This paper is © IEEE and appeared reformatted in IEEE CG&A 35(1): 33-45, 2015, DOI: 10.1109/MCG.2014.80

The Reality Deck - Immersive Gigapixel Display

Charilaos Papadopoulos, *Student Member, IEEE,* Kaloian Petkov, *Member, IEEE,* Arie E. Kaufman, *Fellow, IEEE* and Klaus Mueller, *Senior Member, IEEE*

Abstract

We designed and built a next-generation visualization facility, the Reality Deck, that simultaneously offers state-of-the-art aggregate resolution and immersion. The *Reality Deck* is a 1.5 gigapixel immersive tiled display with a full $360°$ horizontal field of view. Comprised of 416 high density LCD displays, it is built to tackle today's big data problems while providing users with $\frac{20}{20}$ visual acuity for the majority of the visualization space. In this article, we discuss the motivations, design principles and engineering challenges behind the Reality Deck. Additionally, we showcase several techniques that were developed on the facility, focused on enabling natural exploration and supporting the visual analysis of big data.

Author Contact Details:

- Charilaos Papadopoulos

  Affiliation: Dept. of Computer Science, Stony Brook University, NY, USA
  Address: Dept. of Computer Science, Stony Brook University, Stony Brook, NY, 11794-4400
  Tel: +1 (917)859-0527
  Email: *cpapadopoulo@cs.stonybrook.edu*
- Kaloian Petkov

  Affiliation: Dept. of Computer Science, Stony Brook University, NY, USA
  Address: Dept. of Computer Science, Stony Brook University, Stony Brook, NY, 11794-4400
  Tel: +1 (631)216-7142
  Email: *kpetkov@cs.stonybrook.edu*
- Arie E. Kaufman

  Affiliation: Dept. of Computer Science, Stony Brook University, NY, USA
  Address: Dept. of Computer Science, Stony Brook University, Stony Brook, NY, 11794-4400
  Tel: +1 (631)632-8428
  Fax: +1 (631)632-8445
  Email: *ari@cs.stonybrook.edu*
- Klaus Mueller

  Affiliation: Dept. of Computer Science, SUNY Korea, Incheon, Korea and Dept. of Computer Science, Stony Brook University, NY, USA
  Address: Dept. of Computer Science, SUNY Korea, 187 Songdo-Dong, Yeonsu-Gu, Incheon City, Korea
  Tel: +82 32.626.1200
  Email: *mueller@cs.stonybrook.edu*

C. Papadopoulos, K. Petkov, A. E. Kaufman and K. Mueller are with the Department of Computer Science, Stony Brook University, Stony Brook, NY.

The Reality Deck - Immersive Gigapixel Display

INTRODUCTION

Fig. 1: Synthetic, to-scale, view of the immersive gigapixel Reality Deck facility displaying a geometric model of "future" New York City (approximately 40 million triangles with hundreds of materials).

A vast number of data sources, such as supercomputers and sensors, have become "fire hoses", generating information at a far greater rate than it can be digested. For example, the AWARE-2 camera system can capture a 1.47 gigapixel photograph [1]. In cosmology, the Large Synoptic Survey Telescope (www.lsst.org), will feature a 3.2 gigapixel sensor and capture approximately $200,000$ images annually.

One approach towards visualization is to fit vast quantities of data on a single display. Various summarization, abstraction and focus + context techniques aim at accomplishing that, while providing users with the overall patterns and structure of the data.

Maximizing available screen real estate demonstrably affects the visualization process [2]. Tiled display arrays (or powerwalls) are a realization of this concept, offering a large, high-resolution, collaborative workspace. For applications that benefit from "physical immersion" (e.g., surrounding the user with visuals), CAVEs [3] are a suitable visualization setting, potentially offering a fully immersive field of view (FoV) and stereoscopic 3D. Immersive Virtual Environments (IVEs), however, present a fundamental dichotomy. When large-scale datasets are being visualized, total screen real estate and resolution must be maximized, making powerwalls more appropriate. However, powerwalls lack immersion and their resolution can be limited by physical constraints (e.g., a $100'$ planar powerwall is unwieldy to construct and utilize). Conversely, CAVEs offer immersion but their maximum resolution is roughly 100 megapixels per eye for state-of-the-art systems. Furthermore, their total workspace can be somewhat constrained (approximately $9' \times 9'$ for typical setups).

Motivated by the current landscape, we designed and built the Reality Deck, the world's first gigapixel resolution display in a fully enclosed setting. A synthetic, to-scale, rendering of the facility is depicted in Fig. 1, while various applications are shown in Figs. 2 and 4-9. The Reality Deck offers more than 1.5 gigapixels worth of resolution, a $360°$ horizontal FoV and a workspace of approximately $33' \times 19' \times 11'$, allowing multiple users to naturally explore data at different scales by approaching or distancing themselves from the displays while maintaining the panoramic context. This display real estate can be abstracted in planar or immersive configurations. In this article, we examine the motivation and design process behind the Reality Deck. We expose findings and benchmarks that arose during the engineering process. Finally, we present a number of novel techniques that evolved while utilizing a gigapixel resolution display.

IMMERSIVE VIRTUAL ENVIRONMENTS

The design of IVEs is largely driven by two main factors; the visual acuity and the degree of immersion. Visual acuity is a metric used to quantify the quality of the visuals that a display can deliver. Maximizing visual acuity allows users to approach the display surfaces and naturally perform multi-scale exploration (rather than perceive the resolution limits of the display technology). For details on visual acuity, see the sidebar.

Immersion is defined as the degree of suspension of disbelief that a visualization system provides to the observer. Cruz-Neira et al. [3] define five key factors for maximizing immersion: *view-centered perspective* (head tracking), *panorama* (surrounding the viewer with visuals), *body and physical representation* (user's awareness of the physical constraints of the interactive workspace), *intrusion* (restricting the user's senses) and *field of view* (display portion that the user can observe without rotating her head). The CAVE's popularity is a testament to its ability to optimally combine these five factors. While maximizing visual acuity allows for natural "zooming" through the data, maximizing physical immersion enables a wider range of "panning" by looking around.

CAVEs are routinely used for data visualization in a number of scenarios but the concept is arguably not keeping up with the growth in dataset sizes. For example, the CORNEA (kvl.kaust.edu.sa/Pages/CORNEA.aspx), a high-end CAVE installation, offers a total of approximately 100 megapixels per eye. Yet, modest examples of panoramic images from www.gigapan.org can exceed 1 gigapixel. Another consideration for current immersive facilities is their ability to deliver quality visuals as a factor of the observer′s position within the visualization space. As shown in Fig. 3, the CORNEA provides a visual acuity metric of approximately $\frac{20}{34}$ but only at a "sweet spot" in its center. For a single user, moving away from this sweet spot can occur naturally in a head-tracked 3D application.

