JOHNS HOPKINS APL
TECHNICAL DIGEST
Volume 35, Number 3

Exploring Immersive Technology at APL

FOLLOW THESE 4 STEPS TO LAUNCH THE XR EXPERIENCE:

1. Print this page for the optimal experience.
2. Using Safari on an iOS device or Chrome or Safari on an Android device, go to https://www.jhuapl.edu/techdigest/XRapp.
3. Point your device's camera at the black "Tech" marker below.
4. Have fun! Use the buttons to view the 3-D models and learn more about the associated articles.

Tech

Having trouble? Make sure the marker sits flat and is evenly lit. Try removing the marker entirely and bringing it back into the frame. Distance the camera at least 12 inches from the marker. If you're tracking the marker from your screen rather than from this printed page, toggle the rotation button on the top left to reorient the models. *Have feedback?* Email us at TechnicalDigest@jhuapl.edu.

Simulated X-Ray Vision Using Mixed Reality

Stephen A. Bailey, Miguel A. Rufino, Rodrigo-Rene R. Munoz-Abujder, and Hirsh R. Goldberg

AR VR MR

ABSTRACT

A Johns Hopkins University Applied Physics Laboratory (APL) team created an application for the Microsoft HoloLens, a mixed reality (MR) head-mounted display (HMD), that serves as a proof of concept for a capability that would allow warfighters to observe their surroundings beyond their immediate line of sight. The approach uses a high-fidelity 3-D reconstruction of the beyond-line-of-sight (BLOS) environment in the form of a point cloud based on data from remote sensors and overlays it with the physical surfaces in the user's surroundings as seen through the HMD. This approach allows the user to observe areas beyond their line of sight without ever physically occupying or directly observing the space.

INTRODUCTION

What if we could provide our warfighters with the ability to see through physical surfaces in their immediate surroundings? The effect would be like Superman's x-ray vision, which enables him to see through surfaces at will. Warfighters could know the precise location of enemy and friendly assets, move through their surroundings without surprises, and share a common operating picture with other friendly forces. Mixed reality (MR) head-mounted displays (HMDs) have the potential to provide this capability, and the authors have created an application for the Microsoft HoloLens to serve as a proof of concept.

Sensors on the HoloLens allow the device to build a spatial mapping mesh representing the physical surfaces surrounding the device.[1] This mesh provides the inspiration for how the HMD could be used to simulate x-ray vision for warfighters. This 3-D reconstruction of the HMD's surroundings is aligned and overlaid with its real-world counterpart when viewed through the HMD. The mesh expands to include more surfaces as the user moves through their surroundings and can grow to include several distinct areas such as individual rooms in a building. Visualizing the spatial mesh for each distinct area would allow the user to observe the physical surfaces present in a nearby area without actually occupying or directly observing that space.

When visualized, the spatial mesh can have sufficient fidelity to allow the user to determine the contents of the space. The mesh in Figure 1 shows a room with a chair with a small table to the left and stairs to the right. The authors reasoned that a higher-fidelity 3-D reconstruction of the environment could be collected by a different sensor platform and used for visualization instead of the spatial mesh produced by the HoloLens. This would enable the user to observe spaces they have never occupied since the sensor collecting the data does not need to be colocated with the HMD. The data can be displayed in the HMD as an ordinary hologram and overlaid with the physical surfaces in the user's surroundings in a manner similar to the spatial mesh. Once overlaid, the hologram can be configured to always be visible regardless of the location of the HMD, providing the x-ray vision effect.

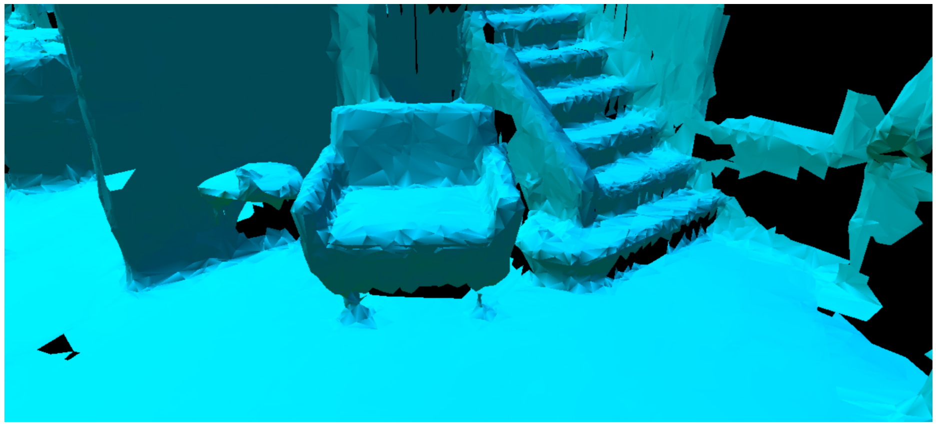

Figure 1. Spatial mapping mesh visualized. The mesh shows a room with a chair with a small table to the left and stairs to the right.

