk with a sufficient number of such tiny devices, their fields of view covering at least all usable (nonvertical) areas of the work/display surface at extremely high imaging resolution. In a practical system, it may even be necessary and useful to image the entire display surface with such cameras as they may assist in recalibrating the display for constant performance, thus compensating for mechanical deformations, as well as for color changes caused by ambient illumination and aging projector bulbs. The same set of cameras may also help track all user motion and interaction in the vicinity of the display surfaces, implementing multitouch by taking advantage of the known geometry of the display surface (it would admittedly be challenging to determine exact contact events using this method, and nearly impossible to obtain a measure of contact pressure). Furthermore, multiple cameras may enable digital "flattening" of warped or degraded documents,[23] as well as full reconstruction of 3-D objects.

To conclude, we appreciate the opportunity to present the evolution of ideas that led to our rather unusual A-Desk work/display environment concept and to discuss why we believe in its possible relevance and applicability to future designs.

ACKNOWLEDGMENTS

We are grateful to John E. Thomas (collaboration on physical construction and projector mounting), Banu Kutlu (collaboration on software development and testing), Peter Lincoln (collaboration on telepresence prototype), and Michael B. Gilbert at the NGA (collaboration, input). Funding: Future Analyst Workspace (A-Desk), Principal Investigators Henry Fuchs and Greg Welch, U.S. Air Force Office of Scientific Research, Agency Number FA8750-08-2-0209, IARPA Analyst Workspace for Exploitation (A-SpaceX) Program, Program Manager Dr. Jeffrey G. Morrison.

REFERENCES

1. [Online]. Available: https://www.iarpa.gov/. Accessed on: Sep. 15, 2019.
2. R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs, "The office of the future: A unified approach to image-based modeling and spatially immersive displays," in *Proc. 25th Annu. Conf. Comput. Graph. Interact. Techn.*, 1998, pp. 179–188.
3. G. Bishop and G. Welch, "Working in the office of 'real soon now'," *IEEE Comput. Graph. Appl.*, vol. 20, no. 4, pp. 76–78, Jul./Aug. 2000.
4. T. Johnson and H. Fuchs, "A unified multi-surface, multi-resolution workspace with camera-based scanning and projector-based illumination," in *Proc. Eurograph. Symp. Virtual Environ./ Immersive Projections Technol. Workshop*, Jul. 2007. [Online]. Available: https://www.uni-weimar.de/medien/vr/ipt-egve/IPT-EGVE_2007
5. P. Lincoln, G. Welch, A. Nashel, A. State, A. Ilie, and H. Fuchs, "Animatronic shader lamps avatars," *Virtual Reality*, vol. 15, pp. 225–238, 2011.
6. C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti, "Surround-screen projection-based virtual reality: The design and implementation of the CAVE," in *Proc. 20th Annu. Conf. Comput. Graph. Interact. Techn.*, vol. 27, pp. 135–142, Aug. 1993.
7. A. Majumder and R. Stevens, "LAM: luminance attenuation map for photometric uniformity in projection based displays," in *Proc. ACM Symp. Virtual Reality Softw. Technol.*, Nov. 2002, pp. 135–142.
8. T. Johnson, G. Welch, H. Fuchs, E. La Force, and H. Towles, "A distributed cooperative framework for continuous multi-projector pose estimation," in *Proc. IEEE Virtual Reality Conf.*, Mar. 2009, pp. 35–42.
9. T. Johnson, "A cooperative approach to continuous calibration in multi-projector displays," Doctoral dissertation, Univ. North Carolina at Chapel Hill, Chapel Hill, NC, USA, Dec. 2009.
10. *Working Through Synthetic Worlds*. C. A. P. Smith, J. G. Morrison, and K. W. Kisiel (Eds). Burlington, VT, USA: Ashgate Pub. Company, 2009.
11. A. State, G. Welch, and A. Ilie, "An interactive camera placement and visibility simulator for image-based VR applications," in *Proc. Eng. Reality Virtual Reality (3D Imag., Interaction, Meas.; IS&T/SPIE 18th Annu. Symp. Electron. Imag. Sci. Technol.)*, vol. 6055, pp. 640–651, Jan. 2006.
12. Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
13. Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image Vision Comput.*, vol. 10, no. 3, pp. 145–155.
14. R. Raskar *et al.*, "Multi-projector displays using camera-based registration," in *Proc. Conf. Vis.*, 1999, pp. 161–168.
15. [Online]. Available: https://en.wikipedia.org/wiki/Compiz. Accessed on: Sep. 16, 2019.
16. [Online]. Available: https://www.projectorcentral.com/pdf/projector_manual_9773.pdf. Accessed on: Sep. 15, 2019.
17. [Online]. Available: https://www.forbes.com/sites/johnarcher/2017/01/19/angry-3d-fans-start-petition-against-lg-oled-tvs. Accessed on: Sep. 15, 2019.
18. [Online]. Available: http://www.lumagen.com/docs/Radiance_Pro_Manual_071819.pdf. Accessed on: Sep. 17, 2019.

