Luna: An AR Musical Instrument on the Meta Quest 2

Samuel Dietz
Australian National University
Canberra, Australia
samuel.dietz@anu.edu.au

Charles Patrick Martin
Australian National University
Canberra, Australia
charles.martin@anu.edu.au

Abstract

Head-mounted augmented reality (AR) computers present the opportunity to develop new musical interfaces that would be impossible to build physically or with conventional computing devices. Unfortunately, typical computer music tools have not been easy to apply within AR development tool chains. Integrating standard computer music tools in AR development would allow more rapid prototyping of new instrument ideas and transfer of knowledge from experienced computer musicians. The goal of this paper is to demonstrate that AR digital musical instruments can be developed using libpd, the library version of the standard computer music environment Pure Data. We present a case study of an AR instrument developed for the Meta Quest 2 integrating libpd in the AR development tool-chain for the interactive audio components. The iterative development process was tracked through autoethnographic reflections and analysed with thematic analysis. We found that Pure Data was an effective way to develop audio interactions on the Quest 2 and that the hand tracking on this platform was capable of complex gestural interactions. This work could enable a broader community of computer musicians to explore AR NIME development, taking advantage of the unique affordances of this medium.

Keywords

Augmented Reality, DMI, Pure Data, hand tracking

1 Introduction

Virtual and augmented reality devices unlock the capability to pass the limitations of the real world and create new forms of expression [10]. Augmented reality (AR) musical instruments can blend real-world elements with virtual ones. This allows performers to integrate their musical setup with any environment without removing themselves from their audience, while also removing the physical barriers associated with real-world instruments, as demonstrated by current AR instruments [15].

Present exploration of these instruments for musical applications is limited. Devices such as the Microsoft HoloLens [15] being expensive and the open-source Project North Star headset are complicated to construct [3]. Less expensive consumer-grade VR devices now support AR and hand-tracking through cameras [1]. It may be that more affordable devices can remove many of the barriers associated with AR musical instruments leading to a wider variety of DMI designs and artistic explorations.

This paper describes a case-study DMI for a more affordable self-contained VR headset, the Meta Quest 2. The goal of the research is to demonstrate the feasibility of this process and to study the capability of the hardware and development workflow in terms of creating an effective audio experience with free-hand gestural interaction. Our AR instrument, Luna, shown in Figure 1, uses libpd for audio processing (on the Meta Quest device), uses hand-tracking for free-hand interaction with virtual objects in real space. Luna was evaluated through an autoethnographic iterative design process. This involved the development of a proof of concept to demonstrate libpd [6] on the Meta Quest 2 hardware and set of three prototypes. Each prototype was tested with a 5 minute solo performance, followed by a reflection over each instrument's performance which were analysed with thematic analysis. We found that the Meta Quest 2's operating system, processing power and hand tracking is a viable solution for a standalone AR instrument with complex free-hand gestural interaction. Using libpd and Pure Data to create audio interactions on this platform can enable wider explorations of AR NIME designs.

Figure 1: Luna, an AR musical instrument created using Unity and Pure Data running on a Meta Quest 2. This instrument uses free-hand interaction to generate objects representing continuous synthesis processes.

2 Related Work

A variety of recent research has examined how virtual or augmented reality can be applied to musical tasks such as providing a concert experience [7], modular synthesis [14], laptop ensemble performance [2], and creating "impossible instruments" [17]. Boem et al. recently explored the potential for musicians to use networked AR/VR technologies for education, performance and composition [4]. The possibilities of augmented reality in

NIME '25, June 24–27, 2025, Canberra, Australia

particular are articulated throught the concept of "Hyperreal-Instruments" [17] that by incorporating real-world objects into virtual instruments, can "defy our sense of audiovisual reality" while satisfying our "haptic expectations" [17]. While many AR and VR systems involve a tethered computer, newer standalone AR or VR headsets can have advantages in movement freedom, interaction and portability [12]. Standalone head-mounted AR computers such as the HoloLens 1 and 2 have been used for sonic interventions with sculptures [9] and free-hand musical performance [16].