HIGH-LEVEL APPROACH

The team's approach to simulated x-ray vision involves several distinct steps. First data are collected from the warfighters' surroundings via remote sensor. Some amount of processing, filtering, and characterization of the data is performed. The data are then sent to the HMD and displayed so that they are overlaid with the user's surroundings. An additional calibration step establishes a common frame of reference between the sensor collecting data and the HMD. This is needed to accurately map sensor data into the HMD frame of reference. Once a common frame of reference is established, the sensor and HMD must keep track of any pose changes to maintain alignment.

Although the spatial mesh served as inspiration for this work, humans can often readily interpret visualizations of high-fidelity point cloud data without the extra overhead needed to convert the points to a mesh. The authors made use of preprocessed lidar data covering the APL campus and real-time sources of point cloud data from lidar and depth-sensing stereo cameras in place of a mesh for visualizing the beyond-line-of-sight (BLOS) environment surrounding the HMD.

When working with real-time sources of data, fiducial markers present on the sensor provided the means to establish a common frame of reference between the sensor and the HMD. When working with a preprocessed, georectified data set, such as the one shown in Figure 2, alignment was achieved through the use of preset calibration locations. This involves the user standing at a certain physical location in their surroundings and running the calibration algorithm. The calibration process is an area of ongoing research.

The Unity game engine served as the development platform for creating the software application for the HoloLens.[2] Robot Operating System (ROS)[3] was used

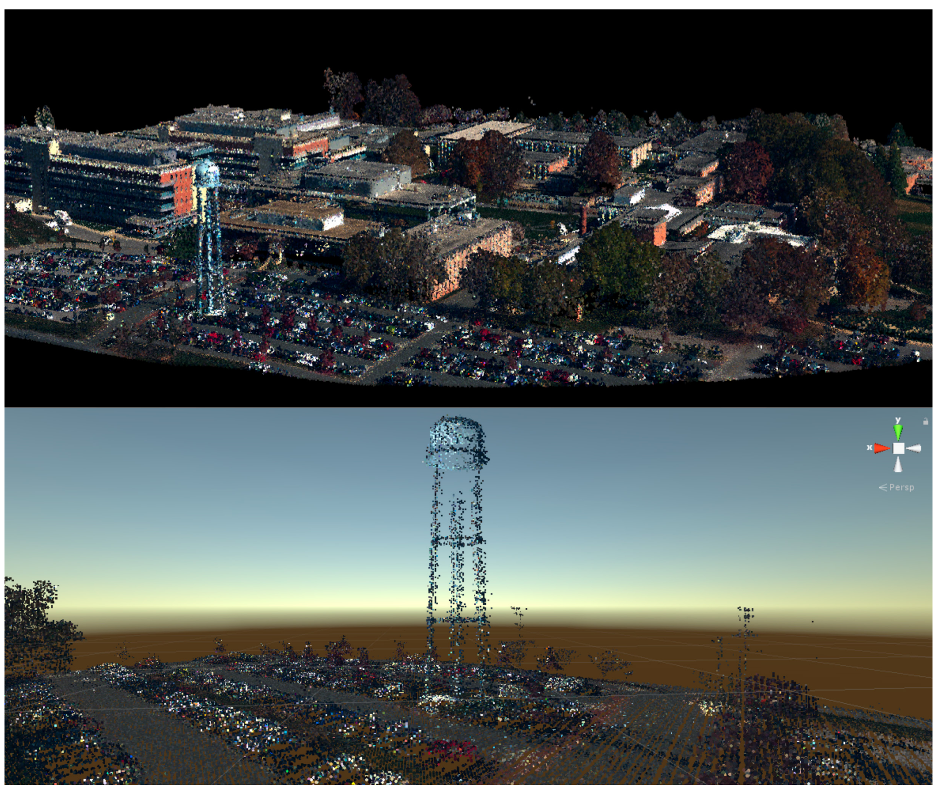

Figure 2. APL campus lidar point cloud. Top, APL campus point cloud visualized from above. Bottom, APL water tower and surrounding data viewed in Unity.[3]

to create a data pipeline supporting the collection, filtering, classification, packaging, and delivery of sensor data to the HMD. ROS runs on both the HMD and the compute resources used by the remote sensor to support the data pipeline. Mobile ad hoc network (MANET)-forming radios such as the Silvus StreamCaster[4] have been successfully used to send data between a remote sensor and the HMD. During testing, a variety of small, inexpensive Wi-Fi routers have also been used.