Fig. 2: Interior view of the Reality Deck, showing the front wall and partial side walls. A user is exploring a fused GIS data set of downstate New York, which incorporates elevation information, road networks and detailed 3D geometry for select areas. Insets illustrate the detail that can be resolved by approaching the facility walls, with building facades and architectural details easily distinguishable.

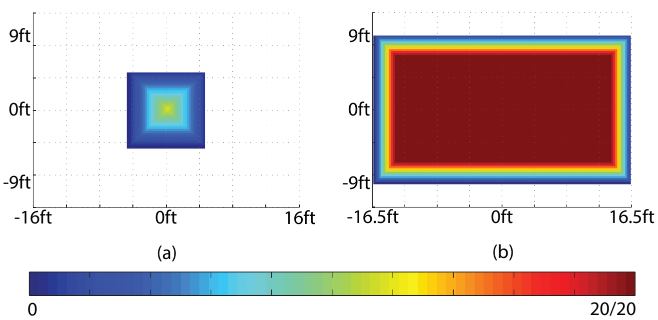

Fig. 3: Visual acuity heatmap (floor plan view) for the CORNEA (a). The facility offers $\frac{20}{34}$ visual acuity in a $10' \times 10'$ workspace. However, the maximum visual quality is achievable only at the very center of the facility. In contrast, the Reality Deck (b) offers a large workspace of $33' \times 19'$. Given its monitor configuration, $\frac{20}{20}$ visual acuity can be achieved approximately $31''$ from the displays. This translates to a total workspace area with $\frac{20}{20}$ visual acuity of about $384'^2$.

However, this movement results in the user approaching display surfaces that already do not saturate her visual system, creating a suboptimal visual experience. Even worse, in multi-user scenarios, only one user can occupy the location where visual acuity is maximized, forcing lower quality visuals on other collaborators.

Conversely, powerwalls are targeted at visualizing high resolution data. High end systems, such as the Stallion Powerwall (tacc.utexas.edu/resources), can reach 300 megapixels of aggregate resolution. However, these facilities materialize as a single planar surface, resulting in two issues. First, they offer only a small degree of physical immersion to the user (whose FoV is saturated by the large display, but without any panorama). Additionally, planar designs can be unwieldy to scale due to spatial constraints and potential ergonomic issues of traversing large distances during visual exploration. Finally, hybrid designs contemporary to the Reality Deck, such as the CAVE2 [4] bridge the gap between an immersive CAVE and a powerwall, with a larger emphasis on CAVE-like characteristics, such as stereoscopic 3D.

"Immersifying" a Tiled Display Wall

As a higher-level goal, we felt that the Reality Deck should provide the high pixel density of tiled displays but also the full FoV of CAVEs. As a next generation facility, it should offer a significantly leap in aggregate resolution (with the gigapixel milestone being an obvious choice). Additionally, $\frac{20}{20}$ visual acuity should be available for the majority of a large visualization space, to promote physical navigation. Two further constraints imposed that the facility fit within the available $40'$ by $30'$ lab space and the budget of $\$1,000,000$.

Our Reality Deck defines an enclosed space, surrounded by high pixel density displays. The arrangement of the displays presented an open design problem. After considering different placements of display surfaces, we opted for a rectangular arrangement with four walls, spanning approximately $33' \times 19'$. This layout enables interesting usage scenarios, depending on the nature of the data and collaborative situation. It is different to most CAVE systems that use a cube-like arrangement of displays. Our rectangular layout allows for more flexibility in operation and also maximizes usage of available lab space.

The four walls of the Reality Deck serve as configurable viewports into visualizations. The straightforward mapping allows the four walls of the system to serve as a 4-viewport configuration into the virtual world, akin to a CAVE. Alternatively, the display can be interpreted as a single continuous

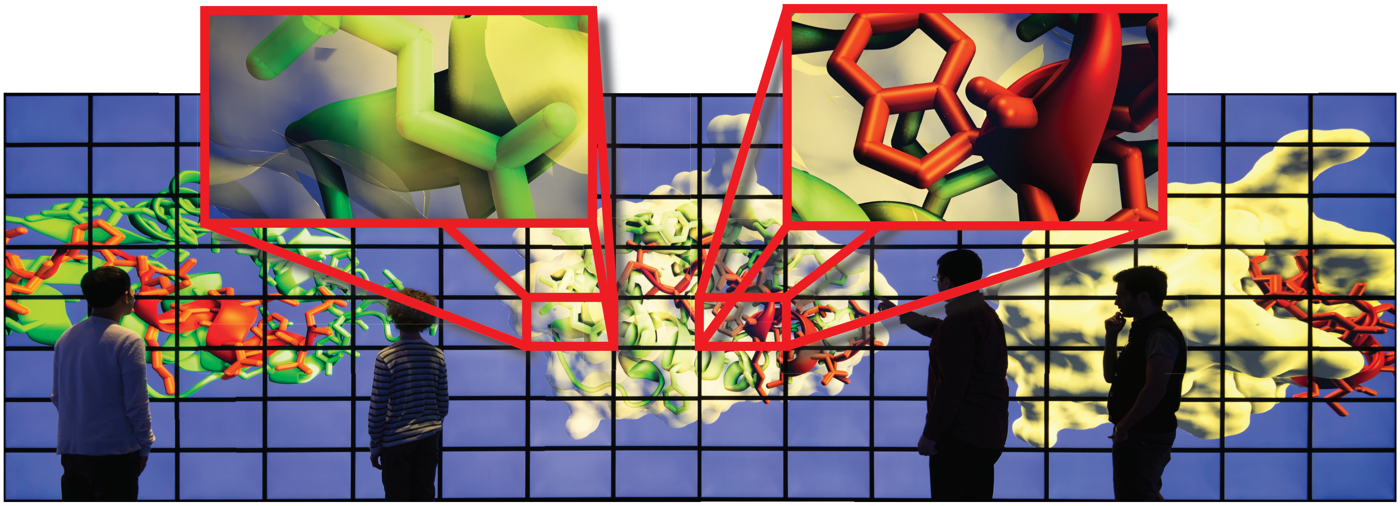

Fig. 4: Visualization of the MDM2 oncoprotein binding to the p53 tumor suppressor. In this example, users are simultaneously exploring different visualization modalities that bring out various features of the data-set. The insets show zoomed in crops of the visualization, highlighting the effects of ambient occlusion which, in combination with alpha-blending, provides additional depth queues to the users.

planar viewport. This configuration is useful for large scale two-dimensional data (e.g., GIS, parallel coordinates). A visual discontinuity exists at the point where the two extremes of the logical frustum meet on the physical display surface (typically in the middle of the rear wall). We employ our *Infinite Canvas* technology (described later and in detail in [5]) to naturally ameliorate this shortcoming - this "continuous viewport" display mapping addresses scalability problems of traditional, planar, powerwalls.

