Clusters to Single Nodes: Advanced Visualization Systems

White Paper

Author: Brent Dingle, Ph.D.

Last Updated: November 2019*

Abstract

This document discusses the design and development of a single-node advanced visualization system. The current implementation uses a standard engineering HP Z8 with three Nvidia P6000 cards. The system was tested using ten Barco F70 projectors mosaicked together to support stereoscopic viewing. The paper begins with a brief history of advanced visualization systems and continues into a description of a traditional cluster-based CAVE. This establishes a baseline system. Using this baseline as a reference, a single-node system is then described. Through these descriptions it will be shown the single-node system is lower in cost, provides more robust capabilities, and supports more intuitive user-interfaces than the baseline system.

Contents

1 Motivation and History

Throughout time humans have employed visualization systems. Need and usage covers a near countless number of areas ranging from maps, scientific drawings, historical records, data plots, and the list continues to grow. Visualization has been used to convey abstract and concrete ideas and it takes many forms such as: visual analytics, scientific visualization, information visualization, visual communication, and educational visualization.

Visualization may be defined as any technique for creating diagrams, images, or animations to communicate a message. Almost every field of study makes use of visualization systems, many rely heavily on computer graphics. The importance and need for such systems are well documented throughout history and the design, development and improvement of such systems remains critical today.

In 1965, Ivan E. Sutherland published the following:

> *"We live in a physical world whose properties we have come to know well… We sense an involvement with this physical world which gives us the ability to predict its properties well… A display connected to a digital computer gives us a chance to gain familiarity with concepts not realizable in the physical world… The user of one of today's visual displays can easily make solid objects transparent – he can "see through matter!" Concepts which never before had any visual representation can be shown.*
>
> *By working with such displays of mathematical phenomena we can learn to know them as well as we know our own natural world. Such knowledge is the major promise of computer displays [1]."*

This publication inspired many to pursue and create advanced visualization hardware and software systems. Early efforts coming out of this paper focused on head-mounted displays (HMDs) and virtual reality (VR) systems.

In 1987 a panel report to the National Science Foundation (NSF) stated:

> *"…Visualization in Scientific Computing promises radical improvements in the human/computer interface… It can bring advanced methods into technologically intensive industries and promote the effectiveness of the American scientific engineering communities… [It] will also provide techniques for exploring an important class of computational science problems, relying heavily on cognitive pattern recognition or human-in-the-loop decision making… The university/industrial research and development cycle is found to be inadequate for Visualization in Scientific Computing. The programs and facilities are not in place for researchers to identify and address problems far enough in advance… [2]"*

Following this report various conferences, workshops and funding sources became more amiable towards supporting and conveying research done in Visualization.

In 1992, Cruz et al. introduced the concept of a CAVE visualization system [3]. It was designed as a visualization tool for scientific application. In 1993, within a follow-on paper, they describe the CAVE setup as a cube shaped apparatus designed to project images onto three walls and a floor [4]. The cube shape was used to approximate a sphere. Thus, the complexities of creating a spherical projection were avoided. The display supported stereoscopic 3D visualization.

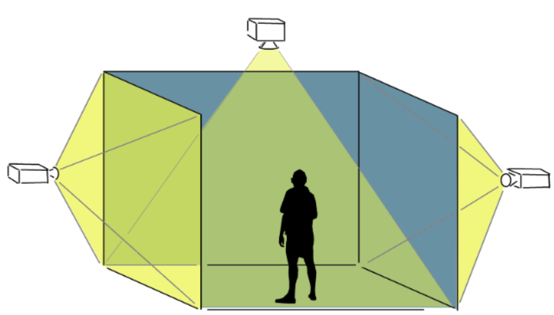

Figure 1: CAVE setup with 3 walls and a floor projection.

