SVOrk: Stanford Virtual Reality Orchestra

Kunwoo Kim
kunwoo@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Andrew Zhu Aday
azaday@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Eito Murakami
eitom@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Marise van Zyl
marise@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Yikai Li
yikaili@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Max Jardetzky
mjardetz@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Ge Wang
ge@ccrma.stanford.edu
CCRMA, Stanford University
Stanford, California, USA

Figure 1: SVOrk Premiere Concert, June 2024.

Abstract

This paper chronicles the creation of the Stanford Virtual Reality Orchestra (SVOrk), a new computer music ensemble where both performers and audience engage in a shared, fully immersive virtual reality (VR) chamber-esque concert experience. Motivated to explore group-based live performance within VR, SVOrk has designed and crafted virtual musical interfaces, fantastical 3D-modeled environments, and a network infrastructure to support real-time shared participation. Inherent within this initiative is a reimagining of conventional concert experiences, introducing virtual lobbies, customizable avatars, and immersive audience interactions. These experimental features explore new forms of social presence, audience identities, and expressive communication to help address the overarching question, "What does it mean to participate in a VR musical performance?" SVOrk's premiere concert took place in June 2024 with five performers and approximately 60 audience members (across five sessions), featuring a program of six musical works. This paper describes the motivations behind SVOrk, its research and development process—including designs for networking, avatar, and audience interaction—and takeaways from the premiere concert. We also present audience feedback and reflect on our experiences in curating group VR performances.

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

Keywords

Virtual Reality, VR Concert, VR Orchestra, VR Instrument Design.

1 Introduction

Stanford Virtual Reality Orchestra (SVOrk) is a computer music ensemble where both performers and audience engage in a shared, fully immersive virtual reality (VR) chamber-esque music performance. Founded in 2024 at Stanford University's Center for Computer Research in Music and Acoustics (CCRMA), SVOrk aims to explore group musical performance and audience engagement in VR environments. The ensemble makes use of head-mounted displays (HMDs), laptops, and custom software; it integrates real-time computer-mediated music, bespoke VR instruments, and 3D-modeled visual environments to offer possibilities for fantastical, expressive, and immersive experiences.

SVOrk is also a *concert-going* experience in VR, featuring two virtual lobbies (Private Waiting Room and Public Waiting Room) and a series of live, group computer music performances. In the Private Waiting Room, the audience members are introduced to their virtual environment through brief tutorials on locomotion, system settings, and virtual reactions, and can customize their avatar's appearance. The Public Waiting Room serves as a virtual meeting space, where participants can socialize, read the virtual program notes, and musically interact with the environment until the rest of the attendees arrive. Once the concert program begins, the audience is immersed in a series of musical works set within their respective virtual environments.

SVOrk uses Chunity [7] (ChucK [31] in the Unity game engine [1]). ChucK is a strongly-timed computer music programming language that enables temporally deterministic control over dynamic sound synthesis. The audio parameters in ChucK can be precisely mapped with graphical and interactive parameters, enabling systems that closely integrate programmable audio, graphics, and interaction. The framework built using Chunity, called the SVOrk Toolkit, is to be released as open-source software and a library within Unity.

Virtual reality, specifically that which uses HMDs, affords a heightened sense of presence and immersion, allowing users to engage with simulated environments in ways that mirror their perception of and interactions with the physical world, even if they defy physical laws [9, 20]. With 3D-modeled graphics and spatialized audio rendered live through HMDs and headphones, users are transported into completely computer-generated, interactive worlds that can also be connected using computer networking. These distinctive features of VR invite new possibilities for group-based live performances.

To explore these potentials, this project focuses on three key design areas: networking, avatar, and audience interactions. First, SVOrk requires a local-area network system to connect performers and audience members in real-time by communicating parameters for audio, graphics, and interactions. Second, SVOrk needs avatars to represent the participants' presence and movement within the shared virtual environment. This leads us to explore the forms and dynamics of how participants exist in and navigate through our VR environments. Third, SVOrk seeks to create mechanisms for audience participation throughout the performance. This area investigates audience communication, including new forms of reactions.

After developing the basic infrastructure and presenting two demo concerts, SVOrk premiered in June 2024. Five performers showcased six original musical works to approximately 60 audience members spread across five sessions. These works, composed and crafted by the authors, incorporated fantastical worlds and musical interfaces: a journey through the clouds; a dandelion within machines; uncanny facial expressions; subconscious memories; vulnerable inner spaces; and commentary on a world grappling with digital technologies. In addition to the musical pieces, we developed a method to record and replay the concert in VR, as well as ways to curate the concert for a physical venue. Lastly, we gathered feedback from the audience members who attended the premiere concert; we highlight these in Section 4.4.