The tools required for creating a VR or AR musical instrument vary depending on device. For tethered VR, the computer music software can operate on a regular computer, e.g., The Hyperreal Instruments use Max for audio with Unity for graphics with an HTC Vive VR headset [17]. An AR instrument, "polaris̃" involves a tethered Project North Star AR headset, Unity, and Pure Data for audio [3]. For non-tethered head-mounted systems, workflow choices have been limited by platform which is generally Arm architecture with Windows on Arm, Android, or Apple's Vision OS. The HoloLens uses Windows on Arm and researchers have applied the Disunity Toolkit [13] with Unity to create self-contained DMIs [15]. This approach was demonstrated to work practically, it did not allow authoring audio interactions in standard software such as Pure Data or Max. While Pure Data and libpd [6] have been successfully applied in tethered VR arrangements [3], they have not so far been demonstrated within self-contained Android AR headsets. In this research, the Meta Quest 2, a relatively affordable Arm Android based VR headset is the target for development.

3 System Design

Following a proof-of-concept showing that libpd on the Meta Quest was feasible, Luna was developed in three iterations. Each iteration was explored through solo performance and reflection to determine the direction of the next iteration.

3.1 Proof of Concept

We set an initial goal to create a proof of concept that would demonstrate use of Pure Data to develop an audio interaction locally on a Meta Quest 2 device. We had three goals for the proof of concept: 1) display an empty project through the Meta Quest 2 display, 2) Allow the user to have basic head movement throughout the scene, and 3) emit an audible 440hz tone using Pure Data. Unity was selected as the game engine for development as there was an existing project, LibPdIntegration, that promised to allow Pd integration within a Unity project that had been used in existing AR music projects [3]. Unfortunately, LibPdIntegration did not support the Android platform and had been abandoned. We forked the project, added additional binaries and modifications to add support for the Unity projects on Android platform allowing the criteria to be met on the Quest 2. Our fork of LibPdIntegration with Android support is available online here: https://github.com/Samson026/LibPdIntegrationForAndroid

3.2 Iteration 1

This iteration focused on implementing an initial patch and interface design, shown in Figure 2. A simple sphere represented the interface with a plain lit material and a simple trail particle effect. The synthesis engine was a frequency modulation synthesizer with modulation and index inputs. A pinch motion could be used to grab and hold the sphere within the scene. While grabbing the sphere, the movement of the sphere changed the index of modulation, while distance between the hands changed the modulation amount. Iteration 1 also implemented a spatial audio effect, so that audio appeared to be emitted directly from the sphere object.

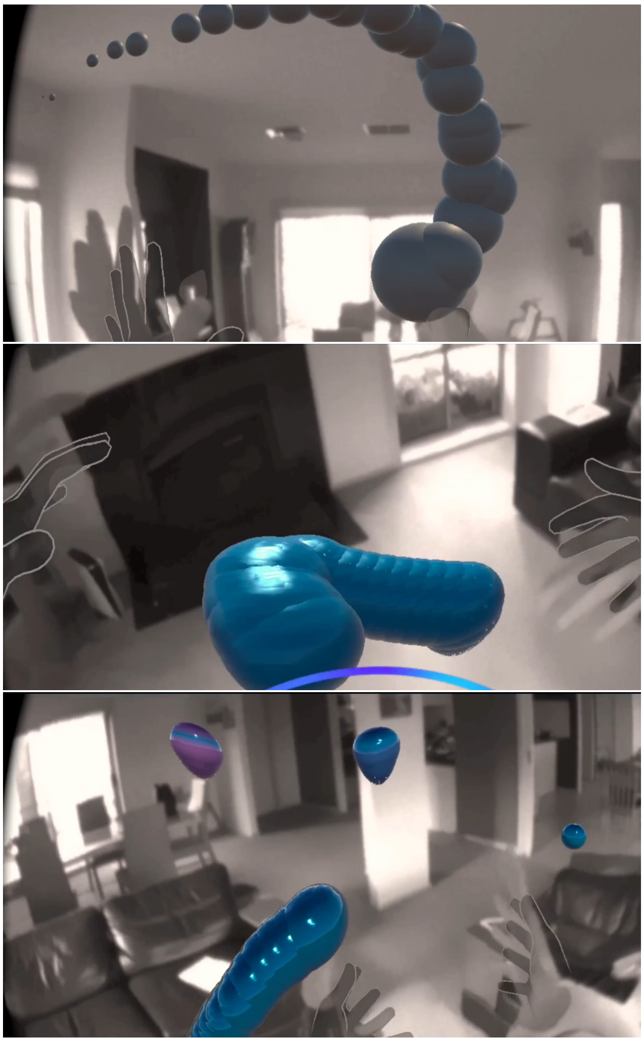

Figure 2: Iterations of Luna, top-to-bottom: 1, 2, 3. Iteration 1 introduced the free-hand interactive concept. Iteration 2 enabled multiple objects in the scene and refined the visual design. Iteration 3 refined the mapping from gesture to synth parameters and visual feedback.