2 Cluster-based Systems

The early CAVE visualization systems, as described by Cruz et al. were powered by a cluster of computers [3]. While other combinations of walls, ceiling and floor were possible, most of these systems were setup as a cube with a projection on three walls and the floor [Figure 1]. This typical setup used five computers. One computer for each projector and a head node to coordinate between them [Figure 2]. The system was equipped with wired head and hand tracking to allow the primary user to change perspective and interact with virtual objects. Each projector ran at a resolution of 1280x512 pixels at 120 Hz. To allow stereoscopic 3D viewing, users wore glasses that worked in conjunction with the projectors. The projectors displayed left-eye and right-eye images in sequence. The glasses were designed to allow the user's left eye to see only the left image and the right eye to see only the right image. The concept of this left-eye-right-eye display is based on stereopsis, findings on which date back to 1838 [5], if not earlier.

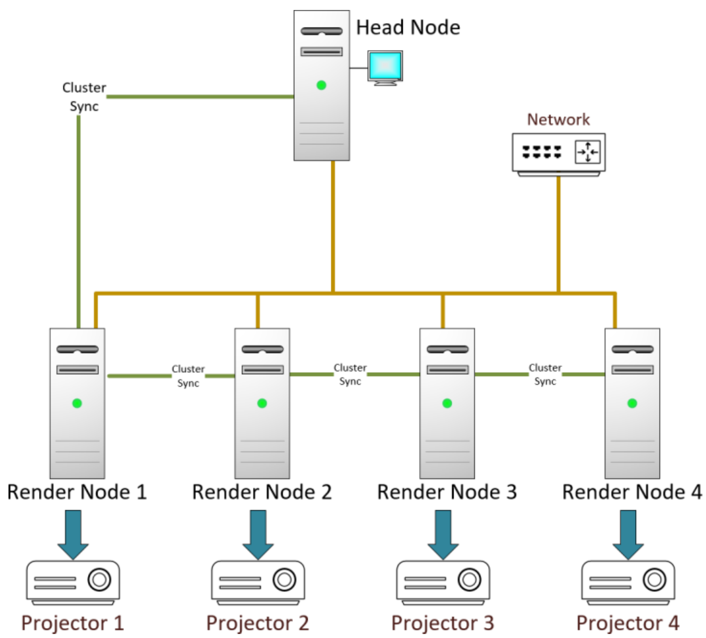

Figure 2: Cluster-based CAVE system.

The first CAVEs, and those that followed, focused largely on the stereoscopic capability of the displays. The first paper introduced this as new type of VR and described 3 major benefits of such [3]: suspension of disbelief, viewer centered perspective, and immersion. These benefits remain today. Equally important today is a goal identified in 1992, and inherent in the definition of visualization:

> *"One of the most important aspects of visualization is communication. For virtual reality to become an effective and complete visualization tool, it must permit more than one user in the same environment. [3]"*

This identifies a CAVE-based system as being designed to be a collaborative environment. Interestingly this goal may have been intended as a point to distinguish the CAVE system from head-mounted display (HMD) systems. However, with the massively multiplayer online (MMO) networking capabilities of today, it is possible to achieve this goal and benefit using HMDs. This, in a sense, makes HMDs an interesting, highly distributed, highly dynamic, cluster-based visualization system. It also means the HMDs can become part of a CAVE-based system – acting as a remote component. This in turn leads to the possibility of collaboratively networked CAVE-based systems. This extends beyond the intent of this paper but will be mentioned again in the conclusion.

Another adaptation of the original CAVE systems is the use of stereoscopic capable LCD monitors instead of projectors. An excellent example of such a system is the CAVE2. This system was originally developed at the University of Illinois at Chicago (UIC) in the Electronic Visualization Laboratory (EVL). It was funded by the National Science Foundation's Major Research

Instrumentation program and the Department of Energy. The CAVE2 debuted in October 2012. It is a trademark of the University Of Illinois Board Of Trustees [6]. It is now marketed by Mechdyne Corporation, and they are the only licensed integrator of the CAVE2 [7].