As a roadmap for the rest of the paper, Section 2 describes our motivations and related works. Section 3 details the research and development in networking, avatar, and audience interaction. Section 4 chronicles the premiere concert. Section 5 provides future work ideas and conclusion.

2 Motivation and Related Works

Research on VR music concerts has grown rapidly in recent years in the contexts of performance, audience, and venues [23]. While VR concerts have expanded across various domains, here we focus on those that are fully immersive (requiring an HMD to experience) and feature 3D-modeled environments (excluding VR experiences that use 360° video to capture live performance in physical reality). In particular, we have identified three defining characteristics of SVOrk and will address their respective literature: 1) SVOrk uses computer-mediated music that is synthesized and generated in real-time, often in response to human interaction; 2) SVOrk provides a shared VR environment for both performers and audience; 3) SVOrk concerts take place in a physical venue, fostering a shared sonic and spatial intimacy, akin to a chamber performance in Western art music.

First, SVOrk adopts an audio-first approach, where sound and music are fundamental to the functionality and aesthetics of a user-experience [28]. This is often achieved by controlling audio parameters in real-time through dynamic sound synthesis, live user interactions, and spatial information [8, 9, 16, 34] in what Serafin et al. (2016) have called Virtual Reality musical instruments (VRMIs) [26]. This contrasts with many existing VR music platforms that utilize VR to experience live musical performance done in physical reality, potentially with virtual representations of the performer, as well as augmented visuals and interactions [11, 12, 18, 19]. While these works explore the VR medium from unique perspectives, SVOrk focuses specifically on how the medium interacts with computer-mediated music.

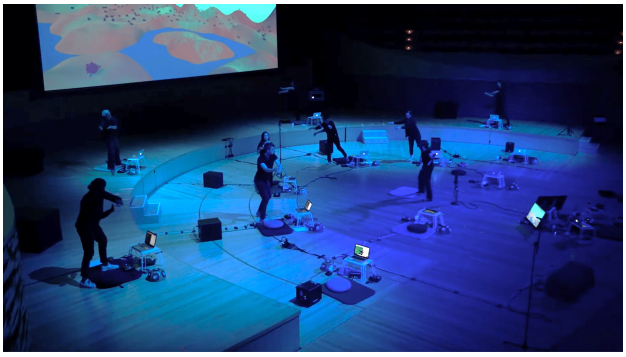

Figure 2: *Resilience* (2019) by Jack Atherton.

Second, SVOrk creates a shared VR environment for both performers and audience in real-time. While some existing works incorporate performances in VR with dynamic sound synthesis [8, 16, 17, 25], these performances are typically projected on a

2D screen for the audience to experience (Figure 2). SVOrk builds upon the ethos of the VR computer music ensemble performance of these works, but it aims to further embrace the medium by placing the audience within VR alongside the performers. The VR environments of SVOrk are designed to be fantastical, expressive, whimsical, and social, and champion an artful design approach to tool-building [29].

Third, like its spiritual predecessors, the Stanford Laptop Orchestra (SLOrk) [30] and the Princeton Laptop Orchestra (PLOrk) [27], SVOrk is an "electronic chamber ensemble" in VR. The concert occurs in a physical venue, where all participants—performers and audience members—are keenly aware that they share a space, sometimes shoulder-to-shoulder. This differentiates SVOrk from larger scale VR concert platforms such as TheWaveVR [6] and Sensorium Galaxy [5]. While VR offers advantages like remote access and large-scale participation, a survey by Onderdijk et al. (2023) found that 50% of virtual concertgoers felt the absence of social interaction was a drawback [22]. SVOrk was designed to take place in a physical venue that fosters social intimacy (dressing for the occasion, socializing in lobbies, being ushered to seats, engaging in in-person discussions, etc.), ensures accessibility to necessary equipment, and allows structured ways to respond to technical difficulties. Therefore, SVOrk prioritizes physical co-presence over remote participation.

To the best of our knowledge, SVOrk is first-of-its-kind in that it combines all of these characteristics. A defining aspect of SVOrk is its ability to create vastly different reality contexts with dynamic group-based audio interactions that can be intimately shared among all participants (Figure 3).

Figure 3: SVOrk concert in shared physical and VR space.

3 Research and Development

This section outlines the three key design tasks in the research and development of SVOrk. Our team of six was divided into three groups, each dedicated to a specific task.

3.1 Networking

The first technical requirement for SVOrk is the synchronization of graphics, audio, and concert states across all devices. For example, audience members need to see each other's avatars, while performers need to update the audio being played through everyone's headsets, not just their own. Updating concert-related information must happen synchronously in real-time to ensure that all participants experience the performance simultaneously from the perspective of their personal avatars. Moreover, when designing pieces for SVOrk, composers should be able to focus on crafting audiovisual experiences, rather than troubleshooting network issues. To meet these demands, we developed "SVOrk Netcode," a custom networking solution tailored to the design requirements of the VR Orchestra.

