

Extending Virtual Reality Display Wall Environments Using Augmented Reality

Arthur Nishimoto

University of Illinois at Chicago
Chicago, Illinois

Andrew Johnson

University of Illinois at Chicago
Chicago, Illinois

ABSTRACT

Two major form factors for virtual reality are head-mounted displays and large display environments such as CAVE® and the LCD-based successor CAVE2®. Each of these has distinct advantages and limitations based on how they're used. This work explores preserving the high resolution and sense of presence of CAVE2 environments in full stereoscopic mode by using a see-through augmented reality HMD to expand the user's field of regard beyond the physical display walls. In our explorative study, we found that in a visual search task in a stereoscopic CAVE2, the addition of the HoloLens to expand the field of regard did not hinder the performance or accuracy of the participant, but promoted more physical navigation which in post-study interviews participants felt aided in their spatial awareness of the virtual environment.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI); Mixed / augmented reality; Virtual reality.**

KEYWORDS

Augmented Reality; Virtual Reality; Display Wall; Immersive Displays; Presence

ACM Reference Format:

Arthur Nishimoto and Andrew Johnson. 2019. Extending Virtual Reality Display Wall Environments Using Augmented Reality. In *Symposium on Spatial User Interaction (SUI '19)*, October 19–20, 2019, New Orleans, LA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3357251.3357579>

1 INTRODUCTION

Virtual reality (VR) exists in two major forms: CAVE automatic virtual environment (CAVE) or similar large display wall environments and head-mounted displays (HMDs). Each of these form factors has distinct advantages and limitations based on the use case. CAVE and in particular the LCD-based successor CAVE2 [5] provides high resolution, high contrast visuals with low latency to a group of participants wearing lightweight stereoscopic glasses. These types of large display environments have been shown to be

conducive to collaborative work where high resolution digital content or multiple smaller digital artifacts aid in the group discussion [7–11, 14, 17].

HMDs provide immersive visuals by filling a user's field of view (FOV) and allowing a user to look around at the virtual world by completely filling their field of regard (FOR). This immersion comes at the cost of blocking out any sense of presence of collaborators around them without the use of virtual representations such as avatars. While resolution is increasing with each generation of VR HMDs (currently at a horizontal resolution of 4k pixels), they have yet to match the clarity of LCD display walls such as the 46k horizontal resolution of CAVE2. Augmented Reality (AR) HMDs combine the real world with virtual elements with many of the same limitations as VR HMDs - cumbersome headsets, lower resolution, and for most AR displays limited FOV.

In our 2.3 meter tall cylindrical CAVE2 display, the approximate vertical FOV from the center is 40 degrees. CAVE2 forms a 320 degree cylindrical display, and with its detached backwall can provide the user with a 360 degree horizontal FOR. In comparison, VR HMDs tend to have FOVs of 110 to 115 degrees. High-end AR HMDs costing around \$10,000 such as the NVis SX111 has 64 degrees vertical FOV while cheaper alternatives such as the original HoloLens has approximately 17 degrees vertical FOV and the Magic Leap at 40 degrees.

VR HMDs are successful in cases where FOR is more important for immersion than resolution while large VR capable display walls are successful where resolution and group presence are more important than immersion. This work combines AR and large display wall VR environments to leverage the best of both form factors. Even with the limited FOV in current see-through AR HMDs like HoloLens, we believe framework developed in this work and the study conducted illustrates a first step in combining these technologies to explore the potential of augmenting virtual reality environments in full stereoscopic mode.

2 RELATED WORK

This paper builds on existing work of collaboration on large display environments involving both VR and non-VRs, comparisons with HMD-based VR systems, and AR with display walls.