The system remains a cluster-based computer system. In the described design [8], the CAVE2 system is a cylindrical setup made up of 18 columns with 4 displays in each column. Each display pair is controlled by a separate computer. Thus, the system is a cluster of 36 computers plus a head node [Figure 3].

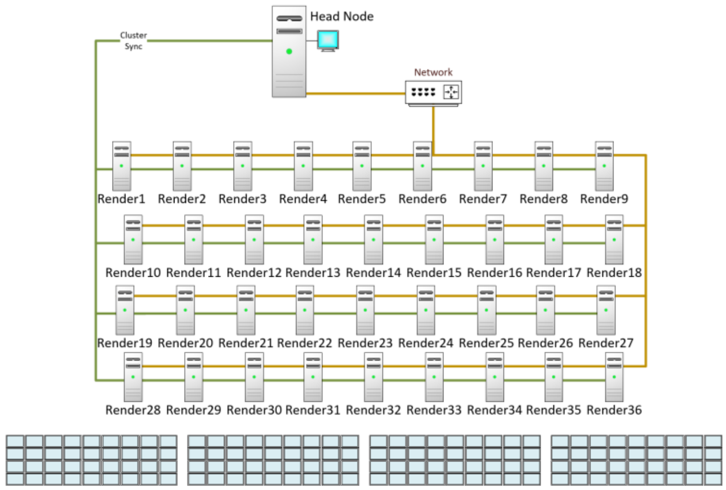

Figure 3: CAVE2 visualization system with 36 computers, 1 head node and 72 displays, networked together.

The CAVE2 system allows for a variety of configurations and reduces the physical footprint needed [Figure 4]. It provides a hybrid capability to switch between 2D and 3D or use both at the same time. The system includes an optical tracking system for user interfacing and interaction. This tracking works via the placement of tracking balls on a set of glasses and on control wands. The input and tracking data from the wands and glasses is then processed by a dedicated input server machine and streamed to CAVE2 applications. For stereoscopic display users still wear specialized glasses. However, the glasses are wireless and make use of polarization and specialized display screens which allow the user to correctly see the left images with the left eye and the right images with the right eye [8]. CAVE2 systems also include sophisticated audio options.

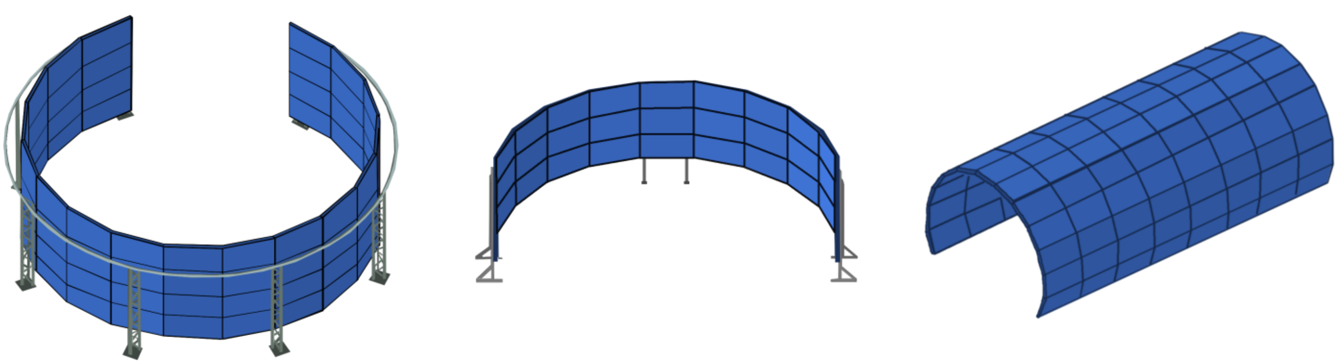

Figure 4: Possible configurations of a CAVE2-based system.