In multi-user scenarios, the facility display space can be subdivided in various ways. First, it can be split into 4 planar tiled-displays, one per "wall". Here, each of the two long walls offers approximately 471 megapixels of resolution while each narrow side wall has 295 megapixels. For additional immersion, we can create two CAVE-like systems, with three walls each, operating independently. Each one of these CAVEs offers roughly 0.76 gigapixels worth of resolution, depending on the "border" space that separates the viewports of the two CAVEs. This pixel count is several times larger than that offered by state of the art CAVE systems, at the expense of bezels and anaglyph-only stereo (due to the selected panel technology).

Building an Immersive Gigapixel Display

Display Selection and Customization

Arguably the most critical component of a visualization environment is the display subsystem. CAVEs are usually based on projectors while tiled display arrays are constructed using both projectors and LCD monitors.

The main benefit of projectors is that they can create a nearly seamless image. On the other hand, projectors require regular maintenance. Based on our experience with our 5-sided Immersive Cabin [6], maintenance must always be followed by a manual recalibration of the system, which can be time consuming for the 10 projector CAVE-like Immersive Cabin and unmanageable for a system that utilizes hundreds of projectors. Additionally, projectors produce significant amounts of heat and noise and the space requirements are affected by the need to accommodate the throw distance (as much as $1.5'$ for short-throw lenses). Finally, projectors are generally much more expensive than LCDs to acquire and maintain. Due to these drawbacks, projectors were eliminated early in our design process.

We then considered different types of LCD monitors based on the following criteria:

Resolution target: Based on the available space and super-gigapixel resolution target, the monitors should provide approximately 100 PPI.

Bezel size: Ideally smaller than 5mm, however the size should not exceed 8mm for a 23" display and 15mm for a 30" display. These metrics were based on bezel dimensions of commercially available monitors with potential structural modifications.

Display size: Larger monitors are preferable as long as they can deliver the required pixel density.

Image quality: The monitors should use high quality panels with good contrast, backlight uniformity and viewing angles.

Stereo support: Stereo is a very desirable feature, but not at the cost of image quality or significantly reduced pixel density.

Based on these characteristics, we evaluated a number of displays with panel sizes from 23" to 60", IPS, PVA and PLS panel technologies and various bezel sizes. We also considered secondary factors, such as power consumption and weight, which affect the requirements for the mounting, power and cooling infrastructure. The different tiled display designs were simulated in our Immersive Cabin (see Fig. 5) in order to

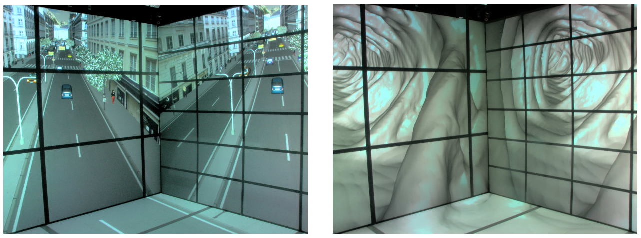

Fig. 5: We evaluated the perceptual effects of the screen and bezel size for different LCD monitors in the Immersive Cabin. Left wall in both images simulates a large-format Planar monitor (60” diagonal); right wall is the Samsung S27A850D with modified bezels (27” diagonal).

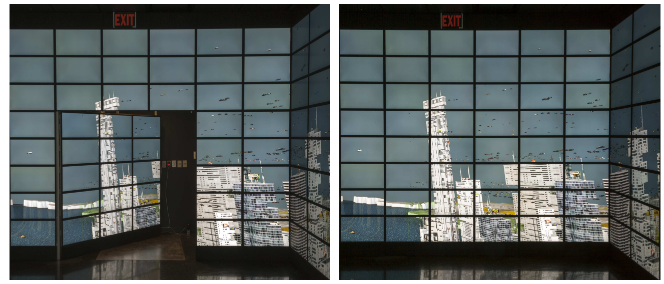

(a) Door Open (b) Door Closed

Fig. 6: (a) A 3×5 section of displays is mounted on an aluminum subframe and rotates around a spring-loaded hinge. The mechanized door can be operated from within the Reality Deck. (b) When closed, the door blends into the rest of the structure, showing a futuristic view of New York City.

analyze the perceptual effects of the bezels on different visualization tasks, including medical and architectural visualization. The results of an informal user study were then used as a contributing factor in the display selection process.

We considered several offerings from a variety of vendors, including ultra-narrow bezel LCDs and monitors with stereoscopic 3D support. We outline the reasoning behind rejecting these offerings in the Discussion and Lessons Learned section later on in this article. At the time of construction no commercially available display satisfied all five of these criteria, however the Samsung S27A850D provided a good balance. It is a professional 27” PLS panel with 2560×1440 resolution and excellent contrast, color saturation and viewing angles. In contrast to CCFL monitors, the S27A850D uses LED backlighting, which significantly reduces the weight and power requirements (46W for the S27A850D versus 134W for a Dell U2711). Finally, while the original bezel is relatively large, the monitors were easily modified with a custom mount that reduces the bezel to 14mm.

Given the available physical space, we arranged the monitors in four orthogonal surfaces. The front and back walls are 16 displays wide while the left and right span 10 displays. All four walls are 8 displays tall for a total of 416 tiled monitors.

The mass-produced nature of commercial monitors entails certain variation in image quality, even for products from the same batch. We evaluated every monitor before modification, looking primarily at image uniformity when displaying a full white and a full black signal, as well as identifying issues with color reproduction. Three photographs were taken of each monitor from a fixed camera position and with a standardized set of camera and display settings. Based on the stacks of images a total of 98 from the first batch of 441 monitors were replaced. After testing the second batch, we selected the best 416 displays, as well as a set of spares, for modification and use in the Reality Deck.

Using lightweight monitors allowed us to design custom mounting brackets and a simple aluminum frame so that individual monitors can be aligned with sub-millimeter accuracy (confirmed via laser leveling) and can also be replaced by a single person. The plastic cover of the S27A850D houses the circuit board, user controls and power supply. These components have been moved to the rear bracket, resulting in a uniform black frame around the display with no visual distractions. The door to the facility is a section of the frame that is mounted on a hinge and holds a 3×5 grid of monitors. It is power operated but can also be opened manually in case of an emergency. When closed, the door rests completely flush with the rest of the wall and it is visually indistinguishable from the other displays (Fig. 6). The displays are offset from the floor of the facility by approximately 7” to allow the installation of tracking cameras and sound speakers.