With SVOrk Netcode, performer and audience computers are connected to a server written in ChucK, using a network switch (2 gigabit switch was used in the premiere concert). All messages are relayed via the server using Open Sound Control (OSC) [32], such as avatar transforms, audio parameters, event triggers, and conductor cues. No raw audio data are transmitted; instead, because audio is being synthesized on each machine in real-time, we network only high-level synthesis parameters and timing events. This approach significantly reduces the network load and mitigates the need for complex data compression/decompression schemes. In addition, being on local area network kept client latency at a minimum; in practice there were no noticeable latency issues during our concerts.

After much consideration, we decided not to use existing networking libraries in Unity, including Netcode for GameObjects [4], Fish-Net [2], and Mirror Networking [3]. These general-purpose libraries are designed for multiplayer video games and address networking challenges such as cheaters, unstable network conditions, and networked physics. However, these libraries are complex both in their implementation and usage. Most of these challenges do not apply to SVOrk. Furthermore, SVOrk has novel requirements for audio synchronization not supported by existing networking solutions (e.g. networking data in Chunity/ChucK). Therefore, we designed our own custom networking system, built around ChucK, OSC, and Chunity.

The complete SVOrk Network implementation—including the server, client, and record/replay system—comprises approximately 2,500 lines of code (LOC). In comparison, Unity's Fish-Net library clocks in at 147,654 LOC. By focusing on our particular networking needs, we were able to develop a streamlined solution within a single quarter—a small fraction of the complexity of existing alternatives. While we had our fair share of performance-time challenges, networking was not one of them.

The design of SVOrk's lightweight networking system arose from our observation that modern programming culture tends to admire complexity and discourages doing things from scratch. The phrase, *"Don't reinvent the wheel,"* which is used to justify this stance, is often misunderstood. In reality, the wheel actually needs to be reinvented for every new context—bicycles, roller blades, all-terrain, tractors, etc. Moreover, ownership matters; having your own bike, as opposed to renting a bike (this is analogous to software licensing), implies you have the deep knowledge to repair and adapt your bike, especially when riding into new terrains. Lastly, the "wheel" is symbolic of a perfect solution. In practice, however, perfect solutions do not exist in complex systems—like networking. Rather than calcifying the status quo,

we took the opportunity (and effort) to develop original tools. SVOrk Netcode emerges from and reaffirms this ethos.

3.2 Avatar

The second design task for SVOrk is to create an avatar that digitally represents the participants' embodied presence, movements, and interactions with the virtual environment. Avatars play an important role in VR, particularly in shared spaces. They allow participants to personalize their digital presence, enhance self-expression, and promote empathy [10, 24, 33].

To foster these qualities, our design requires avatars to represent the participant's head orientation, locale, and basic gestures. SVOrk uses a blob-shaped humanoid form in which the participant's head, body, and hand movements are kinetically mapped to the tracking data from the HMD and VR hand controllers. This transform data is then communicated across the network, enabling shared social presence and interactions among all participants.

Studies show that avatars can inadvertently reinforce various biases related to gender, race, and socioeconomic status [10, 15]. Therefore, SVOrk keeps the height and shape consistent among all avatars. However, to foster a more personalized connection between the audience and their digital representations, we include customizable options such as choosing two accessories (sunglasses, hats, backpacks, belts, etc.) and a material for their avatars (Figure 4).

Lastly, SVOrk implements a data structure to store various avatar properties, including avatar customizations (selected accessories and material), system settings (volume, movement speed, and turn angles), and locomotion options (enable/disable movement, turn, gravity, and flying maneuvers). Any changes made to these parameters are saved to a local file, to ensure consistency throughout the concert, even in cases of restarting the program.

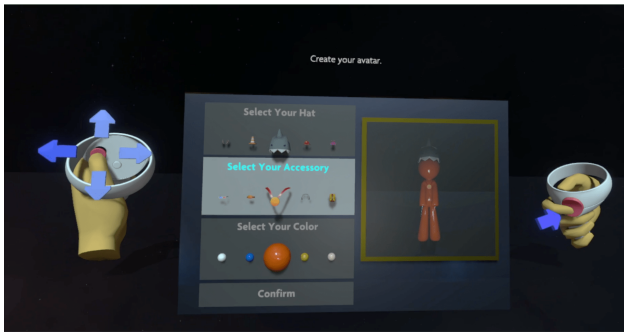

Figure 4: SVOrk's avatar customization process.