While the CAVE2 improved upon the original CAVE design, and a 320° panoramic view is striking, there remain several criticisms of cluster-based systems. From a business standpoint, the need for customized middleware is a significant cost. Recall, this middleware is necessary for any application to run and make full use of the displays. It is specifically designed and written for every application. To run any new software requires a modification of the middleware to support it. This requires time, money, and specialized experts to run the software. From a user standpoint, the bezels between the LCD screens limit the effectiveness of the stereoscopic view. The bezels prevent the illusion of 3D from popping "out" of the screen, thus the 3D effects are limited to "behind" the displays. Another point of user concern is one of view-ability of information. A 320° view is amazing for images and scenes. The head rotates around the CAVE2 room and the user feels immersed. However, if the data is mostly 2D, the need to rotate around the room to see all of it is hindering to absorbing and seeing it as a collective whole of information.

In sum, the advancement of cluster-based CAVE systems has focused largely on improving the visual display devices. The fundamental computer setup remains unchanged. The system design still requires multiple computers – in fact the CAVE2 requires more than in the original CAVE design [Figure 3]. Operating such systems requires a significant amount of skill and knowledge in networking to run. It mandates a controlling head node and the continual development and maintenance of middleware software, or the development and maintenance of software specifically designed for running on a CAVE-based system. Likewise, the interfacing remains largely unchanged. The early setups used wired gloves and wired head tracking. The newer cluster-based systems have migrated to wireless tracking but still require expensive, specialized equipment.

3 Single-Node Systems

The cluster-based, advanced visualization systems, in 2013, offered cost reduction points compared to previous CAVE cluster systems. The original systems relied on expensive lamp-based projectors. Replacing those with LCDs reduced some of the initial cost and removed much of the maintenance costs of calibration and bulb replacement. The LCDs also provided higher resolutions and could operate at higher light levels.

As of 2019, laser projectors offer resolutions and operate at roughly the same room light levels as LCD displays. When considering the total square footage of display capability, the costs become reasonably comparable. The footprint of projectors has also been reduced with the advancement of short-throw, or ultra-short throw, lenses. Maintenance becomes almost equivalent, with additional options of automatic keystone correction, blending, and pixel alignment now built into many projector systems. The potential return to using projectors provides several benefits. They reduce the number of bezels visible to the user, provide a gentler form of light on the user's eyes, and offer stereoscopic effects that "pop out" of the screens. In essence the case of projectors versus LCDs, as of 2019, becomes a discussion of user preference. The choice may be influenced by intended use or physical constraints, but not so heavily by display cost.

This moves much of the variable cost of an advanced visualization system into other aspects of the design. Specifically, it is not the cost of the display technology, but rather the operating hardware, software and complexity that drive cost. In the case of cluster-systems, with the advancement of graphics cards, some reduction of cost can be gained by controlling four displays instead of two per computer. This would cut the computer cost in half. There is also the possibility of using head-mounted displays. This would reduce the number of displays and number of computers needed to equal the maximum audience size. However, neither of these solutions removes the complexity of setting up and using the system. Both still require experts to configure and run, as well as specialized hardware and software.

This leads to the desire and need for a system that requires a lower expertise level to run, and preferably less hardware. Fortunately, with the advancement of graphics cards and capabilities this is possible to achieve as a single-node system. A single-node system uses one computer to control all the displays. There is no middleware required and no local network required.

The creation of a single-node system is possible due to advancements in video graphics cards. Specifically, this type of system relies on the ability of Nvidia cards to mosaic a set of displays together to function as one display. It also makes use of warp and blend technologies, either via the Nvidia drivers, or using the built-in capabilities of newer projectors. It is preferred to use the projectors' native abilities as this offloads the processing to the projector. However, the Nvidia drivers are capable should the projectors not have the built-in ability, or if projectors are not being used. Another advancement useful to single-node systems is the ability of Nvidia drivers to send sync signals with the DisplayPort (DP) video signal to allow stereoscopic viewing across multiple displays. This removes the need for extra cables to send this sync.