The facility provides a visualization space of approximately $33' \times 19' \times 11'$ ($W \times D \times H$). Fig. 3 shows a heatmap of the visual acuity within our facility, illustrating the 384 sq. ft. space in which the system achieves $\frac{20}{20}$ or better visual acuity.

Visualization Cluster and Peripherals

The vast number of displays presented a challenge when designing a cost-efficient high-performance visualization cluster. We evaluated a number of different configurations, at various GPU and display per node densities. Our final setup consists of 18 ExxactCorp nodes, with dual hexcore Intel Xeon E5645 CPUs. Each node contains four AMD FirePro V9800 GPUs. The head node is a similarly configured machine with a single GPU. The majority of the cluster nodes drive 24 displays, six per GPU, in a 3×2 monitor grid. The displays of two render nodes in the front-right and back-left corners of our facility operate in groups of 1×4 to ensure that no display group “straddles” the corners of the facility, which would necessitate two rendering passes when the facility operates in “immersive” mode. Each display group is abstracted as a single framebuffer using AMD Eyefinity functionality, simplifying software development and improving performance. The cluster is located in a machine room adjacent to the facility and connects to the displays using Gefen DisplayPort fiber optic extenders. All nodes are interconnected via Ethernet and 40 Gbps InfiniBand networks. A rough total of seven miles of cables were utilized in the facility.

The Reality Deck is also equipped with a 24-camera tracking system from OptiTrack, based on the S250e IR camera. A number of research techniques, described later in the paper, utilize this tracking system for both user interaction and performance optimization. Additionally, we have deployed a 24.4 surround sound system with Genelec 6010A speakers and JBL LSR4312SP subwoofers. The total material cost of the facility, including spare monitors and a "hot-swappable" spare visualization node was approximately $\$950,000$.

Visualization Software and Performance Benchmarks

We have created two visualization frameworks for driving the Reality Deck, using a number of third party libraries (e.g., NVIDIA SceniX, Equalizer [7]) and custom extensions for interaction, tracking and out-of-core rendering. Since each node in our cluster is driving 4 Eyefinity display groups, a total of 72 instances of the rendering application are running at the same time. Thus, we distribute input and per-frame variable data (such as physics-enabled scene graph object state) to each node and the visualization is handled locally, in contrast to systems based on OpenGL command-stream distribution (e.g., Chromium).

In terms of GPU performance, the most critical aspect is having sufficient GPU memory for the large framebuffers required when driving a 6 display Eyefinity group with a resolution of 7680×2880 from a single GPU. In such situations the OpenGL buffers and a couple of multisampled fullscreen render targets can require 2GB or more of memory. Therefore, the minimum we considered to be suitable is 4GB per GPU.

We evaluated the low-level rendering performance with the SPECviewperf 11 benchmark. The observed performance drop when moving from a single WQXGA display to the full 6-display Eyefinity group was between 3% and 20%, depending on the application. Moving to 8x multisampling resulted in an additional 24% to 40% performance loss, while the visualization remained interactive even at the increased image quality and resolution. For shader-bound workloads, performance scales linearly as we move from one to six displays per GPU. Given the high resolution per GPU, some of the traditional rendering algorithms (e.g., volume rendering and GPU raytracing) need to be redesigned for interactive performance in the Reality Deck. Others, such as virtual/gigapixel texturing, can still run at 60Hz provided that there is sufficient memory, memory transfers are managed asynchronously, and no fullscreen render targets or compositing at the OS level are used. We expect that newer generations of GPUs will lift some of these performance restrictions.

When dealing with high resolution data, a major bottleneck during rendering is the fact that the entirety of the data does not fit into GPU memory. Our visualization pipelines support real-time out-of-core texturing, similar to Sparse Virtual Texturing [8], by decomposing data into fixed size tiles which are loaded from external memory (or a network share) based on the results of a visibility determination pass. For a virtual texture comprised of approximately $60,384$ tiles, precompressed to the DXT1 format at a resolution of 520×520, we can achieve approximately 730 MPixels/sec of total CPU-GPU streaming bandwidth, to each GPU in the system, effectively providing 33 fps for bandwidth-bound applications. Alternatively, we can choose to limit the number of tiles uploaded per frame to achieve higher frame rates and our research on acuity-driven gigapixel visualization (described below) can be used to further optimize data transfers, based on the user's location within the Reality Deck.

SUPPORTING TECHNIQUES

The Reality Deck is the only display to offer more than a billion pixels in a horizontally immersive setting and a large workspace that encourages physical navigation. The continuous display surface enables unique data exploration techniques, such as the *Infinite Canvas*. Meanwhile, the gigapixel resolution presents unique challenges in the rendering of different types of data. Our *acuity-driven gigapixel visualization* framework accounts for the distance between the user and the display when making streaming choices. We have also developed a novel *frameless visualization* to enable high resolution volume rendering at interactive frame rates.

The Infinite Canvas

The Infinite Canvas is a novel physical navigation technique targeted at the physical exploration of high resolution data that extends arbitrarily along a single dimension. We begin by placing the virtual camera within a closed, curvilinear, surface with dimensions that approximate the floor aspect ratio of the facility. Imagery can then be mapped on this geometric surface. The Infinite Canvas interactively manipulates the mapping based on the user's position and orientation within the Reality Deck. Through this manipulation, the user is presented with the illusion of a surface that extends arbitrarily along one dimension, as the wrap-around discontinuity is kept outside their field of view.

More formally, using the tracking system, we obtain the position $\mathbf{p}$ and two-dimensional orientation $\mathbf{d}$ of the user, as well as the aggregate rotation σ with respect to a fixed reference vector. Additionally, we intersect the vector $\mathbf{p} - \mathbf{d}$ with the geometry of the canvas in order to obtain $\mathbf{p_{back}}$, the point that lies directly behind the user on the canvas surface. We then calculate an angular offset σ_{start}:

$$\sigma_{start} = 2\pi \lfloor \frac{\sigma}{2\pi} \rfloor - \mathrm{atan2}(\mathbf{p_{back_x}} - \mathbf{p_x}, \mathbf{p_{back_y}} - \mathbf{p_y})$$

If the data to be visualized spans a $[0, 1]$ range on the horizontal axis, we can extract a region $\alpha \cdot [\sigma_{start}/2\pi, \sigma_{start}/2\pi + 1)$, with α being a scale factor equal to the ratio between the circumference of the mapped geometry and the total width of the canvas (e.g., if the data is 3 times larger than the immersive display, $\alpha = 1/3$). This section is then mapped onto the geometry, starting and ending at $\mathbf{p_{back}}$, placing the mapping discontinuity behind the user, as illustrated in Fig. 7. This mapping is computed inside a GLSL shader at runtime and it is used to sample an out-of-core texture. For additional implementation information, we refer the reader to [5].

