Matthew Hadick
HCI5800 - Spring 2025

Sonic Sculptures in Space - A VR Synthesizer Experience

Final Report

Executive Summary

Category	Details
Team Members	Solo (Matthew Hadick)
Project Type	VR (Virtuality Reality)
Project Category	Arts & Entertainment
Goal	Create a VR app that reimagines sound design as an embodied, spatial practice by transforming virtual reality into an intuitive, sculptural environment for musical exploration. The project aims to: ● Leverage VR's immersive potential to make spatial audio a primary medium of creative interaction. ● Lower the technical and cognitive barriers to musical expression for non-musicians. ● Replace abstract interfaces (e.g., DAWs, code-based tools) with direct, sensory engagement: walking around, shaping, and responding to sound in real time. ● Create a playful and accessible space that blends composition, performance, and learning into a single interactive experience.
Vision	To evolve Sonic Sculptures in Space into a flexible ecosystem for spatial sound exploration across VR and AR. Future development will include: ● **Creative Modalities:** Serve as both a meditative environment and a platform for performance art and improvisation. ● **Educational Utility:** Teach users sonic principles (waveforms, frequency, spatialization) through embodied play. ● **AR Integration:** Allow users to create soundscapes in their physical surroundings via augmented reality. ● **Multiplayer Support:** Enable collaborative performances and co-authored sound environments. ● **Expanded Toolset:** Introduce additional sound object types, real-time modulation tools, and gesture-responsive controls.
Technology Used	**Software:** Unity (XR Toolkit, Audio Synthesis), Built-in shaders and interaction assets **Hardware:** Windows 11 PC (development), Meta Quest 2 (headset)
User Interactions	**Glowing Orbs** – Grabbing orbs changes the waveform and associated color; if they collide with one another, the ground, or walls, the pitch changes to a random pitch within an programmed musical mode. **Moveable Walls** – Can be grabbed and used as paddles or shields; dampen the sound and emit their own low frequency hum in the same musical mode as the spheres. **Tiltable Map** - Users can tilt the map using the right analog stick to move spheres using gravity. **Proximity-Based Interaction** – User movement around the map influences spatial audio mix, enhancing immersion.

Goal and Literature Review

Digital audio interfaces and workstations have tended to resemble traditional software and have existed almost entirely in two dimensions. In order to create music, users have to either play virtual instruments, attach instruments via MIDI, or program sound using code or graphical user interfaces. While these traditional modalities have been the foundation of some incredible music and sound design experiences, and there is still yet untapped potential in their use, they still require substantial expertise to use effectively, making the barrier to entry high for non-musicians.

In digital music production, there is often an emphasis on emulating space using digital components (Virtual Studio Technologies, or VSTs, typically) that generate effects that change the character of the sound. Reverb, delay/echo, stereo panning – effects like these create an illusion of space that musicians and sound designers can take advantage of for their compositions.

Virtual reality expands the possibility space for interacting with sound by allowing designers to expand the scale of instruments to the scale of environments emulated by these effects. Instead of creating the sense of a room with sound, virtual reality allows designers to place sound in the room, anchoring it to interactive, spatially-aware objects that users can approach, manipulate, and inhabit.This shift transforms sound design from a symbolic or representational process into an embodied and tactile experience.

The goal of this project was to create a system where the user doesn't simply sculpt sound through parameter changes on a screen – they walk around it, reach into it, reshape it with gestures, and hear it respond in real time based on their position in space. In this way, VR doesn't just simulate spatial audio, it makes spatiality a central medium for musical expression, dramatically lowering the cognitive and technical requirements for musicians and non-musicians alike to engage creatively with sound. It can be used as a way of simultaneously making and enjoying music, creating an opportunity to learn more about properties of sound and its symbiotic relationship with space.

The synergy between virtual reality (VR) and music has been a focal point since the early days of VR development. Jaron Lanier, regarded as the godfather of VR, focused on this intersection throughout his pioneering work with the medium. In a VR performance piece titled *The Sound of One Hand*, Lanier utilized a DataGlove and head-mounted display to create and manipulate virtual instruments in real-time, allowing for a unique, improvisational musical experience. He described the instruments as "somewhat autonomous, and occasionally fight back," highlighting the dynamic interaction between performer and technology .

Further elaborating on the role of music in technology, Lanier authored an article titled "How Music Will Save the Soul of Technology," where he posited that music serves as a medium to humanize technological advancements, ensuring they remain connected to human emotion and creativity.