SIMULATING X-RAY VISION IN MR

Point cloud data sets such as the one shown in Figure 2 may be too large to be displayed inside the HoloLens at once. In these scenarios, the data set was split into spatially tiled regions to be displayed one at a time, and the user had the ability to select which tiled region was displayed. This may be achieved algorithmically or through the use of an input device such as a controller. Figure 3 shows a lidar tile containing the APL water tower overlaid with the real water tower viewed through an office window. This is an example of displaying a preprocessed, georectified data set inside the HMD.

Figure 4 shows real-time lidar streaming and alignment. The left side of the image shows the lidar point cloud overlaid with its real-world counterpart, and the right side shows the same data viewed from behind a wall, demonstrating the simulated x-ray vision. Points closer to the sensor have been colored green and those farther away have been colored red. The Ouster OS1-64[5] Lidar used in Figure 4 produces too much data to be displayed efficiently in the HoloLens, so voxel filtering was applied to the data before delivery to the HoloLens. Voxel filtering has the effect of averaging together nearby points to produce a point cloud that has lower fidelity but is easier to display for the HMD. Using this approach, a lower-fidelity version of the data collected by the sensor can be displayed in the HMD at once, but as shown in Figure 4, the display includes many points, corresponding to office furniture, ceiling, and walls, that clutter the view and may not be of interest to the user. Therefore, it may be desirable to avoid showing those points inside the HMD and instead focus only on the green silhouette representing a person, for example, or other objects of interest.

Figure 5 shows an alternative approach that only displays certain portions of the point cloud inside the HMD without voxel filtering. This figure still shows point cloud data overlaid with its real-world counterpart on the left and a similar image viewed from behind a wall on the right. This demonstration makes use of the ZED[6] depth-sensing stereo camera, which can produce a

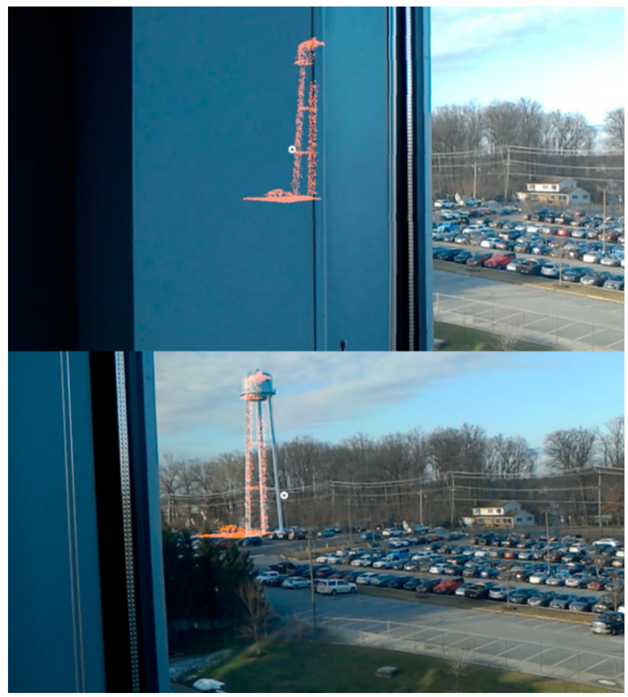

Figure 3. Display of a preprocessed, georectified data set inside the HMD. Top, A lidar point cloud containing the APL water tower visible through a wall. Bottom, The same lidar point cloud aligned with the actual water tower when the user moves to look through window. Lidar data are shown in orange for increased visibility.