3.3 Iteration 2

Iteration 2 focused on implementing support for multiple instances of one Pd sound source, more appealing visual appearance for the sphere and expanding user interaction. Our goal was to allow multiple spheres with different audio sources to be present around the user simultaneously. This required modification of the Pd binary included within Luna and the included wrapper within LibPdIntegration. A water-like material texture was added to the sphere shown in Figure 2 including dripping and trail particle effects. This iteration also expanded on how users interact with the FM synthesizer's base frequency, modulation, and index. To address tracking issues found in iteration 1, iteration 2 allowed users to attach a sphere that will follow midpoint between the user's two palms. Users could alter the parameters by changing the distance, angle, and position of their hands relative to a sphere.

3.4 Iteration 3

Iteration 3 was focused on improving and refining existing designs. As precise manipulations were difficult without physical controls, iteration 2 limited the range of possibilities allowing more intentional inputs to the device. The limitation presented by this change was offset by a much larger input stage. Remapping the inputs for index, modulation and amplitude modulation to distances from an origin on the x, y and z axis. This allowed Luna to provide both minor sound adjustments as well as more diverse outputs, as users could use their hands for small adjustments and move around the environment to make more significant impacts. This feature also encouraged moving around the environment to make significant sound changes, leading to more engagement between digital elements and the real-world setting. Visual feedback was added providing a better understanding of how inputs were manipulating the sound. The scale and colour of each sphere's material was changed to dynamically respond with the inputs given by the user. Finally, an ambient element to the sound was added so that once a sphere was detached a triangle wave envelope loops over each sphere changing its volume. This led to a more interesting sound experience from each sphere while a user is not interacting with it.

4 Results

The evaluation process in this project was autoethnographic [8] first-person research. As the developer was also the primary user of AR instrument, an autobiographical design process [11] was appropriate, with both the development process and final software outcome the subject of research. Each iteration of the software was tested in solo improvisation in a lounge room. These performances occurred throughout the development process were around 5 minutes and were followed by a written reflection on the experience. The data collected from these reflections was used for two purposes. Each reflection was used to set the direction of the next iteration as outlined in the section above. The data was subsequently subject to a process of thematic analysis [5]. This process led to three themes identified within the data discussed in the following sections. These sections are written in the first person from the perspective of the first author.

4.1 Interacting with Spatial Audio

The term *clear sound quality* was consistently mentioned throughout the reflection of each iteration. Each reflection stated that the sound quality was clear with no noticeable audio artifacts present throughout the performance. With iteration 1 running on Pd's single-instance mode and iteration 2 and 3 running in multi-instance mode with no limit to the number of instances a user could spawn into their environment, it can be inferred that there is no noticeable impact of running Pd in multi-instance mode on the Meta Quest 2 headset.

Another finding within both iteration 2 and 3's reflections were poor spatial audio quality and its impact on immersion. Iteration 2 implemented support for Pd's multi-instance features. This allowed users to spawn and manipulate different spheres throughout their environment. It was noted that the 3D audio present on the Quest did improve the immersion of the experience. While there were in some circumstances difficulties, I was able to distinguish the spheres from one another and it heightened the experience when walking around the environment.

4.2 Consistency and Free-hand Interaction

It was noted in the reflection of iteration 2 that a lack of visual feedback to inputs made by the user took away from the impact of smaller movements made through the interface. A lack of visual feedback created doubt if inputs were registered correctly by the device. This is highlighted on small inputs made by the user. These inputs can be hard to recognise and may lead users to rely on larger more exaggerated movements to obtain noticeable feedback from the device. Iteration 3 would increase the scale of the sphere and material of a sphere based on the user's spatial and gestural inputs. This had a positive influence on the instrument with iteration 3's reflection stating that colour and scale changes give a good visual feedback to the user on how their inputs are influencing the sound. This highlights the importance of user feedback for augmented reality instruments, and how providing feedback to the user via a combination of different senses can create a more intuitive form of interaction.

A takeaway of one of the limitations of the Meta Quest 2 was the limitations in hand tracking. The design of the first iteration involved grabbing and moving around a sphere by pinching fingers. Reflection 1 states that the sphere is dropped from the hand quite often, this was due to a loss in hand tracking. This note influenced the design of iteration 2 and 3 to toggle a sphere to follow the user's movements. Reflection 2 states that sound consistency was greatly improved and that attaching and detaching the sphere apposed from grabbing it solved the disconnection issue. This new approach introduced a new issue noted in reflection 2. It stated that the sphere teleports to 0,0,0 on loss of hand tracking. While hand tracking is re-established almost immediately, a loss of tracking would sent the bones used in tracking to the 0,0,0 position, sending unwanted inputs to the instrument. Both reflection 2 and 3 mentioned the hand tracking interfering with the instrument's pose detection.