Aligning with this philosophy, contemporary research increasingly explores the convergence of virtual reality and musical interaction. For example, Hamilton et al. (2021) present a methodological framework for evaluating social presence in collaborative music-making within VR environments, emphasizing embodiment and interpersonal coordination as key dimensions of the experience. Similarly, Shvets and Gillies (2023) propose a model for volumetric music composition, redefining the listener's role as a spatially-embedded co-creator and highlighting VR's capacity to transform musical authorship through immersive design.

As Buckley and Carlson (2019) argue, music-led VR composition must consider not only traditional elements like timbre and structure but also the phenomenology of embodiment—what they term the "techno-somatic dimension," where technology and the body fuse to create seamless, intuitive interaction. Sonic Sculptures in Space fully embraces this paradigm, treating interaction not as a means to manipulate sound but as a form of participation within a living acoustic system. In this environment, the body becomes an instrument of exploration, and the space becomes both score and stage. This stands in contrast to many VR music composition tools, even Lanier's, which often replicate existing paradigms – whether digital audio workstations or physical instruments – without reimagining the relationship between scale, gesture, and musical emergence.

These explorations show the potential of VR as a transformative medium for musical expression, and align with Lanier's vision of technology that enhances, rather than diminishes, the human experience. My goal is to create a system where sound surrounds the user, responds to them, and grows with them, making musical exploration feel less like using a tool and more like entering a living collaboration. Ultimately, this project aims to lower the barrier to musical expression by transforming abstract sound design into a direct, sensory experience.

Vision

In its fullest form, Sonic Sculptures in Space aspires to become a dynamic, open-ended playground for environmental sound creation

Core Applications

- Meditative and Performative Space: Functions as both a contemplative environment and a canvas for experimental performance art, encouraging users to engage in ambient creation through presence and gesture.
- Music and Sonic Education Tool: Offers embodied learning experiences that introduce users to the properties of sound—such as waveform, frequency, envelope, and spatialization—through intuitive, hands-on interaction rather than abstract explanation.

Future Iterations

- AR Integration:Expand into augmented reality to allow users to paint and place sound in their actual environments, merging physical space with generative soundscapes.
- Multiplayer and Social Co-Creation: Introduce synchronous multi-user functionality, enabling collaborative sound composition, shared performances, and social immersion in persistent spatial audio environments.
- Expanded Interaction Vocabulary: Add new sound object types (e.g., rhythmic loops, generative textures, reactive instruments), manipulable parameters (e.g., filter, LFO, envelope shaping), and interaction modalities (e.g., gesture-based modulation, body-tracking).

Software/Hardware Used

The project was developed in Unity using the XR Interaction Toolkit and built-in audio synthesis tools. It was tested on a Windows 11 PC and deployed to a Meta Quest 2 headset using Oculus Link. No external assets or packages were used.

Final Product

The project leverages Unity's XR Hands setup with ray and direct interaction support to enable intuitive, embodied manipulation of audio-emitting objects in a fully immersive VR environment. No third-party assets were used; all visual, auditory, and interaction behaviors were implemented from scratch using custom C# scripts and Unity's built-in systems. The scene hierarchy features an XR Origin rig and contains spatialized sound spheres, interactive walls, a map-tilting mechanism, and environmental structures that work together to produce an evolving, reactive soundscape.

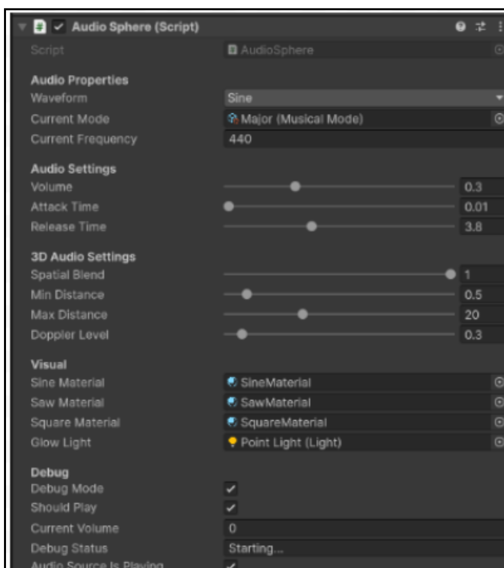

AudioSphere Component within Sphere Object, inspector view

At the core of the project are the AudioSpheres. These are procedurally generated objects that emit real-time synthesized sound based on waveform type (sine, saw, square) and pitch. Sound is created via Unity's low-level OnAudioRead() callback, bypassing pre-recorded samples in favor of dynamic audio synthesis. The spheres respond to collisions with walls, the floor, and other spheres by altering pitch, adding audio effects, or shifting octaves. Waveform changes are triggered when the player grabs and releases a sphere using XR's XRGrabInteractable component. Each AudioSphere includes settings for attack, release, spatial rolloff, and doppler simulation, ensuring a smooth blend of musical responsiveness and perceptual realism. Spatial blend and volume attenuation are tuned using custom rolloff curves, reinforcing the sonic materiality of each object as it moves through space.