19. [Online]. Available: https://www.asus.com/us/Monitors/ROG-SWIFT-PG27VQ/specifications/. Accessed on: Sep. 20, 2019.
20. [Online]. Available: https://en.wikipedia.org/wiki/Comparison_of_virtual_reality_headsets. Accessed on: Sep. 20, 2019.
21. J. P. Rolland, F. Biocca, F. Hamza-Lup, Y. Ha, and R. Martins, "Development of head-mounted projection displays for distributed, collaborative, augmented reality applications," *Presence Teleoperators Virtual Environ.*, vol. 14, no. 5, pp. 528–549, Oct. 2005.
22. [Online]. Available: https://www.tiltfive.com/. Accessed on: Oct. 23, 2019.
23. M. S. Brown, M. Sun, R. Yang, L. Yun, and W. B. Seales, "Restoring 2D content from distorted documents," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 11, pp. 1904–1916, Nov. 2007.

Andrei State is a Senior Research Scientist with the UNC Department of Computer Science and Cofounder of InnerOptic Technology Inc., which creates medical 3-D guidance systems. His research interests include stereoscopy, virtual/augmented reality, and display technology. He received a Dipl.-Ing. aer. degree from the University of Stuttgart, Stuttgart, Germany, and a M.S. degree in computer science from the University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. Contact him at andrei@cs.unc.edu.

Herman Towles is a Senior Systems Engineer and Distinguished Member of Technical Staff at Applied Research Associates. His research interests include augmented reality, 3-D pose tracking, computer vision, and image sensing. He received B.S. and M.S. degrees in electrical engineering from the University of Memphis, Memphis, TN, USA. Contact him at herman.towles@gmail.com.

Tyler Johnson is a Principal Engineer with Scalable Display Technologies, Inc., Cambridge, MA, USA. He received a Ph.D. degree in computer science from the University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, in 2009 for his dissertation work demonstrating continuous calibration of projection-based displays using intelligent camera-projector units. Since joining Scalable Display Technologies in 2010, he has become an industry expert in the calibration of multiprojector displays and has contributed to several patents in the field. Contact him at tyler@scalabledisplay.com.

Ryan Schubert is the Special Projects Coordinator and Lab Manager with the Synthetic Reality Lab (SREAL), the University of Central Florida, Orlando, FL, USA. His research interests include virtual and augmented reality, interactive computer graphics, and human–computer interaction and interfaces. He received a M.S. degree in computer science from the University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. Contact him at rschuber@ist.ucf.edu.

Brendan Walters is a Senior Backend Engineer with Pendo.io, Inc., where he works primarily on fault-tolerant high-throughput data processing. His research interests include using machine learning for medical image analysis, and extracting insights from user behavior in consumer and business applications. Contact him at brendan@ithildinconsulting.com.

Greg Welch is a Professor and the AdventHealth Endowed Chair in Healthcare Simulation with the University of Central Florida, Orlando, FL, USA. As a Computer Scientist and Engineer, he has appointments in the College of Nursing, the College of Engineering and Computer Science, and in the Institute for Simulation & Training. His research interests include human–computer interaction, human motion tracking, virtual and augmented reality, computer graphics and vision, and training related applications. He was the recipient of the the IEEE Virtual Reality Technical Achievement Award in 2018 (VR 2018), and the Long Lasting Impact Paper Award at the 15th IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2016). Contact him at welch@ucf.edu.

Henry Fuchs is the Federico Gil Distinguished Professor and Computer Science and Adjunct Professor of Biomedical Engineering at the University of North Carolina at Chapel Hill. He has been active in computer graphics since the early 1970s, with rendering algorithms (BSP Trees), graphics engines (Pixel-Planes), virtual environments, tele-immersion systems and medical applications. He is a member of the National Academy of Engineering and a Fellow of the American Academy of Arts and Sciences. He was the recipient of the 2015 ACM SIGGRAPH Steven Anson Coons Award, and in 2018, an honorary doctorate from TU Wien. Contact him at fuchs@unc.edu.

Contact department editor Kyle Johnsen at kjohnsen@uga.edu or department editor Christian Sandor at chris.sandor@gmail.com.