2.1 CAVE vs HMD

Prior work has examined how pairs or small groups of users explore or analyze the virtual environment in a CAVE or CAVE2 environment and compares the same task performed in a VR HMD. In 3D spatial visualization tasks performed in a CAVE2 and VR HMD between two collaborating users, it was found the HMD performed faster since the two users were able to perform independent tasks on their own HMD instead of using the shared CAVE2 display. There

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SUI '19, October 19–20, 2019, New Orleans, LA, USA

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6975-6/19/10...\$15.00
<https://doi.org/10.1145/3357251.3357579>

were no significant differences found in task accuracy or leadership role between the two systems. Users reported that CAVE2 provided a more satisfying experience because it was easier to communicate with their collaborator since they could see themselves and their collaborator. The lack of control for the non-head tracked user was the main downside of using CAVE2 vs the HMD [4].

2.2 AR and Large Displays

Low cost see-through HMDs and large display wall frameworks have been proposed to aid in interactive data visualization. An AR HMD was used to control and view subsets of the visualization displayed on the wall. This allowed the HMD to interact with the visualization in real-time and preview different parameters before re-rendering the full dataset on the large display [13].

A number of other works have examined using large immersive displays with see-through HMDs in order to simulate outdoor AR in a controlled environment. These examined using the VR display to manipulate individual factors that affect AR performance such as latency, different FOV, FOR, and scene complexity [2, 12, 15].

In these environments, the VR system was used in monoscopic mode instead of stereoscopic since these cases focused on large virtual outdoor environments where virtual objects were projected far away from the user. Our work focuses on virtual objects closer to the user and includes elements drawn in front of the projection plane where stereoscopic cues are more useful and critical to be properly projected to the user's perspective.

The Dataspace environment combines a cylindrical large display environment with AR HMDs to create a collaborative workspace where multiple participants can view information on the large displays in addition to having a center table for viewing an AR representation projected from each HMD wearer's perspective [3]. The AR representation viewed on the table can be related to the data on the large 2D displays, but are still two distinct independent views. Our work examines how these two views can directly interact with one another and in stereoscopic 3D.

3 IMPLEMENTATION

The CAVE2 hybrid-reality environment was designed as a mix between the more immersive six-sided projection based CAVE and high resolution, high contrast thin-bezel LCD screens of modern tiled display walls. The result is a 320 degree horizontal stereoscopic display wall capable of higher resolution and brighter colors than a classic CAVE. Due to the limited stereoscopic viewing arc of passive stereoscopic displays and challenges of tiling curved displays horizontally and vertically, the vertical FOV is limited by the height of the CAVE2 displays.

Powering the 72 passive stereo displays is an 18 node display cluster, a 14 camera Vicon optical tracking system connected to a dedicated tracking PC, and a single master node.

The HoloLens was retrofitted with retro-reflective markers as shown in Figure 1 and tracked using the CAVE2's optical tracking system. This was done to more accurately align the HoloLens optics with the CAVE2 displays by using the same system as normally used with the tracked passive stereo glasses. The HoloLens was also fitted with the same lenses as CAVE2's stereo glasses. In order to avoid the HoloLens internal optics interfering with the polarizing



Figure 1: Standard CAVE2 glasses (L) and retrofitted HoloLens (R) with optical tracking markers and passive stereo lenses

filter on the passive stereo lenses, the lenses had to be attached on the curved exterior of the HoloLens headset instead of the flat surface inside the headset closest to the wearer's eyes.

In addition to having a 6.4 meter diameter tracking space for the user to physically walk around the virtual environment, the user was also given a tracked wand controller for limited virtual translation and rotation.

The Unity game engine and Mechdyne's getReal3D for Unity plugin was used to run and synchronize the Unity scene across all 18 nodes controlled by the master node. A separate server-client system based on Unity's networking framework was used to specifically connect a Unity application running on the HoloLens to the master node. From the master node we updated the HoloLens optics with tracking data from the CAVE2 tracking system as well as updating the state of the virtual environment.