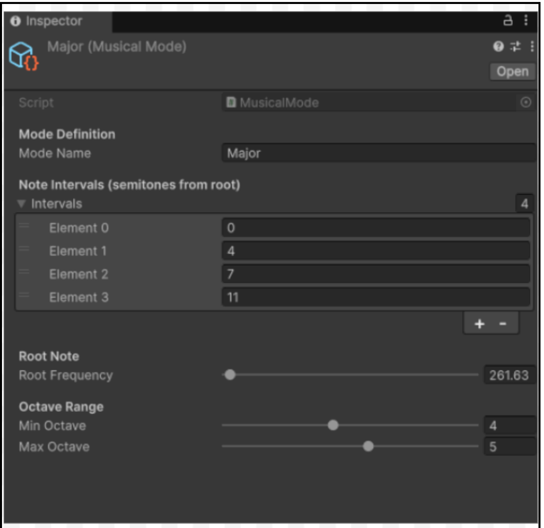

Modal settings, prefab asset, inspector view. Users can create custom modes established by root frequency and intervals, and set the octave range.

Musical pitch is governed by a ScriptableObject prefab called MusicalMode, which defines intervals, root frequency, and allowable octave ranges. This framework enables consistent harmonic output while still allowing for procedural randomness. Visual feedback is tightly linked to sonic properties: each waveform corresponds to a distinct material and emissive lighting effect, and sphere glow intensity is modulated in real-time by a global LFO system. The GlobalLFO script uses an AnimationCurve to drive oscillations across time, which are then broadcast using UnityEvent<float> to all subscribed visual elements, linking sound, light, and motion in an ambient rhythm.

MoveableWall components extend the interactive sonic palette by emitting low-frequency pads that generate ambient bass drones. Like the AudioSpheres, these walls synthesize sound in real-time, using their own OnAudioRead() methods to produce a continuous tone. Walls can be repositioned in space via XR ray or direct interaction, enabling users to structure the sound environment physically. Importantly, walls also act as dynamic dampeners, modulating the volume of nearby AudioSpheres based on distance and a configurable dampening strength. This system introduces a novel spatial mechanic: walls don't simply play sound, they shape it by conditioning the acoustic environment through proximity-based filtering and sound dampening.

Additionally, the player can use the right control stick to "tilt" the map entirely, which allows for creating cascading interactions between spheres and surfaces, effectively turning gravity into a compositional tool.

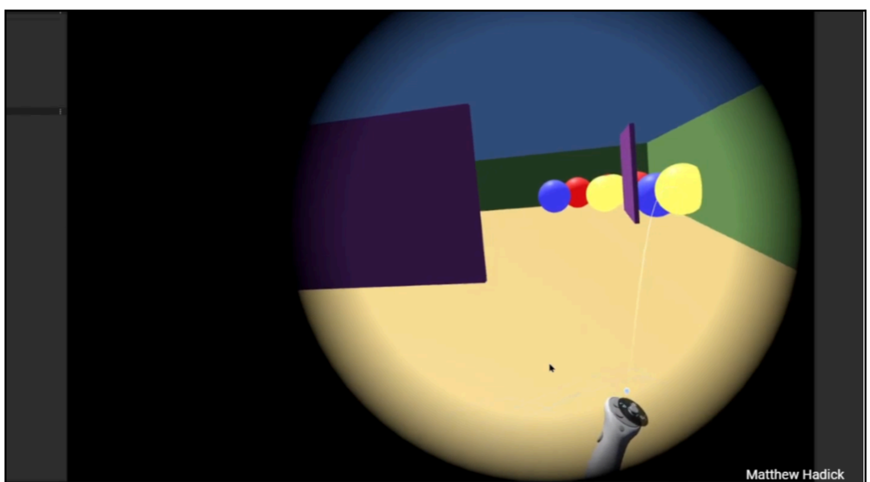

Screenshot from Play Mode. Players can move the right joystick to tilt the environment, allowing gravity to affect the pitch emitted by spheres.