We have found that the Infinite Canvas is a technique that complements well the traditional approach to data exploration within the Reality Deck (walking along the displays, looking

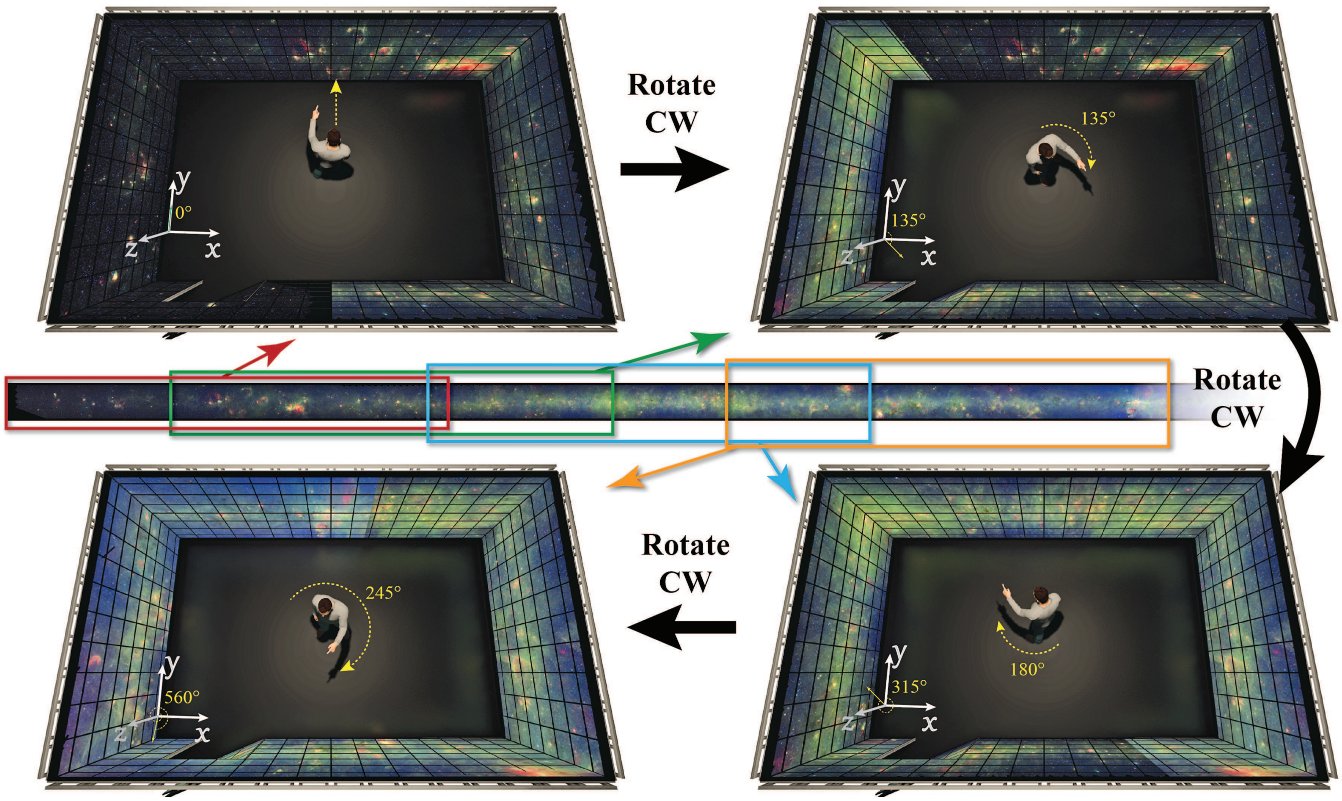

Fig. 7: Synthetic view of the Reality Deck illustrating the operation of the Infinite Canvas technique for the exploration of the GLIMPSE/MIPSGAL survey of the Milky Way galaxy. The user starts by looking at the front wall then proceeds to rotate clockwise while the aggregate rotation σ is used to calculate the new data mapping. After a $> 360°$ aggregate rotation (bottom-left subfigure), the user is presented with a new view into the data. All wall images are captured directly from the Reality Deck and show that the remapping discontinuity is always behind the user.

for points of interest). Indeed, when the Infinite Canvas is not utilized, users would eventually encounter the 2D frustum discontinuity at the rear wall of the facility and then utilize a controller to further "scroll" the data along its major dimension. The Canvas, while still permitting users to skip through sections of the data with an external input device, allows them a greater degree of "mental immersion" by obviating the need for a context-switch between manipulating the visualization virtually and navigating physically by walking.

Acuity-Driven Gigapixel Visualization

During the exploration of gigapixel imagery, data transfer overhead to the visualization nodes can be substantial. To texture a 1.5 gigapixel display at full detail using 256^2 tiles, more than 13 gigabytes per second of bandwidth are required for a 30 fps refresh rate! In the above example, all the displays of the Reality Deck were textured at full detail (assuming a 1-to-1 mapping between display resolution and texels). However, this detail can only be perceived when the user is standing at an *optimal* distance from the displays (in our case, roughly 31"). At larger distances, adjacent pixels become increasingly indistinguishable as their projections within the user's visual system begin to overlap. Consequently, it makes sense to select the image level of detail dynamically, based on the distance of the user from each display.

Most commodity rendering pipelines make a determination of the image level of detail based on the projection of a particular pixel into texture space. In our acuity-driven visualization framework [9], we scale this projection based on the distance D' of the user from a particular display. Specifically, if we assume an original texture space projection A, we calculate $A' = \frac{D'}{D_{opt}} A$ (the further away the user, the larger the texture space projection, resulting in a higher selected level in the LoD pyramid). Based on this notion, we calculate an acuity-based LoD offset m_{acu}:

$$\frac{1}{2^{m_{acu}}} = \frac{D_{opt}}{D'} \Rightarrow 2^{m_{acu}} = \frac{D'}{D_{opt}} \Rightarrow m_{acu} = log_2(\frac{D'}{D_{opt}})$$

This offset is biased for quality, for a resulting LoD level $m' = m_{MIP} + \lfloor \max(0, log_2(\frac{D'}{D_{opt}})) \rfloor$, where m_{MIP} is the normal LoD selected by the rendering system. The operation