Reflection 2 noted that hand poses were not consistent impacting the ability to interact with the instrument, so hand poses were simplified in iteration 3. However, iteration 3's pose usage was still affected by the Quest 2's hand tracking, with reflection 3 stating that tracking is not consistent and that pose detection went off unprovoked. This described the tracking issues on the Quest from two different angles. Both complex and simplified hand poses were heavily impacted by errors in the Quest's hand tracking either through miss fires or user difficulties to trigger poses. While a middle ground may exist, where neither of these issues have a major impact on the experience, I believe that this is unlikely. There was not a large difference between the original and updated hand pose, therefore the line between complex and simple may be too small for an optimal middle ground.

4.3 Real and Digital World Synthesis

An important element of AR applications is their ability to blend real world and digital elements into a single environment. It is evident through the reflections that the limitations of the Quest 2 have an impact on Luna's ability to create a convincing environment which includes both of these elements. The Quest 2 provides a gray scale video see-though AR experience. It was noted in reflection 1 that the gray scale background with the default lit material created an interesting atmosphere. The lack of colour and sounds produced from the instrument created a very alien experience. I believe this demonstrates that the gray scale pass though did not take away from the immersion of the instrument as a whole, but impacts the user's interaction with the outside world. This was further exaggerated in reflections 1 and 2 with improvements to the sphere's material. While it is noted that the improvements to the material had positive impacts on the instruments immersion, as the sphere was more realistic, I believe it brought more focus onto the spheres and away from the user's environment. While focus on the spheres within the instrument are positive, it further divides the real world and virtual elements within the scene.

With the changes to parameter control in iteration 3, values were set by modifying the x, y and z axis of the sphere, improvements to the real world and digital synthesis were noted. This change required users to make larger movements throughout the scene to have larger impacts on the instrument's sound. I found this encouraged me to walk around my environment and place spheres in different corners of the room. This improved the connection between the real and digital world. By creating a need to interact with the real world around them, it was much harder for users to become fully immersed in the digital world, greatly improving Luna's ability to blend both realities.

5 Conclusion

This research project has investigated the viability of an AR musical instrument running on the Meta Quest 2 through an autoethnographic iterative design process. This project focused on embedding Pd into Unity on the Quest 2 using the LibPdIntergration package, the aim being to investigate the viability of using a widely available and affordable headset in conjunction with a well known open-source audio library, to create a standalone AR musical instrument. Examining reflections on the iterations of our DMI design, we identified three themes, interacting with spatial audio, consistency and free-hand interaction, and real and digital world synthesis indicating successes in our design and areas where improvements might be made.

Our project identified and demonstrated a software development process for creating standalone AR musical instruments on the Meta Quest 2 with Pure Data. This would allow rapid prototyping of future AR instruments for this platform. Our evaluation made it clear that Pure Data was an appropriate an effective platform for AR audio development leading to interesting and stable sound experiences. The results suggested that the Quest 2's hand tracking can allow for complex gestural interactions; however, tracking errors have an impact on the user's experience. The free-hand user interaction scheme was refined iteratively throughout our development process demonstrating that effective free hand controls are possible although careful development and testing may be necessary.

Augmented reality musical instruments still appear to be under-explored despite the many affordances of these platforms and potential for creative applications. We have argued in this work that this may be due to a lack of typical computer music tools within this space. Our work has introduced a practical workflow for integrating Pure Data computer music development into the process of creating a Meta Quest 2 AR application. Our evaluation confirmed the feasibility of this process and we identified themes that distill the concerns of a DMI developer and performer using the platform. Our work could make DMI prototyping and development easier on AR platforms, contributing to the potential for exploratory research in AR music performance.

6 Ethical Standards

This research was conducted within applicable ethical standards. As autobiographical research involving only the authors, this work did not require institutional review.