Figure 5: Single-node system at Raytheon Space and Airborne Systems.

As a result of these technological advancements, it is possible to run an advanced visualization system, which supports stereoscopic display and regular 2D display, using <u>one</u> engineering level machine, with the simple addition of video cards. This saves time, money, and improves usability. This has been done at several Raytheon locations. One such system is located at Raytheon Space and Airborne Systems (SAS). As of 2019 the primary system being used in the SAS Immersive Design Center is formed from an HP Z8 with three Nvidia P6000 cards, displaying on ten Barco F70 projectors [Figure 5].

To set up the system requires the ability to run video cables, and a general understanding of computer hardware and the display hardware being used. It also requires an understanding of how warp and blending works and how Nvidia mosaic works. These are skills a typical IT, Computer Science or Engineering undergraduate can learn and apply. In fact, this was verified at Raytheon SAS in the summer of 2018, where two interns were taught these skills, and required to apply them. This teaching-and-learning experiment was repeated with four other interns in the summer of 2019. All resulted in success.

To run the system requires no more than being a Windows user. Once configured any user can log into the single-node computer, just as any other computer. Whatever software is available will run on a single-node system. This opens a world of possibilities. Beyond being able to natively run any Windows software, the single-node can be subjected to most security policies that are typical of businesses today. More specifically it does not require any specialized Windows build beyond what would normally be applied to an engineering workstation (e.g., HP Z8). This means the single-node can be placed on whatever business LAN network is available. This provides direct access to network data and applications.

All these single-node features contrast with a cluster-based system. Cluster systems require specialized, high-end, server computer systems and server racks. Because of the specialized Windows builds required to operate they usually are firewalled off from any other business networks and data. Cluster systems also require middleware to run most applications, which significantly limits what software is available to the user. They also require localized networking and cabling to function. These in turn require their own level of maintenance and expertise.

Clearly the usability of a single-node system greatly exceeds that of a cluster-based system. In addition, the single-node systems reduce cost by reducing the computer count of 37 computers down to 1 [Figure 6]. By using projector displays a single-node system further reduces the cost by decreasing the number of video cards, from 37 down to 4, or less.

By reducing the complexity of the system, the cost and time of maintenance is also reduced. A single-node system has no networking required, beyond a typical connection to the business LAN. The projectors require an initial setup to achieve the desired warp and blend and mosaic effect desired. They may require bi-annual check-ups to ensure the pixel alignment stays correct, but many projectors are now offering automatic warp, blend, and alignment features. The lifespan of most projectors extends into the 30 to 50 thousand hours of life. This equates to about five years or more of use, which is comparable to most stereoscopic enabled LCDs.