All audio is routed through a centralized AudioManager using Unity's AudioMixer system, which allows grouped control of global effects, reverb settings, and master output levels. The AudioManager provides runtime access to key parameters, ensuring that sonic environments can be scaled, tuned, or modified globally without interfering with per-object behavior. The XR Interaction Toolkit plays a critical role in maintaining a seamless user experience, offering built-in grab detection, pose tracking, and interaction events across both hands.

Relationship to Course Content

This course was inspired and influenced by a wide range of topics covered in class, including:

3D Spatial Audio & Embodied Perception:
Drawing from Week 11's focus on spatialized audio and HRTF-based positioning, this project leverages Unity's spatial audio capabilities to ground sound sources in 3D space. Users hear sound shift naturally as they move, reinforcing spatial awareness and embodiment in line with VR presence theory (Week 5).

Presence, Interaction, and Immersion:
As discussed in Weeks 4 and 5, immersion and presence are enhanced when sensory inputs are harmonized with intuitive interaction. By replacing traditional sliders and MIDI inputs with gestural, spatial controls, the project repositions users as embodied performers within a dynamic sonic environment—echoing principles from Heeter's model of personal and environmental presence.

Synthetic Sound Design & Real-Time Feedback:
In line with Week 11's breakdown of synthetic audio and waveform synthesis, Sonic Sculptures in Space employs procedural sound generation from sine, saw, and square waves. Each interaction becomes a real-time act of composition, emphasizing sound as emergent, embodied feedback rather than symbolic input.

Summary

While this working prototype only scratches the surface of what such a system could offer, I believe it represents a strong first step toward reimagining how people interact with sound. The project saw considerable success, though not without its share of challenges.

On the success side, I focused on creating a smooth and intuitive play experience, and I believe this was achieved, even if the final presentation video doesn't fully reflect it due to capture lag. Grabbing and manipulating objects feels fluid, waveform changes are clearly conveyed through sound and color, and interactions are easy to interpret. I intentionally avoided relying on external audio middleware like Wwise to keep the workflow simple; while future iterations would benefit from deeper integration, I was pleased with what Unity's native audio engine could accomplish. The AudioSphere script allowed for clean, modular design through public class structure, enabling flexible frequency control and waveform switching. From a development standpoint, the MusicalMode ScriptableObject proved especially elegant. It's lightweight, intuitive, and musically rich.

From a UX perspective, I aimed to keep the stakes low and the environment open to exploration. While new VR users might benefit from a more explicit onboarding layer, leaning into existing interaction conventions allowed for a self-guided, discoverable experience. That said, both UX and UI will need iteration to enhance accessibility and support more advanced control.

There were also challenges. The project originally began as an AR prototype centered on visual painting, but technical limitations forced a pivot to VR. While I preserved the spirit of the idea centered on 3D audio, the visual expressive tools had to be tabled. I also attempted to include in-world UI sliders on the moveable walls for real-time audio manipulation. Although the logic worked, they caused repeated crashes in play mode, despite extensive and time-intensive debugging, so I ultimately disabled them before the final demo. While these will be essential in future versions to deepen control and expression, if I were to start the project over, given the time frame, I would have spent much less time troubleshooting these diegetic and Meta UI features and focused more on the spatial UX, interactions, and general functionality.

All things considered, this prototype lays a meaningful foundation for a system that transforms sound design into a spatial, embodied experience – one that invites users not just to shape sound, but to inhabit it. As a researcher and designer without an extensive programming background, I learned so much about building procedural systems from scratch, structuring modular scripts, and translating conceptual interaction models into functional, real-time experiences.

REFERENCES

Buckley, K., & Carlson, K. (2019). Towards a framework for composition design for music-led virtual reality experiences. Proceedings of the Audio Mostly 2019 on Interaction with Sound, 1–8. https://doi.org/10.1145/3356590.3356617

Hamilton, A. F., Murphy, D. T., & Wanderley, M. M. (2021). A framework for the evaluation of social presence in musical interactions in virtual reality. Frontiers in Psychology, 12, 663725. https://doi.org/10.3389/fpsyg.2021.663725

Lanier, J. (n.d.). How music will save the soul of technology. Retrieved May 9, 2025, from https://www.jaronlanier.com/knittalk.html

Lanier, J. (n.d.). VR page featuring The Sound of One Hand. Retrieved May 9, 2025, from https://www.jaronlanier.com/vr.html

Shvets, A., & Gillies, M. (2023). Volumetric composition: An embodied approach to composing music in virtual reality. EVA London 2023 Conference Proceedings, 166–173. https://www.scienceopen.com/document?vid=97bb9981-3ae8-4814-bacb-7ac737dfeaff