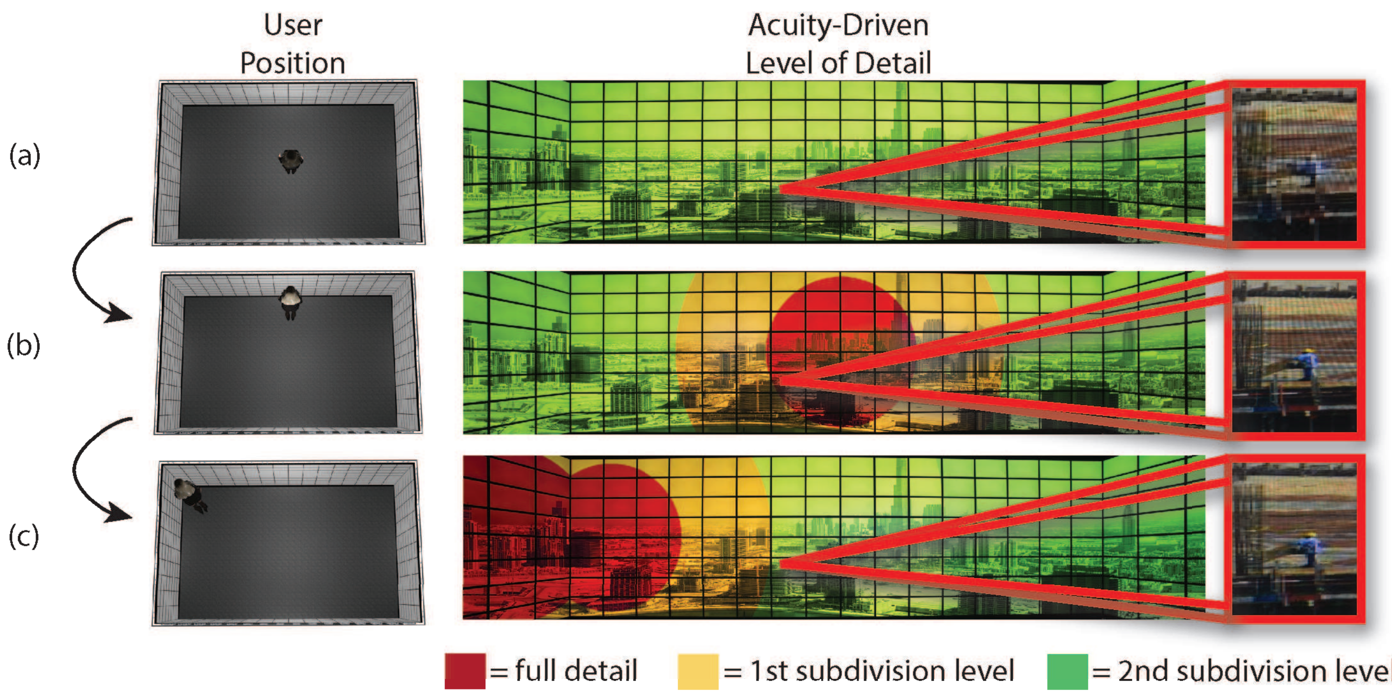

Fig. 8: Our acuity-driven gigapixel visualization technique used during the exploration of a gigapixel panorama of Dubai from Gigapan. (a) The user starts in the middle of the Reality Deck (her position is shown in the renderings on the left). Based on her distance from the displays, our system selects the appropriate LOD level for each pixel on the display. The current LOD is color-mapped and visible in the wide-angle photographs from the interior of the Reality Deck during operation. As the user approaches the front wall (b), the system adapts the LoD, delivering the full detail of the data at the displays that are closest to the user. The LOD selection is updated as the user moves within the facility (c). The insets on the right show zoomed in views of a small section of the front wall, illustrating the change in resolution as the user's position changes.

of our technique (which happens per-pixel, on the GPU and with minimal performance overhead) is illustrated in Fig. 8.

Since this acuity-driven texturing approach results in less than maximum detail being delivered to the displays, we evaluated potential impact on user performance. We designed a user study that had subjects search for targets of different sizes within a gigapixel resolution image. Our analysis (exposed in detail in [9]) indicated that users did not have to make special accommodations in their distance from the displays when searching for targets while using our technique. Finally, we benchmarked our technique on synthetic usage sessions stemming from real tracking data and observered a substantial 70% reduction in data transfer overheads.

Frameless Visualization

Most applications use the GPU to produce a set of pixels over a regular 2D grid, which is then displayed using a double buffering scheme. As the resolution increases, so does the latency associated with the generation of a single frame. Frameless rendering has been proposed as an alternative [10], in which pixel computations are decoupled from the display system by approximating full resolution images from sparse sample sets, at the expense of temporal coherency.

We have developed a novel system for the reconstruction of high resolution and high frame rate images from a multi-tiered collection of samples that are rendered framelessly. In contrast to the traditional frameless rendering techniques, we generate the lowest latency samples locally. These initial points also guide the sampling of more complex and computationally expensive effects (e.g., global illumination).

Our system is based on a distributed single-pass ray-casting volume renderer. At the level of a single GPU and its attached displays, raycasting is used to generate samples at sub-native resolution which are asynchronously reconstructed into a low resolution temporally upsampled preview image. This image is also used during the creation of the priority map that guides the remote rendering of higher-quality unstructured samples. This map can also be modified based on the user's position within the Reality Deck.

An external source can provide a stream of rendered samples, which are combined with the ones stored locally to progressively refine the final image. The illustration in Fig. 9 uses a GPU cluster as the sample source, however, depending on the display configuration, the source can be an auxiliary GPU in the same system or even a cloud-based service. The resulting sample stream is broadcast across the network and each visualization node stores samples that belong to the local