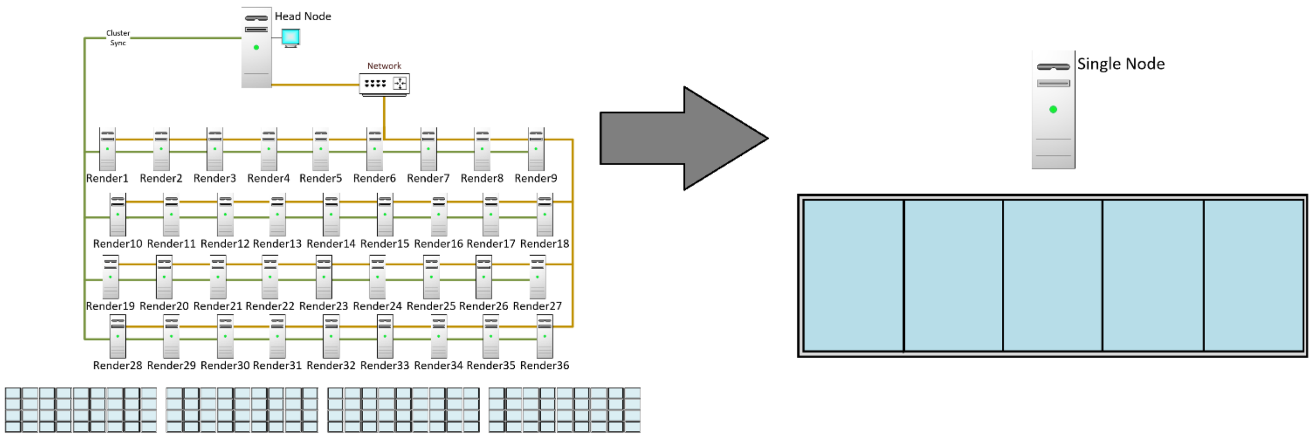

Figure 6: Cluster-based to Single-node reduces complexity, saving time and money, improving usability.

In general, single node systems allow for cheaper and easier investigations to further advance visualization development and usage. Such systems support: scalability through simplification and standardization, portability by decreasing the size and number of components, and collaboration by making systems easier to access and use.

Because a single node system is independent in implementation from the display devices, the display does not have to be enormously huge. It can be just 3 or 2 or even 1 projector or LCD display. This means a smaller system can be standardized. The major difference being the number of video cards, and the amount of RAM and disk space included in the smaller systems. With the creation and standardization of smaller systems it becomes possible to move advanced visualization into the average conference room, or into portable systems. While the display size will be limited, this would allow a larger number of employees to become more familiar with using advanced visualization systems by making them more accessible and available. It also opens the possibility to network these systems together for shared, remote experiences.

The concept of portable visualization systems with stereoscopic capability has been explored in several ways. Some of the early attempts made use of Mechdyne's Rapidly Operational Virtual Reality (ROVR) system and similar [9]. These types of systems were still extremely large and required significant effort to transport. They also required a significant amount of expertise to set up and run. They have been replaced by smaller backpack-sized systems, designed and developed by the author and his colleagues. These systems typically include a laptop or very small desktop computer and a short throw projector, such as an HP Z2 Mini-G4 and an Optoma GT5600. Investigations to further reduce the size and weight of these systems are in progress. The items being considered are even smaller form factor computers and mini-sized projectors, such as the WOWOTO H10 or ELEPHAS RD606. By using single-node systems, these portable solutions bring the system cost to under $3000. They provide most of the functionality of the larger systems and little to no expertise in setting up and running. They also allow further remote collaboration possibilities and easy access to visualization systems.

Another portable advanced visualization system is in the continued development of head-mounted displays (HMDs). In essence these are another form of a single-node advanced visualization system. These VR and AR devices offer an option to connect to larger systems on an individual basis. While it is not useful to put HMDs on twenty people sitting in the same room, it may be useful for those individuals not in the room. Effectively it will be possible to transmit the entirety of a large-scale visualization system (i.e., a CAVE) to a remote user wearing an HMD. Thus, the remote user will see the same screen as the twenty people in the room would see. This could further be enhanced by overlaying a VTC window into the HMD's view. Notice this possibility is not far from what massive multiplayer online (MMO) games do today. It is merely a shift in the game environment. Instead of a game world, the player is placed in a much simpler room with a large display. On that display is copied the content of what is in a real, physical room.

4 Conclusion

Single-node advanced visualization systems offer significant benefits compared to cluster-based systems. They are less complex. Require less expertise to use and maintain. They provide easier access to networks and data. They support better user experiences, and they reduce costs. For example, estimating a $5000 cost for each machine in a cluster of 37 machines compared to a $20,000 cost for a single node machine, the savings estimate on computers would be $190,000.