3.3 Audience Interaction Design

The third design task is to craft VR audience interactions and the contexts around them. The audience needs: 1) methods to react to performances (i.e. VR-equivalent applause) and adjust system settings; 2) instructions on how to use the controls and navigate through the user interfaces; and 3) social and environmental context for interactions.

In conventional concerts, applause serves as a way to express appreciation, encouragement, and approval of the performance. In VR, however, this form of traditional clapping does not easily translate as the participants are holding their controllers, while visually and sonically isolated by their HMD and headphones. Therefore, SVOrk includes three different visual representations of applause: star particles emanating from the audience's avatar hands, shooting stars projecting across space, and flower petals gracefully falling around their bodies. Each reaction has its own dynamics. Star particles add expressivity to hand gestures, shooting stars target distant locations, and petals emphasize the audience's presence. These appear as three bubbles left of the audience's avatar, which they hover their hand over to make a selection, then press a button on their controller to trigger the chosen reaction.

Since every audience member uses an individual HMD and software to experience the concert, enabling adjustments to system settings may provide a more personalized experience. To facilitate this, SVOrk includes an "audience toolbox" that lets them modify volume, movement speeds, and turn angles (Figure 5). This toolbox also allows them to view a tutorial and program notes at any time during the concert. Instead of a conventional panel interface, SVOrk uses an animated toolbox that sits beside an audience member, opening and closing upon use—adding a layer of whimsy to the fantastical experience.

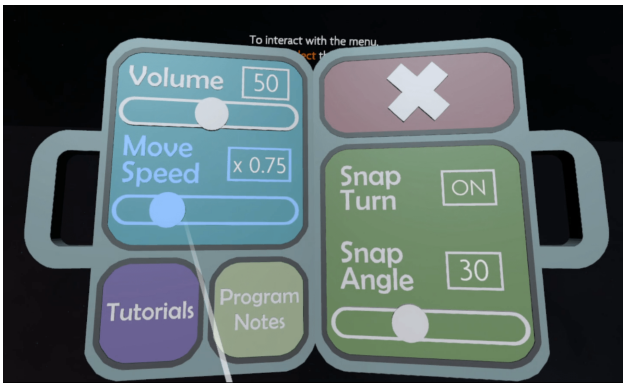

Figure 5: Audience toolbox used to adjust system settings.

The audience learns about these interactions through a virtual tutorial space called the Private Waiting Room, the first VR environment they encounter in the concert. This three-to-five minute session offers a series of a guided onboarding experience, including a welcoming message, safety precautions, avatar customization, locomotion, system settings, and audience reactions. The user interface was designed to accommodate people with no experience in VR by limiting controls to a single button on each hand. Instructions are conveyed through a combination of narration, text, and graphical animations. (Figure 4).

After the Private Waiting Room, the audience gathers in a space called the Public Waiting Room (Figure 6). The 3D-model of this area resembles the physical venue to serve as a transitional space between the physical world and the virtual world. In this room, performers and audience members can freely move around and meet each other for the first time in their avatars. Once everyone has arrived from the Private Waiting Room, performers trigger a recorded verbal announcement, signaling the start of the concert.

During development, we made two key design choices: 1) we decided to minimize any form of verbal or text communication among attendees, such as voice chat, text messages, or IDs appearing above their avatars, to preserve anonymity and encourage free expression. This approach prevented any hints that could reveal non-virtual information about the attendees.

Figure 6: Public Waiting Room filled with audience members.

2) We aimed to maintain our ethos of a computer-mediated music environment even within the Private Waiting Room and the Public Waiting Room. The background music in both spaces is generated through programmable audio. In addition, in the Public Waiting Room, the audience can interact with the curtains to change audio parameters such as volume, tempo, instrumental densities, and groove of the music.

4 Premiere Concert

SVOrk premiered on June 1st and 2nd, 2024, with five one-hour concerts. Each session included five performers, three staff members, and approximately 12 audience members. The authors each composed and crafted six musical works, over the course of eight weeks. The concert was held at a physical venue (CCRMA Stage). Before the premiere, we held two demo concerts—one in March 2024, and one in May 2024—to test out the equipment and network stability.

In this section, we describe the musical works (artist statement and musical interface) performed during the premiere concert, concert procedure (hardware, software, safety, performer panel, program notes, running and recording the concert), and post-concert audience feedback.

4.1 Musical Works - Artist Statements

4.1.1 CCRMA Dreams. Designed by Ge Wang, this piece takes the audience to a dream world of CCRMA through a whimsical flight through the clouds, ushering through a field of chickens, flying whales, a virtual building modeled after CCRMA, and images of CCRMA's people (Figure 7). "If CCRMA were to dream, it would dream of its people." was the main inspiration of the piece.