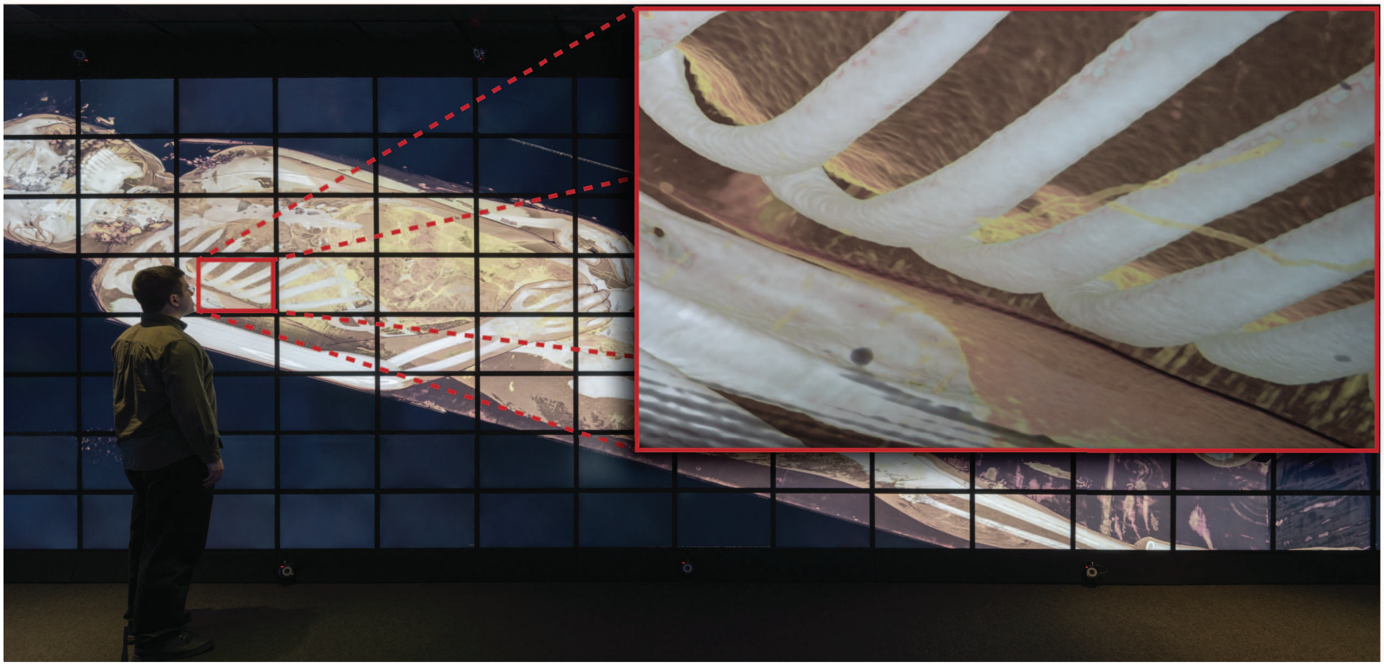

Fig. 9: Exploration of the Visible Human data set using our frameless visualization technique on a 471 megapixel section of the Reality Deck. Despite the massive rendering resolution, the system remains interactive, achieving a stable 5 fps. Traditional volume rendering requires approximately 10 seconds for generating one frame, resulting in a non-interactive experience. The inset shows the image quality (including high quality filtering, ambient occlusion and global illumination) that can be achieved by allowing the frameless visualization algorithm to converge.

viewport in a buffer that matches the full resolution of the display.

During compositing, various sample attributes (e.g., age) are used to determine which samples are utilized during interaction. Depending on the proximity of the rendering source to the local GPU, the latency can vary widely. Local, low-latency samples can be reliably used during scene changes. Contrarily, samples generated remotely can have very high latencies depending on the network configuration and the rendering parameters, and can only be used to improve the resolution and image quality of static scenes.

We demonstrate our frameless visualization technique in the Reality Deck using the Visible Human volumetric dataset ($512 \times 512 \times 2048$ resolution). In this example, all render nodes use the same rendering modality but differ in terms of reconstruction filter and step size. Fig. 9 shows the result of our technique on an 471 megapixel section of the Reality Deck, when backed by a 30-GPU cluster. Double-buffered volume rendering at full resolution requires approximately 10 seconds per frame, during which the display is frozen.. In contrast, our system maintains interactive response (at a steady 5 fps), while image convergence can be achieved in between 1 and 2 minutes depending on the view. Network bandwidth is a major limiting factor in further scaling of the system and we plan to investigate the performance impact of low-latency Infiniband networking.

Discussion and Lessons Learned

Building a facility such as the Reality Deck is an act of balancing costs, features and feasibility. For example, early on in the design process, we planned on incorporating stereoscopic 3D in the platform. However, an examination of the hardware landscape resulted in two conclusions. Either we would have to utilize high-frame rate active stereo panels or a passively polarized solution. Active-stereo monitors generally offered high pixel densities (e.g., 1080P resolution at 23" diagonal) but were implemented using Twisted-Neumatics or similar panel technologies that compromised on color consistency and viewing angles, two features that are of utmost importance for a scientific visualization display that promotes physical navigation, potentially placing users at grazing angles in relation to some monitors. On the other hand, passive solutions half the display resolution when stereo is enabled and can suffer from "polarizer ghosting" when not viewed head-on [4]. The ghosting issue can be resolved by adjusting the polarizer layer on the display, but this obviously eliminates the option of utilizing commodity products and results in increased cost. More advanced solutions can provide individual stereoscopic image pairs for multiple users by combining time-multiplexing, polarization altering and color channel decomposition from multiple projectors [11]. While the collaborative potential of such an approach is substantial, its practicality for deployment within a large visualization system is limited, since it requires

6 DLP projectors to provide multi-user stereo frame buffer to six persons.

Since a primary driver in the construction of our facility was pushing the boundaries in aggregate resolution, we opted to forego stereoscopic 3D and rely on depth cues provided via motion parallax (which have been demonstrated to be stronger than stereo alone for some tasks). Given the collaborative nature of the Reality Deck, handling head-tracking in multi-user scenarios is important. The solution depends on the application at hand. Approaches such as view-vector clustering followed by image blending [12] can be deployed if the computational cost associated with the multiple rendering passes can be afforded.

Similarly, in our design process we balanced between bezel size and total resolution, landing on the side of visual fidelity. A number of companies provide "bezel-less" or ultra-narrow bezel display products (based on LCD or DLP technologies). However, such products are generally targeted at the digital signage industry and not optimized for pixel density (e.g., it is common for such offerings to provide a resolution of 1920×1080 at a 55" diagonal, resulting a pixel density of 40.05 versus the 108.79 PPI provided by the Reality Deck, substantially reducing the space within which users receive $\frac{20}{20}$ visual acuity).

The design process of the Reality Deck commenced in 2011 but future system builders will face similar dilemmas (with higher pixel density quantifiers). For example, 4K monitors are becoming increasingly affordable, offering approximately 160 PPI at 28" diagonals. A system comprised of such displays could conceivably be constructed today (with performance and connectivity considerations being potential caveats). The question becomes whether the already-observed benefits of large displays scale to such extreme resolutions and display footprints (prior research was conducted on significantly smaller systems). We feel that the Reality Deck is the first platform capable of answering such questions and are already in the process of quantifying the scalability of various visualization tasks when the display spans upwards of 1 gigapixel in aggregate.

On a closing note, the Reality Deck (and other facilities such as the CAVE2) are fundamentally different from prior visualization systems in that they not only define a display "surface" but also a clearly outlined "space" that is large enough to promote movement for users. The techniques described above rose organically from patterns in this movement. The acuity-driven level of detail scheme leverages a fundamental property of large displays (that users may be looking at portions of the visualization that lay past their visual acuity threshold), which is amplified by the sheer size the Reality Deck. The Infinite Canvas was inspired by observing the tendency of users to scan the visualization by physically navigating from one end of the display to the other. Our techniques are motivated by relatively basic observations, but they can still enable seamless interactions and substantially improve system performance. While physical navigation has been an active research area for multiple years, we feel that its evaluation within the context of specific applications in extremely large immersive systems presents fertile ground for innovation.