Figure 2: HoloLens view inside CAVE2

Our initial tests with the HoloLens worked with aligning the visuals on the HoloLens with the CAVE2 displays. Given the limited available information on the HoloLens optics we were able to get a decent registration between the visuals HMD and the display wall in head tracked stereoscopic mode. For the purposes of the study, we focused on just the content visible outside of the FOV of the CAVE2 screens. We applied a screen mask on the HoloLens that would make any graphics visible on both the HoloLens and CAVE2

displays to be transparent. In Figure 2 the cutout appears dark from the HoloLens video capture. On the actual HoloLens optics, the dark lower section of the window appears completely transparent.

4 USER STUDY

This study was designed to examine if an AR device can offset the disadvantage of the CAVE2's lack of a floor and ceiling in an immersive task. By using AR techniques to extend the CAVE2 displays we wanted to investigate: (1) would users be able to more effectively search a 3D spatial dataset with the additional of AR, (2) how would physical movement of the user be affected by the use of AR, and (3) how would the virtual navigation be affected with the addition of the AR headset.

4.1 Participants

10 users participated in the study. The participant pool consisted of 6 female and 4 males. 7 of which were graduate students, 1 undergrad, and 2 staff. In the pre-study survey, 8 of the participants reported prior experience with CAVE2 or VR HMDs. 3 participants reported having never used HoloLens before.

4.2 Procedure

Participants were asked to search 32 individual network graphs (16 per condition) and count the number of triangles formed by exactly 3 nodes. Each graph consisted of 40 nodes with 45-55 edges and 1 to 4 triangles. There were 4 types of graphs that were evenly dispersed across both conditions. These types differed in how they were initially presented to the participants.

- 1: Graph completely on the CAVE2 screens
- 2: Graph on and extends above the CAVE2 screens
- 3: Graph on and extends below the CAVE2 screens
- 4: Graph on and extends above and below the CAVE2 screens

As seen in Figure 3, a type 4 graph (along with types 2 and 3) require the participant to use either the wand to navigate or use the HoloLens to view the rest of the graph to continue the search.



Figure 3: A Type 4 graph in CAVE2.

The control condition of the study uses CAVE2 in its typical configuration: lightweight passive stereoscopic glasses and a PS3 Navigation controller (wand). Both the glasses and wand were retro-fitted with retro-reflective markers to be accurately tracked using the CAVE2's 14 infrared optical tracking system. The experimental condition replaces the glasses with the HoloLens headset equipped with retro-reflective markers for head tracking. The HoloLens's native tracking was disabled and replaced with the CAVE2's tracking information gathered from the markers. This was done to minimize differences between tracking in the two conditions as well as aiding

in the alignment of the CAVE2 and HoloLens displays. In both conditions participants were able to use the wand to virtually navigate the scene by translating or rotating the viewer around the virtual scene. Half of the participants were given the control condition first then the experimental while the other half did the reverse. After each condition participants filled out a survey and answered semi-structured interview questions to evaluate task difficulty, presence, and comfort. All of the head tracking (position and orientation), wand interactions, and event states such as graph position/rotation were all logged. The entire session took an average of 48.3 minutes to complete with a 4.2 minute break between conditions for the survey and interview.

4.3 Metrics

The independent variables in this study are the two device setups: HoloLens and non-HoloLens. The dependent variables of interest are completion time, accuracy of the task (number of correct triangle counts), head gaze direction (on CAVE2 screens or above/below screens), and physical movement (user movement) versus virtual movement (using the wand to move/rotate scene).

4.4 Hypothesis

We would expect to see more physical movement of the user wearing the HMD than the control condition (H1) as they would use the additional FOR in the HMD to physically explore the virtual space. Given that the user would be able to look around instead of moving the scene around we would expect more virtual navigation without the HMD (H2). We also expect to see more head rotation above or below the CAVE2 displays with the HMD than in the control condition (H3) due to the increased FOR with the HMD.