Single node advanced visualization systems offer significant gains in scalability and usability. By using a single computer to run the displays it is much easier to standardize and replicate the system. This leads to standardization in processes, usability options, and maintenance. Single node systems promote the continued advancement and use of visualization systems. This will continue to provide benefits in research, development, and implementation throughout the life cycles of technologically advanced products.

Future work will include the continued pursuit of making the system simpler. This may focus on reducing the number of video cards, scaling down the computer size, reducing the size of the display devices, or similar. Future efforts will also be interested in exploring interfacing options. A single node system as described does not require an advanced tracking system, a simple mouse and keyboard is sufficient. An advanced tracking system could be incorporated, but a better option, currently being investigated, is a simple webcam with depth sensing capabilities.

In the future, better incorporation of standard A/V controls will be done. Already a need has been established and some ad-hoc solutions have been deployed to allow hardware-in-the-loop and simulations to be interfaced to the displays. These options, again, are made easier to achieve by using a single-node system. Doing such things with a cluster-based system was not reasonably easy.

In sum, single-node advanced visualization systems are less expensive and provide more robust capabilities than a cluster system. They also support more intuitive interfaces and are a more scalable solution. Perhaps most significantly, they will continue to be useful into the future, regardless of the specific path technology decides to take. Now is the time to prepare and develop infrastructures to support continued development. The future possibilities are many and missing such opportunities will be costly. There is risk in embracing cutting edge technology, but when the short-term payoff is evident and the long-term likely, in the case of advanced visualization, the risks are easily overcome by the gains.

References

1. Ivan E. Sutherland. 1965. "The Ultimate Display," in Proceedings of the IFIP Congress 65, 2, pp. 506-508.
2. B. H. McCormick, T. A. DeFanti, and M. D. Brown (Eds.). 1987. SIGGRAPH Comput. Graph. 21, 6 (November 1987).
3. Carolina Cruz-Neira, Daniel J. Sandin, Thomas A. DeFanti, Robert V. Kenyon, and John C. Hart. "The CAVE: Audio Visual Experience Automatic Virtual Environment," Commun. ACM 35, no. 6 (June 1992), pp. 64-72.
4. Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A DeFanti, "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," in Proc. Of the 20th Annual Conf. on Comp. Graphics and Interactive Techniques, SIGGRAPH '93 (New York, NY, USA, ACM, 1993), pp. 135-142.
5. Charles Wheatstone, Contributions to the physiology of vision. —Part the first. On some remarkable, and hitherto unobserved, phenomena of binocular vision. 128. Phil. Trans. R. Soc. (January 1838). http://doi.org/10.1098/rstl.1838.0019.
6. CAVE2: Next-Generation Virtual-Reality and Visualization Hybrid Environment for Immersive Simulation and Information Analysis. Electronic Visualization Laboratory, University of Illinois at Chicago. https://www.evl.uic.edu/entry.php?id=2016, last verified on November 7, 2019.
7. CAVE2: The Hybrid Work Environment. Mechdyne Corporation. https://www.mechdyne.com/hardware.aspx?name=CAVE2, last verified on November 7, 2019.
8. Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, J. D. Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, Dana Plepys, Dan Sandin, Luc Renambot, Andrew Johnson, and Jason Leigh, "CAVE2: A Hybrid Reality Environment for Immersive Simulation and Information Analysis," vol. 8649 (The Engineering Reality of Virtual Reality 2013, International Society for Optics and Photonics, 2013), https://doi.org/10.1117/12.2005484.
9. Mechdyne: Portable Visualization Solutions, https://www.mechdyne.com/hardware.aspx?name=Portable+Visualization+Solutions, last verified November 18, 2019.

** Origin Date of paper is circa: February 2016. The content of this paper is based on an unpublished paper originally written by author: Brent M Dingle. The original paper was written in preparation for various presentations and interviews the author was participating in at the time. In 2019 additional detail and support was added to existing. Small updates were made in 2020 and 2023, mostly grammatical. A larger update is pending discussing advancements and usage of LED MicroTiles.*