Figure 7: *CCRMA Dreams* (2024) by Ge Wang.

4.1.2 Dandelion. Designed by Kunwoo Kim, this piece places the audience within a vast system of metallic cogs, a bell tower on the ceiling, falling streams of sand, and a dandelion that blooms in the center of the space (Figure 3 and 8). *Dandelion* expresses the beauty of ephemerality of human existence, symbolized by the life-cycle of a dandelion.

Figure 8: *Dandelion* (2024) by Kunwoo Kim.

4.1.3 The Fragmented Self. Designed by Yikai Li, this piece has the audience floating in front of a gigantic face, whose facial muscles are coordinated by the performers (Figure 9). *The Fragmented Self* expresses how modern people are inundated with information—whether from the internet or the world around them—and how this influx shapes our emotions.

Figure 9: *The Fragmented Self* (2024) by Yikai Li.

4.1.4 RemembeRanch. Designed by Max Jardetzky, this piece is set within the interior of a pulsating, colorful, metallic brain, where a cowboy gang wrangles memory blocks back together (Figure 10). *RemembeRanch* expresses the fragility and fallibility of memory, particularly on the subconscious processes that reinforce and solidify memories.

Figure 10: *RemembeRanch* (2024) by Max Jardetzky.

4.1.5 aSSeRtIon. Designed by Marise van Zyl, this piece is set in a room with a basic, nondescript avatar lying on a bed scrolling on a phone (Figure 11). The character grapples with whether or not to take the anti-depressant (SSRI) to cope with symptoms like anxiety, depression, overthinking, and obsession. *aSSeRtIon* provides a room to reflect about mental health, wellness, and medication, moments that are typically hidden from others.

Figure 11: *aSSeRtIon* (2024) by Marise van Zyl.

4.1.6 Prisoner of the Mind. Designed by Eito Murakami, the piece begins inside a semi-translucent train, which escorts the audience into a dissonant digital world filled with cityscapes, flying surveillance robots, marching humanoids and a giant clock ticking in the sky (Figure 12). *Prisoner of the Mind* expresses the discord and isolation of being immersed in the world of digital technology.

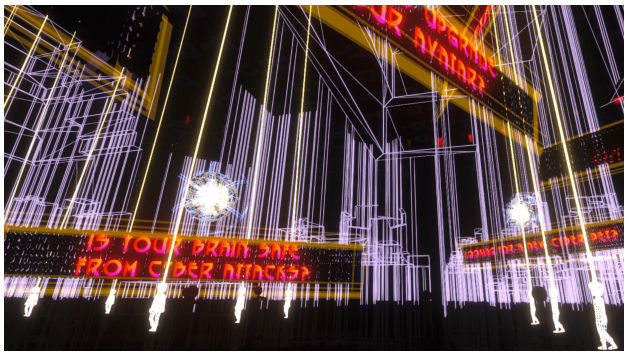

Figure 12: *Prisoner of the Mind* (2024) by Eito Murakami.

4.2 Musical Works - Interface

4.2.1 Audio Spatialization. Spatialized audio is essential to a VR concert experience, as it enhances sense of immersion, while mitigating potential disorientation [9]. Since each audience has individual agency of locomotion within a virtual world, all pieces employ customized binaural spatialization, using interaural level differences. Audio sources can be attached to a virtual object with a custom volume rolloff (distance based amplitude difference) and spread (angle of the source). Thus, each virtual object's distance and angle from the audience member's position and head orientation determine the respective volume levels in the left and right audio channels.

4.2.2 Sound Synthesis. Since SVOrk uses ChucK to generate audio in real-time, any audio synthesis, algorithms, or parameter controls programmable in ChucK are available as part of musical expressions.

For example, *CCRMA Dreams* uses banded waveguides to create drones with pseudo-randomized pitch and physical model properties. *aSSeRtIon* uses various oscillators to create repetitive synthetic sounds, bubbling grumbles, and descending Sheppard tones. *Prisoner of the Mind* uses pre-composed music created using RayTone [21], with dynamic tempo and pitch modulation towards the end of the narrative. *Dandelion* and *Fragmented Self* both use granular synthesis. These sounds are shaped further with various audio filters, reverberation, and effects like echo and delay.

4.2.3 Virtual Reality Musical Instruments. Performers in SVOrk interact with virtual objects to play music as an ensemble. Most interactions are mapped to continuous body movements as they enhance both physical and psychological connections to the virtual environment [20]. Chunity allows synchronized mapping between audio, graphics, and interactions.

In *Dandelion*, two performers use their hand movements to interact with a virtual wheel in front of them. The wheels control the temporality and volume of 20 different machine sound samples in algorithmic polyrhythm. Other performers use arm gestures to alter audio parameters such as vibrato, reverberation, and volume, mapped to the velocities of their hand controllers.