5 ANALYSIS

5.1 Movement Logs

For examining the within-subject logs of the participants, we used paired (dependent) t-tests with an alpha value of 0.01. Using the two conditions CAVE2 (C2) and CAVE2 + HoloLens (C2H) as the independent variables and the physical movement of the participant from the head position distance in meters as the dependent variable, using the HoloLens had no significant effect ($M = 11.22$, $SD = 7.9$) compared to the CAVE2 condition ($M = 11.9$, $SD = 10.7$); $t(160) = 119.85$, $p = 0.361$, $d = 0.07$. These results do not support H1.

For H2 we expected that with the HoloLens, participants will use the wand to virtually navigate less since the HoloLens allows them to see more of the scene without translating or rotating the scene. Participants were allowed to use wand navigation exactly the same in both conditions. Examining the wand rotation and translation input logs we found a significant difference in navigation distance in meters while using the wand in the C2 condition ($M = 19.56$, $SD = 30.07$) and CAVE2 + HoloLens ($M = 11.38$, $SD = 14.7$); $t(160) = 363.48$, $p = 0.001$, $d = 0.301$ thus supporting H2.

Given the additional FOR provided by the HoloLens we found a significant difference in the number of seconds participants looked above or below the CAVE2 displays in the CAVE2 + HoloLens ($M = 13.67$, $SD = 18.0$) compared to the CAVE2 condition ($M = 1.17$, $SD = 3.12$); $t(160) = 237.18$, $p = 0.001$, $d = -0.67$. This supports H3 indicating that participants did use the HoloLens as part of their

search opposed to using just virtual navigation to bring the graph on the CAVE2 displays.

The total gaze distance calculated as a ray drawn on the display surface also indicated that there was significantly less head rotation in the CAVE2 + HoloLens condition ($M = 53.03$, $SD = 29.1$) compared to the CAVE2 condition ($M = 68.69$, $SD = 60.4$); $t(160) = 701.0$, $p = 0.001$, $d = 0.26$. Based on interview responses this was due to less looking back and forth as the HoloLens FOR improved their spatial awareness while examining the graph.

5.2 Task Logs

Using the same analysis for the movement logs, we examined the completion times in seconds and accuracy of each task and found no significance in the task speed or accuracy, although the mean scores in the CAVE2 + HoloLens condition are lower. 3 participants in the CAVE2 condition had perfect scores. 1 participant had a perfect score on both conditions.

While we attempted to mitigate the effects of learning by flipping the starting condition (CAVE2 or CAVE2 + HoloLens) for every other participant, we also conducted an analysis of the completion times comparing when a condition was first versus second for both conditions. For the CAVE2 + HoloLens condition (C2H) we found no significant difference whether it was the first ($M = 66.00$, $SD = 45.90$) or the second ($M = 68.84$, $SD = 28.62$) one presented; $t(80) = 331.10$, $p = 0.495$, $d = -0.077$. We did find a significant difference in the CAVE2 only condition: ($M = 77.87$, $SD = 30.11$) if presented first and ($M = 50.48$, $SD = 33.09$) if presented second; $t(80) = 350.56$, $p = 0.001$, $d = 0.699$. We found that completing the CAVE2 only condition second showed a significantly lower score than if the CAVE2 condition was first; $t(80) = 17.43$, $p = 0.001$, $d = 0.389$. We believe this may indicate that participants were rushing to complete the second task after using the HoloLens.

5.3 Survey Responses

A 12 question survey was given to participants after each condition. These included task difficulty questions from the NASA Task Load Index (1 - low, 5 - high) [6], presence in virtual environments from the Igroup Presence Questionnaire [16], and additional questions on comfort and sickness (1 - strongly disagree, 5 - strongly agree). The questions were analyzed using paired t-tests and generally found no significant differences in the CAVE2 or CAVE2 + HoloLens conditions. The one response with statistical significance was the decreased comfort of using the HoloLens ($M = 2.8$, $SD = 1.4$) compared to the CAVE2 glasses ($M = 4.4$, $SD = 0.91$); $t(10) = 4.68$, $p = 0.02$, $d = -1.185$. This was reinforced by 4 participants who specifically mentioned this in the post-condition interview.