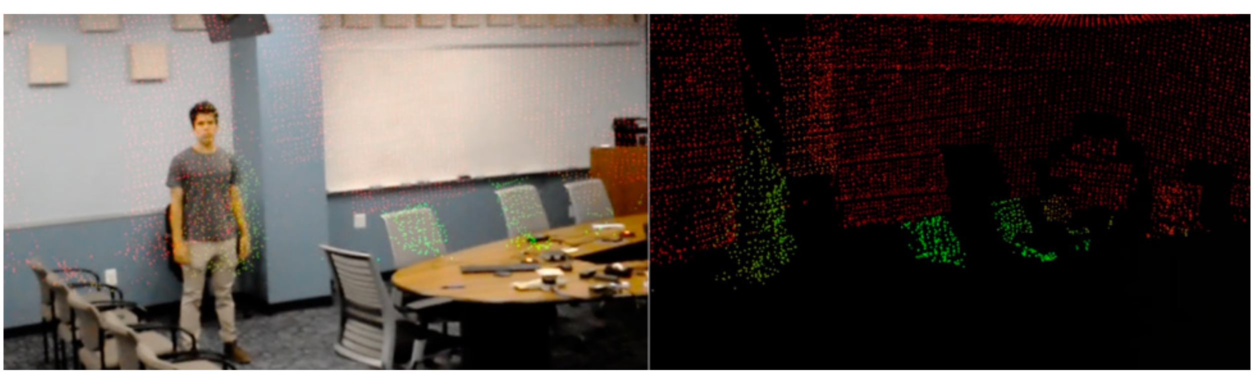

Figure 4. Real-time lidar streaming and alignment. Left, Voxel-filtered lidar data collected and overlaid with real-world counterparts in near real time viewed through the HoloLens. Right, Data viewed through the HoloLens from behind a wall, illustrating the simulated x-ray vision. Points closer to the sensor are colored green; the color shifts to red with distance from the sensor.

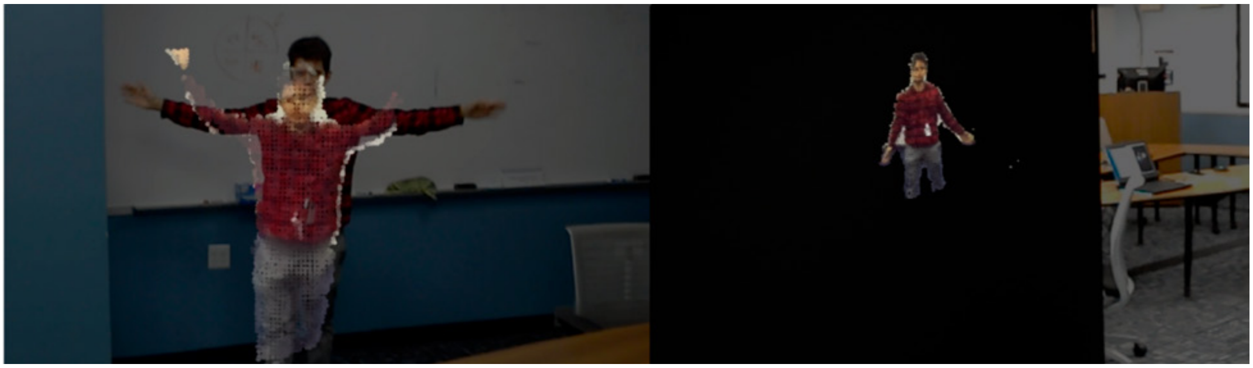

Figure 5. An alternative approach that only displays certain portions of the point cloud inside the HMD without voxel filtering. Left, A semantically segmented point cloud overlaid with its real-world counterpart in near real time viewed through the HoloLens. The slight alignment mismatch is due to data latency. Right, Semantically segmented data viewed from behind a wall, illustrating simulated x-ray vision. Accurate color information has been applied to each point in the point cloud.

point cloud with accurate color information, instead of a lidar. Using machine learning–based computer vision techniques, the point cloud was semantically segmented to extract the points associated with the person present in the scene. Using this approach, significantly fewer data points are sent to the HMD since the segmentation is applied before any data are sent. The decreased quantity of data makes techniques like voxel filtering unnecessary and enables a higher-fidelity point cloud to be displayed in the HMD. However, the semantic segmentation capability adds another step in the data processing pipeline, which negatively impacts the rate at which data can be sent to the HMD.