In *The Fragmented Self*, each performer, like a puppeteer, uses the positions of their hands to control different portions of the dynamic mesh of a human face. While reflecting various facial expressions, the x-y-z hand positions modify various granular synthesis parameters, such as grain volume, length, and position, to manipulate sound samples representing laughter, screams, or cries.

In *RemembeRanch*, each performer swings their lasso, with its physical state (particle system positions and velocities) mapped to arpeggiated synth notes that intensify with the rope's velocity. Every lasso collision with memory blocks builds a real-time musical score, including a breathy tonal sounds, heartbeats, samples from each memory, and a sequence of synth notes that correspond to the memory fragments being reassembled.

4.3 Concert Procedure

4.3.1 Hardware and Software. For hardware, each performer and audience station included a laptop, VR headset (a mix of Meta Quest 2, Quest 3, and Quest Pros), a Meta Link Cable, Sony MDR-7506 headphones, and a LAN cable (Figure 13). The stations were spaced about two meters apart. Based on insights learned from previous demo concerts, we made minor adjustments to the hardware to mitigate technical issues, such as covering the internal camera of the headset to prevent it from entering sleep mode when left unattended and securing the Link Cable with tape to ensure that it does not accidentally disconnect during the concert. In addition to the stations, we had a server laptop and a network switch to manage the local-area network throughout the concert.

For software, we developed separate performer and audience programs built in Unity to run on the laptops of each station. The programs then render audiovisual information on the HMDs, using Oculus Rift connected via Meta Link Cable. Performers and audience members begin with different avatar properties

based on the program, enabling the server to differentiate them accordingly.

Figure 13: Example setup of audience stations.

4.3.2 Safety. Users in VR can be susceptible to motion sickness from the inconsistencies between the visual and vestibular information, especially when the experience implies motion [13]. To mitigate this, we designed the concert's duration to be under 45 minutes in VR, and provided safety precautions before and during the event, advising audience members to remove their headsets immediately if they experienced discomfort. Additionally, we implemented snap turns—discrete angular shifts rather than continuous rotation—as the default locomotion mode to reduce sensory dissonance [14].

We asked the audience to raise their hands in case they required assistance with technical difficulties, discomfort, or interactions. Three staff members monitored the audience and equipment in physical reality throughout the concert. We also requested the audience to remain seated for the entire concert to prevent any accidental physical collisions.

4.3.3 Performer Panel. SVOrk's performer software program included a hidden performer panel, which displayed a timer for each piece, performer-IDs, and conducting messages—features that were invisible to the audience but required for conducting and organizing the performance. First, the timer ensured passages to be introduced at the appropriate moments. For example, in *Dandelion*, performers began playing the bells and chants at the 45-second mark. Second, performer-IDs distinguished individual performer roles within the ensemble. For example, in *aSSeRtIon*, performer 1, 2, 3, and 4 played the audiovisual elements of anxiety, obsession, depression, and compulsion, respectively. Third, conducting messages allowed performers to communicate in a text format. For example, in *Prisoner of the Mind*, the composer/conductor used this feature to indicate when each performer should trigger the humanoids.

4.3.4 Program Notes. We provided program notes in both physical and virtual formats. They included an introduction to SVOrk, composers' notes on their works, and bios of the team. The audience received a printed copy upon being seated, while the virtual version remained accessible through the audience toolbox at any time during the concert.

4.3.5 Running the Concert. With multiple equipment operating in tandem—laptop, HMDs, and networking components—the concert was susceptible to various technical difficulties that could bottleneck the whole procedure. From our previous demo concerts, we encountered critical problems such as network congestion, headset position/orientation errors, VR tracking problems, and even electric overloads. To mitigate these risks, we crafted a detailed, minute-by-minute manual for the five performers and three staff members running the concert. This manual included step-by-step instructions for station setup, seat mappings, audience ushering, server initialization, and troubleshooting protocols. Additionally, we had one backup station for the performers, and three for the audience members, as switching stations proved more efficient than attempting on-the-spot repairs. Fortunately, with lessons learned from the previous demo concerts that resulted in a comprehensive procedural guide, the premiere concert proceeded without technical difficulties.

Beyond technical considerations, we aimed to bring the intimate, social contexts of physical venues and ensure a smooth audience experience. After setting up the stations and opening the venue doors, each performer and staff member ushered audience members individually to their seats. Seniors, glasses wearers, and those prone to motion sickness were assigned to stations with Meta Quest 3s and Pros, which provide clearer visuals and more spacious displays. Once seated, audience members were assisted individually in fitting their HMDs, adjusting orientation, and setting the correct volume levels. After calibration, they were asked to remove their headsets and wait until everyone was ready, at which point the program notes were distributed. When all audience members were prepared, we introduced the ensemble, delivered safety precautions, requested everyone to put on their headsets, and began the concert.