5.4 Interview Responses

After each condition, semi-structured interviews were conducted to gain additional insight into each participant's experiences. In terms of search strategies, participants would generally start at one end of the graph and scan across for the target triangles. A participant reported using the borders of CAVE2's tiled displays as a search grid. 3 of the participants reported that they found the HoloLens useful for maintaining spatial awareness by using the HoloLens to look around rather than virtually navigating with the

wand. This is also reflected in the significantly less head gaze found in the movement logs as participants spent less time looking back and forth. 3 participants specifically reported that they didn't find the HoloLens useful at all citing heaviness of the headset, display update lag, and hard to look around with the headset. Based on video recordings from the HoloLens, the latency between the end of head movement and the HoloLens display alignment with the CAVE2 displays was 375 ms (125 min, 667 max). One of these participants also reported that transitioning from looking at the CAVE2 tiled displays, to the HoloLens was distracting because the tiled display borders would be gone after looking above the displays. While this participant noted that perceptually the borders disappeared while using just the CAVE2 displays (as is commonly experienced in CAVE2), the transition of seeing only the HoloLens display above the tiled display wall made the borders obvious again.

6 CONCLUSION

This study was part of exploratory study into the fusion of high-end, room-sized hybrid-reality environments such as CAVE2 with less expensive AR HMDs. Our preliminary results suggest that this setup promotes more physical navigation, less head movement, and higher head gaze outside the CAVE2 displays - even in participants who stated they found the HoloLens not useful at all. Physical navigation in large display walls have shown to improve user performance [1]. While our study did not show a significant change in performance, we found it encouraging that immersion and presence was not significantly different in our setup even with the uncomfortable nature of HoloLens and the amount of latency in our framework. Both of these limitations we expect to be improved with future models of AR HMDs and our framework.

A surprising finding was that the lack of display borders when visually transitioning from the tiled display wall to the HoloLens was distracting. Given that such borders are generally perceived as a disadvantage in immersive walls, the fact that they were useful as a search grid on the CAVE2 displays and disruptive when they were gone on the HoloLens display would seem that introducing virtual bezels on the AR display might be beneficial.

6.1 Future Work

Continuing to build on the framework described in this paper, we aim to explore these ideas in a more collaborative environment where two or more participants will work with a more complex visualization using AR HMDs or tablets to augment the large stereoscopic display environment. In addition to the extending of the displays, we also want to explore into using the AR device to display private information in before sharing it on the large display and what kind of user interface and interactions would promote using an AR HMD with a large display when one or both displays are the interaction target.

ACKNOWLEDGMENTS

We thank Dr. Debaleena Chattopadhyay and students in her Empirical Methods class for their feedback in the design of the user study. This work used the CAVE2 hybrid-reality environment which was supported by the National Science Foundation (award # CNS-0959053) and the Department of Energy (award # DE-SC005067).