References

[1] Diar Abdlkarim, Massimiliano Di Luca, Poppy Aves, Sang-Hoon Yeo, R. Chris Miall, Peter Holland, and Joseph M. Galea. 2022. A Methodological Framework to Assess the Accuracy of Virtual Reality Hand-Tracking Systems: A case study with the Oculus Quest 2. *bioRxiv* (2022). https://doi.org/10.1101/2022.02.18.481001

[2] Jack Atherton and Ge Wang. 2020. Curating Perspectives: Incorporating Virtual Reality into Laptop Orchestra Performance. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 154–159. https://doi.org/10.5281/zenodo.4813290

[3] Sam Bilbow. 2022. Evaluating polaris~ - An Audiovisual Augmented Reality Experience Built on Open-Source Hardware and Software. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* https://doi.org/10.21428/92fbeb44.8abb9ce6

[4] Alberto Boem, Matteo Tomasetti, Alessio Gabriele, Agostino Di Scipio, and Luca Turchet. 2024. User needs in the Musical Metaverse: a case study with electroacoustic musicians. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, S M Astrid Bin and Courtney N. Reed (Eds.). 221–229. https://doi.org/10.5281/zenodo.13904838

[5] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. https://doi.org/10.1080/2159676X.2019.1628806

[6] Peter Brinkmann, Dan Wilcox, Tal Kirshboim, Richard Eakin, and Ryan Alexander. 2016. Libpd: Past, present, and future of embedding pure data. In *Proceedings of the 5th Pure Data Convention (Pd˜ Con), New York City.*

[7] Ciaran Frame. 2024. Concerts of the Future: Designing an interactive musical experience in VR. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, S M Astrid Bin and Courtney N. Reed (Eds.). 359–366. https://doi.org/10.5281/zenodo.13904880

[8] Annika Kaltenhauser, Evropi Stefanidi, and Johannes Schöning. 2024. Playing with Perspectives and Unveiling the Autoethnographic Kaleidoscope in HCI – A Literature Review of Autoethnographies. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24).* ACM, New York, NY, USA, Article 819, 20 pages. https://doi.org/10.1145/3613904.3642355

[9] Charles Patrick Martin, Zeruo Liu, Yichen Wang, Wennan He, and Henry Gardner. 2020. Sonic Sculpture: Activating Engagement with Head-Mounted Augmented Reality. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 39–42. https://doi.org/10.5281/zenodo.4813445

[10] Desmond Miles. 2023. The Art of VR: How Virtual Reality Is Redefining Creativity and Expression. https://productivityland.com/virtual-reality-is-redefining-creativity-and-expression/. [Accessed 19-04-2024].

[11] Carman Neustaedter and Phoebe Sengers. 2012. Autobiographical design in HCI research: designing and learning through use-it-yourself. In *Proceedings of the Designing Interactive Systems Conference* (Newcastle Upon Tyne, United Kingdom) *(DIS '12).* Association for Computing Machinery, New York, NY, USA, 514–523. https://doi.org/10.1145/2317956.2318034

[12] Nikola Rendevski, Diana Trajcevska, Mario Dimovski, Konstantin Veljanovski, Aleksandar Popov, Naile Emini, and Dimitar Veljanovski. 2022. PC VR vs Standalone VR Fully-Immersive Applications: History, Technical Aspects and Performance. In *2022 57th International Scientific Conference on Information, Communication and Energy Systems and Technologies (ICEST).* 1–4. https://doi.org/10.1109/ICEST55168.2022.9828656

[13] Andrew Sorensen. 2020. Disunity Synthesis Toolkit. https://github.com/digego/DisunityST. GitHub Open Source Code Repository.

[14] Graham Wakefield, Michael Palumbo, and Alexander Zonta. 2020. Affordances and Constraints of Modular Synthesis in Virtual Reality. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Romain Michon and Franziska Schroeder (Eds.). Birmingham City University, Birmingham, UK, 547–550. https://doi.org/10.5281/zenodo.4813182

[15] Yichen Wang and Charles Martin. 2022. Cubing Sound: Designing a NIME for Head-mounted Augmented Reality. In *Proceedings of the International Conference on New Interfaces for Musical Expression.* Article 27. https://doi.org/10.21428/92fbeb44.b540aa59

[16] Yichen Wang, Mingze Xi, Matt Adcock, and Charles Patrick Martin. 2023. Mobility, Space and Sound Activate Expressive Musical Experience in Augmented Reality. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, Miguel Ortiz and Adnan Marquez-Borbon (Eds.). Article 17, 6 pages. https://doi.org/10.5281/zenodo.11189129

[17] Anıl Çamcı and John Granzow. 2019. Hyperreal Instruments: Bridging VR and Digital Fabrication to Facilitate New Forms of Musical Expression. *Leonardo Music Journal* 29 (12 2019), 14–18. https://doi.org/10.1162/lmj_a_01056