This curation process was crucial in running a SVOrk concert, both technically and experientially. VR headsets can be inconsistent in tracking and connectivity. The previous demo concerts showed how technical difficulties can unnecessarily prolong the event, often disrupting the overall experience. Mitigating potential issues in advance, and implementing multiple backup plans proved essential. Post-concert audience surveys revealed that this curated procedure enhanced the audience's overall experience.

4.3.6 Recording the Concert. We developed a system to record and replay the group performance in VR. The system captures all network data and their timestamps as they were routed through the ChucK server. During playback, the server simply replays this data, allowing multiple audience members to experience the performance simultaneously. They retained the agency to move, turn, and react. Since this concert was designed specifically for VR, we thought that any form of documentation and recording should justify the medium as well.

For demonstration videos of the premiere concert:

- Audience: https://youtu.be/n0YC8TXv4ME
- Performer: https://youtu.be/z8O8paXci2A

4.4 Audience Feedback

We sent out an optional, anonymous survey to the audience members after the premiere concert. The survey comprised of a single open-ended prompt asking for their feedback, thoughts, and feelings. 21 individuals provided responses. Below is a selection of results in different categories.

Some attendees found the **audio-first approach** compelling,

- "I appreciate that SVOrk is bringing audio-first experiences to the world of social VR."

- "I feel like the gestures were really well thought out, especially considering the parameters of the setting in virtual reality."

while virtual acoustics could be improved.

- "A distinct virtual acoustic would be great for each piece. Mostly felt very dry."

Attendees found **audience interaction** to be fun.

- "I thought the reactions and avatar creation was a fun detail and great way to get the audience involved."
- "I really liked the format and how the audience was represented by cartoonish avatars and they could see each other during performances. The interactions and movements in the pieces were well thought of and well designed."

Some attendees found **social engagement** in VR satisfying,

- "...it was really interesting and intimate to be a part of someone's thoughts and emotions and feelings without really being able to assign a name to an orange body."
- "In particular some things I found especially effective were seeing the other audience members engage with the dandelions - you could tell their delight even as avatars."
- "Knowing I was part of something with live performers and a live audience made it especially special."
- "There is just something so special about being immersed in the same virtual environment with others. It felt like I was peeking into every SVOrk member's inner imagination—something that I wouldn't be able to experience otherwise."

while some expressed mixed feelings.

- "...an interactive audiovisual experience that is simultaneously creating some kind of celebration of community and individuality...while also feeling lonely and isolated at times due to the nature of being unable to communicate and the nature of some of the content."

Attendees found **fantastical VR experiences** engaging.

- "I enjoyed that the show happened all around me, unlike in a traditional venue."
- "Instead of holding instruments, the performers made sounds with giant bells, buckets of water, laser whips, blinking eyes and mouth, street signs, clock, dark clouds, chickens, anything... This VR orchestra showed me a whole new set of possibilities."
- "This was genuinely one of the most inspiring, beautiful and terrifying experiences I've had recently...it honestly gave me hope of how humans will use technology to teach one another empathy and express the most human notions about us, more effectively. Seeing the frailty of life in the dandelion, the intimacy of struggling with mental health in the bedroom, and the hellscape futuristic city."
- "I loved the introspective and personal nature of many of the pieces, which made the whole experience feel more relatable in a VR environment."
- "Throughout the concert, I felt a mixture of emotions—awe, fear, and empathy. I was in awe of the dandelion growing under the sunshine, fearful of the dystopian future dominated by AI, and empathetic towards the mental struggles someone with anxiety goes through."
- "This was an incredible experience for at least two reasons: 1/ being transported into a dream of an artist; 2/ living art in a 3D dimension where music, visual and experience all converge."
- "I am overwhelmed by the creativity in the performance and really did have a visceral reaction to the emotions meant to be conveyed with the performance."

Some attendees experienced **motion sickness**,

- "My only problem was with the VR headset and nausea, but I knew that may be an issue from past experiences. Closing my eyes for a few minutes if the motion sickness got worse helped me be able to keep the headset on."
- "Motion sickness was only a potential problem in the VR lounge, otherwise not an issue."

while most apparently experienced none.

- "I thought I would feel sick, but I did not (unless I was moving too much or moving sideways). The length was perfect."