FUTURE DIRECTION

Initial progress capturing, filtering, segmenting, aligning, and displaying data in a commercial off-the-shelf HMD has shown promise. The team has successfully demonstrated that point cloud data can be displayed inside an MR HMD in a manner that is readily interpretable by the user, and that a data pipeline can be established to enable the capability in near real time. Aligning the data with the real world is one of the most challenging aspects of this work. Using fiducial markers works well, but over time tracking errors accumulate in both the HMD and the data source. Currently, periodic recalibration is necessary to maintain alignment with the real world, and near-term efforts are seeking to reduce or eliminate this need.

The manner in which the data are displayed inside the HMD is critical to their interpretability by the end user. While additional processing of data allows for more refined imagery to be delivered to the HMD, it also requires more compute resources. Spending additional time refining imagery also means data are delivered to the HMD at a lower rate, which can detract from the accuracy of the imagery being displayed to the end user. Given the real-time nature of the applications for this work, it is critical to strike the right balance between the level of refinement and compute resources necessary to refine the data before sending them to the HMD. The team continually seeks to incorporate end-user feedback to improve the enhanced situational awareness provided by this capability.

This project is being continued under direct sponsor funding in combination with other APL work in the areas of robotics, human–machine teaming, machine learning, and imaging systems.

REFERENCES

[1]"Spatial Mapping." Microsoft. https://docs.microsoft.com/en-us/windows/mixed-reality/spatial-mapping (accessed Feb. 1, 2019).
[2]"Unity." https://unity.com/ (accessed Feb. 1, 2019).
[3]"ROS." https://www.ros.org/ (accessed Feb. 1, 2019).
[4]"StreamCaster Radios." Silvus Technologies. https://silvustechnologies.com/products/streamcaster-radios/ (accessed Feb. 1, 2019).
[5]"Ouster." https://ouster.com/products/os1-lidar-sensor/ (accessed Feb. 1, 2019).
[6]"ZED." Stereolabs. https://www.stereolabs.com/zed/ (accessed Feb. 1, 2019).

Stephen A. Bailey, Air and Missile Defense Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Stephen A. Bailey is a software engineer in APL's Air and Missile Defense Sector. He has a BS in premedical studies from Penn State and an MS in computer science from Johns Hopkins University. He has a diverse set of skills including full-stack web applications development, virtual/mixed reality devices and development, and Unity game engine development. He has worked as the principal investigator for multiple independent research and development (IRAD) projects at APL and serves on an APL committee that funds innovative new ideas across the Air and Missile Defense Sector. His email address is stephen.bailey@jhuapl.edu.

Miguel A. Rufino, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Miguel A. Rufino is a robotics research engineer in APL's Force Projection Sector. He holds a BS and an MS in electrical engineering, both from North Carolina State University. During his time at APL Miguel has worked in a variety of projects ranging from simultaneous localization and mapping (SLAM) and GPS-denied navigation to unmanned vehicle autonomy. His current research interests are distributed localization and mixed reality. His email address is miguel.rufino@jhuapl.edu.

Rodrigo Rene Rai Munoz Abujder, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Rai Munoz is a computer vision research engineer in the Advanced Algorithms Section of APL's Imaging Systems Group. He holds a bachelor of engineering in electrical engineering from the City College of New York and an MS in electrical and computer engineering from Carnegie Mellon University. Rai specializes in deep learning, reinforcement learning, robotics, and computer vision. He has worked on a variety of projects involving steganography, super resolution, and simultaneous localization and mapping (SLAM) improvised explosive device detection. Before joining APL full-time, he worked as a summer intern at APL and in various NASA centers on a variety of projects ranging from space biosciences to AR/VR applications. His email address is rai.munoz@jhuapl.edu.

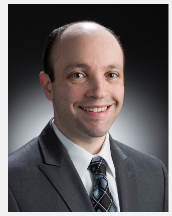

Hirsh R. Goldberg, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Hirsh R. Goldberg is a project manager and the chief engineer for the Imaging Systems Group in APL's Asymmetric Operations Sector. He has a BS and an MS in electrical engineering, both from the University of Maryland, College Park. He has experience in electro-optical/infrared (EO/IR) system technology development; computer vision/machine learning; automatic target recognition; system architecture design; modeling and simulation for ground-based, airborne, and satellite-based remote sensing applications; and lidar imaging systems. Hirsh is the project manager for an Army advanced lidar sensor development program; has supported the development, operation, and analysis of multiple airborne lidar systems; and is currently the principal investigator for independent research and development projects focused on real-time lidar-based 3-D collection and processing on small unmanned aerial vehicle platforms. His email address is hirsh.goldberg@jhuapl.edu.