REFERENCES

- [1] Robert Ball, Chris North, and Doug A. Bowman. 2007. Move to improve: promoting physical navigation to increase user performance with large displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*. ACM Press, 191. <https://doi.org/10.1145/1240624.1240656>
- [2] Doug A Bowman, Tobias Höllerer, Cha Lee, Ryan P McMahan, and Regis Kopper. 2012. Evaluating effectiveness in virtual environments with MR simulation. (2012), 11.
- [3] Marco Cavallo, Mishal Dholakia, Matous Havlena, Kenneth Ocheltree, and Mark Podlaseck. 2019. Dataspace: A Reconfigurable Hybrid Reality Environment for Collaborative Information Analysis. *ArXiv abs/1903.03700* (2019).
- [4] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriott, and B. H. Thomas. 2017. Immersive Collaborative Analysis of Network Connectivity: CAVE-style or Head-Mounted Display? *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 441–450. <https://doi.org/10.1109/TVCG.2016.2599107>
- [5] Alessandro Febretti, Arthur Nishimoto, Terrance Thigpen, Jonas Talandis, Lance Long, J. D. Pirtle, Tom Peterka, Alan Verlo, Maxine Brown, and Dana Plepys. 2013. CAVE2: a hybrid reality environment for immersive simulation and information analysis. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 864903–864903. <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1660803>
- [6] Sandra G. Hart and Lowell E. Staveland. 1988. *Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research*. Human Mental Workload, Vol. 52. North-Holland, 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9)
- [7] Hans-Christian Jetter, Johannes Schöning, Roman Rädle3 Harald Reiterer, and Yvonne Rogers. 2013. Collaborative Interactions in Future Crisis Rooms. In *Submission to ECML Workshop*. http://hci.uni-konstanz.de/downloads/bigwallhci_jetter_et_al.pdf
- [8] M. Klapperstuck, T. Czauderna, C. Goncu, J. Glowacki, T. Dwyer, F. Schreiber, and K. Marriott. 2016. ContextuWall: Peer Collaboration Using (Large) Displays. In *2016 Big Data Visual Analytics (BDVA)*. 1–8. <https://doi.org/10.1109/BDVA.2016.7787047>
- [9] Jason Leigh, Andrew Johnson, Luc Renambot, Tom Peterka, Byungil Jeong, Daniel J. Sandin, Jonas Talandis, Ratko Jagodic, Sungwon Nam, Hyejung Hur, and et al. 2013. Scalable Resolution Display Walls. *Proc. IEEE* 101, 1 (Jan 2013), 115–129. <https://doi.org/10.1109/JPROC.2012.2191609>
- [10] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, and Eric Lecolinet. 2016. Shared Interaction on a Wall-Sized Display in a Data Manipulation Task. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, 2075–2086. <https://doi.org/10.1145/2858036.2858039>
- [11] Thomas Marrinan, Jillian Aurisano, Arthur Nishimoto, Krishna Bharadwaj, Victor Mateevitsi, Luc Renambot, Lance Long, Andrew Johnson, and Jason Leigh. 2014. SAGE2: A new approach for data intensive collaboration using Scalable Resolution Shared Displays. In *Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on*. IEEE, 177–186. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7014563
- [12] Mahdi Nabiyouni, Siroberto Scerbo, Doug A. Bowman, and Tobias Höllerer. 2017. Relative Effects of Real-world and Virtual-World Latency on an Augmented Reality Training Task: An AR Simulation Experiment. *Frontiers in ICT* 3 (2017). <https://doi.org/10.3389/fict.2016.00034>
- [13] K. Nagao, Y. Ye, C. Wang, I. Fujishiro, and K. Ma. 2016. Enabling interactive scientific data visualization and analysis with see-through hmds and a large tiled display. In *2016 Workshop on Immersive Analytics (IA)*. 1–6. <https://doi.org/10.1109/IMMERSIVE.2016.7932374>
- [14] Markus Rittenbruch. 2014. Evaluating the Use of a Very Large-scale Presentation and Collaboration Framework. In *Proceedings of The International Symposium on Pervasive Displays - PerDis '14*. ACM Press, 124–129. <https://doi.org/10.1145/2611009.2611023>
- [15] M. Rodrigue, A. Waranis, T. Wood, and T. Höllerer. 2015. Mixed reality simulation with physical mobile display devices. In *2015 IEEE Virtual Reality (VR)*. 105–110. <https://doi.org/10.1109/VR.2015.7223331>
- [16] Thomas Schubert, Frank Friedmann, and Holger Regenbrecht. 2001. The Experience of Presence: Factor Analytic Insights. *Presence* 10 (Jun 2001), 266–281. <https://doi.org/10.1162/105474601300343603>
- [17] Jürgen P. Schulze, Andrew Prudhomme, Philip Weber, and Thomas A. DeFanti. 2013. CalVR: an advanced open source virtual reality software framework. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 864902–864902. <http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1660802>