CONCLUSION

We presented the motivations and design principles behind the Reality Deck - the world's first immersive gigapixel resolution display. It is a cost-effective device that provides $\frac{20}{20}$ visual acuity with full $360°$ physical immersion in a $33' \times 19'$ venue that can be collaboratively shared with a large number of people. We also described three techniques that provide advanced interaction capabilities and foster real-time display of different types of data - the Infinite Canvas, acuity-driven gigapixel visualization and frameless visualization. The Reality Deck offers significant prospects for future large-scale IVEs with flexible layouts, increased resolution and support for stereoscopic rendering.

ACKNOWLEDGMENTS

This work has been supported by NSF grants CNS-0959979, IIP-1069147 and IIS-0916235. Datasets used to generate the figures in this article are courtesy of Gigapan, NASA, USGS, NYC OpenData, JPL, University of Arizona, National Library of Medicine and the RCSB Protein Data Bank.

REFERENCES

[1] D. Marks, H. Son, E. Tremblay, J. Ford, P. McLaughlin, M. Gehm, R. Stack, S. Feller, J. Kim, and D. Brady, "Optical testing of the aware wide field 2-gigapixel multiscale camera," *Frontiers in Optics*, 2011.

[2] T. Ni, D. Bowman, and J. Chen, "Increased display size and resolution improve task performance in information-rich virtual environments," *Graphics Interface*, pp. 139–146, 2006.

[3] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart, "The CAVE: Audio Visual Experience Automatic Virtual Environment," *Communications of the ACM*, vol. 35, no. 6, pp. 64–72, 1992.

[4] A. Febretti, A. Nishimoto, T. Thigpen, J. Talandis, L. Long, J. Pirtle, T. Peterka, A. Verlo, M. Brown, and D. Plepys, "Cave2: a hybrid reality environment for immersive simulation and information analysis," *SPIE Electronic Imaging*, pp. 864 903–864 903–12, 2013.

[5] K. Petkov, C. Papadopoulos, and A. E. Kaufman, "Visual exploration of the infinite canvas," *IEEE Virtual Reality*, pp. 11–14, 2013.

[6] F. Qiu, B. Zhang, K. Petkov, L. Chong, A. Kaufman, K. Mueller, and X. D. Gu, "Enclosed five-wall immersive cabin," *4th International Symposium on Advances in Visual Computing*, pp. 891–900, 2008.

[7] S. Eilemann, M. Makhinya, and R. Pajarola, "Equalizer: A scalable parallel rendering framework," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 3, pp. 436–452, 2009.

[8] C. Hollemeersch, B. Pieters, P. Lambert, and R. Van de Walle, "Accelerating virtual texturing using cuda," *GPU Pro: Advanced Rendering Techniques*, pp. 623–641, 2010.

[9] C. Papadopoulos and A. E. Kaufman, "Acuity-driven gigapixel visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 12, pp. 2886–2895, 2013.

[10] G. Bishop, H. Fuchs, L. McMillan, and E. J. S. Zagier, "Frameless rendering: Double buffering considered harmful," *SIGGRAPH*, pp. 175–176, 1994.

[11] A. Kulik, A. Kunert, S. Beck, R. Reichel, R. Blach, A. Zink, and B. Froehlich, "C1x6: a stereoscopic six-user display for co-located collaboration in shared virtual environments," *ACM Transactions on Graphics*, vol. 30, no. 6, p. 188, 2011.

[12] J. Marbach, "Image blending and view clustering for multi-viewer immersive projection environments," *IEEE Virtual Reality*, pp. 51–54, 2009.

SIDEBAR: VISUAL ACUITY

When a user looks at a display of fixed resolution, approaching its surface makes the individual pixels increasingly discernible. Visual acuity (VA) is often used to "quantify" the quality of the visualization a display can deliver to an observer. In a hypothetical 2D world, a display of width W (in inches) and horizontal resolution H (in pixels), the dot pitch (distance between two pixels) is $P = \frac{W}{H}$. If the user is looking at a single pixel at the center of the display from D inches away, the angle covered by that pixel on the user's horizontal FoV is $\approx tan^{-1}(\frac{P}{D})$. The VA for this setup is equal to the portion of the pixel that covers $\frac{1}{60^{th}}$ of a degree of the user's FoV or $1/tan^{-1}(\frac{P}{D})$.

This VA metric correlates directly to the more commonly used Snellen fraction which quantifies vision. The fraction $\frac{20}{X}$ corresponds to "this person can see at 20 feet what a person of average vision can see at X feet". A person's Snellen ratio is determined by asking them to distinguish characters (or optotypes) that have been precisely spaced in order to project to certain visual angles at fixed distances. By dividing a Snellen ratio, we can get the previously described VA metric.

Frequently, Snellen's fraction is used to quantify the VA of a display (rather than the less intuitive scalar metric). For example, the display of an iPhone 5 has a horizontal resolution of $H = 640$ and a screen width of approximately $W = 2.17''$. Looking at the display from $D \approx 10''$, we get a VA of ≈ 0.856 pixels per minute or a Snellen ratio of $\frac{20}{23.4}$. The same calculation along the diagonal yields a VA of $\frac{20}{21}$, hence the term "Retina" Display since if the user were to approach the display more, they would not be able to perceive additional visual information.

Recently, the argument was made that VA should not be the primary driver in tiled display design [1]. While doubling the VA metric does not necessarily correspond to doubling in the perceived visual information, there are positive effects (e.g. reduced perceived aliasing). The proliferation of high DPI screens indicates that sharp visuals matter when utilizing a display. More importantly, an increase in resolution allows the visualization of minute details that would otherwise require zooming to resolve. Increasing the VA of the display affords additional levels of detail that the user can reach simply by moving closer to the screen, expanding the potential for physical navigation (which is beneficial to the exploration process [2]). Andrews et al. [3] offer great insights into the effect of resolution on different types of data visualizations. As a result, optimizing the VA metric was a primary concern during the design of the Reality Deck.

REFERENCES

[1] K. Moreland, "Redirecting research in large-format displays for visualization," *IEEE Symposium on Large-Scale Data Analysis and Visualization*, 2012.

[2] R. Ball and C. North, "Realizing embodied interaction for visual analytics through large displays," *Computers & Graphics*, vol. 31, no. 3, pp. 380–400, 2007.

[3] C. Andrews, A. Endert, B. Yost, and C. North, "Information visualization on large, high-resolution displays: Issues, challenges, and opportunities," *Information Visualization*, vol. 10, no. 4, pp. 341–355, 2011.