Some attendees found the framing of a **virtual** concert within a shared **physical** venue meaningful,

- "A lot of attention was given the audience experience i.e. in email communication before the event, being ushered in, fitting the headset etc."
- "I was in a shared experience due to both the visual aspects and the physical ones; from having a relatively good number of audience, being seated close to each other, being greeted and taken to our seats, to bumping into each other in the virtual world and waving at one another, I felt connected; of course in a new way."
- "I really appreciate the thoughtfully curated experience."

while some preferred remote accessibility.

- "I'm not sure about the benefit of being in the same physical space though. Like I almost would of preferred being at home or in a separate physical space."
- "The most pressing need long term, in my opinion, is a way to attend using one's own VR headset from anywhere on the Internet."

5 Future Works and Conclusion

In its current form, SVOrk invites expansion both in scope and depth. First, the SVOrk Toolkit (built in Chunity) is to be released as an open-source software on the Unity Asset Store. This provides a plug-and-play framework that integrates SVOrk's programmable audio, avatars, and audience interaction designs with local-area networking with the expressive aims for creating new shared VR concerts. We envision the SVOrk Toolkit to expand its applications beyond VR concerts to VR classrooms, social therapy, storytelling, and more. For example, in a class focused on audiovisual interaction design in VR, SVOrk could be used as a virtual gathering and laboratory space for lectures, experimentation, and critique.

Second, we aim to explore the concept of VR concerts more deeply, addressing the following questions: What does it mean to perform together in VR? What does it mean to be an audience in VR? What does it mean to compose VR music? What is "good music" in VR? In particular, we seek to push the boundaries of the medium by exploring virtuosic VR performance, explorations in virtual acoustics, and compositions driven by audience participation. For example, *Prisoner of the Mind* allowed audience members to interact with the humanoids by touching them to trigger alarm and glitch-like sounds. Concert attendees reported this interactive element was particularly immersive. Thus delving deeper into audience participation within SVOrk concerts via

pieces that allow for synchronized musical participation presents an interesting direction for future development.

In conclusion, we have presented the design of SVOrk, Stanford Virtual Reality Orchestra, by chronicling its creation from ethos, research and development, to the premiere concert. In so doing, we are also putting forth a new paradigm for group-based computer music composition, instrument design, and live performance. In evolving SVOrk, we hope to further explore the creative and expressive potential of VR technology.

6 Acknowledgment

This work was done with support from Stanford University's Accelerator for Learning, as part of their Virtual Field Trip Design seed grant. Special thanks to all audience and staff members, who participated in the demo and premiere concerts.

7 Ethical Standards

This paper complies with the NIME ethical standards. The post-concert audience survey was anonymous and optional. Audience members consented that their responses on the survey could be used anonymously in a research paper.

References

[1] 2020. Unity. https://unity.com/. (Software) Unity Technologies.
[2] 2024. Fish-Net. https://fish-networking.gitbook.io/docs.
[3] 2024. Mirror Networking. https://mirror-networking.com/.
[4] 2024. Netcode for GameObjects. https://unity.com/products/netcode.
[5] 2025. Sensorium Galaxy. https://sensoriumgalaxy.com/.
[6] 2025. TheWaveVR. https://wavexr.com/.
[7] Jack Atherton and Ge Wang. 2018. Chunity: Integrated Audiovisual Programming in Unity. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. Virginia Tech, 102–107.
[8] 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*. 154–159.
[9] Jack Atherton and Ge Wang. 2020. Doing vs. Being: A philosophy of design for artful VR. *Journal of New Music Research* 49, 1 (2020), 35–59.
[10] Jeremy Bailenson. 2018. *Experience on demand: What virtual reality is, how it works, and what it can do*. WW Norton Company.
[11] Alejandro Beacco, Ramon Oliva, Carlos Cabreira, Jaime Gallego, and Mel Slater. 2020. Disturbance and plausibility in a virtual rock concert: A pilot study. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. IEEE, 538–545.
[12] Louis Cacciuttolo. 2021. *Welcome to the other side: How a social vr concert production broke the world's record of audiences across all media*. 1–2.
[13] Eunhee Chang, Hyun Taek Kim, and Byounghyun Yoo. 2020. Virtual reality sickness: a review of causes and measurements. *International Journal of Human–Computer Interaction* 36, 17 (2020), 1658–1682.
[14] Yasin Farmani and Robert J Teather. 2020. Evaluating discrete viewpoint control to reduce cybersickness in virtual reality. *Virtual Reality* 24 (2020), 645–664.
[15] Victoria Groom, Jeremy N Bailenson, and Clifford Nass. 2009. The influence of racial embodiment on racial bias in immersive virtual environments. *Social Influence* 4, 3 (2009), 231–248.
[16] Rob Hamilton. 2019. Collaborative and competitive futures for virtual reality music and sound. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 1510–1512.
[17] Rob Hamilton. 2019. Coretet: a 21st century virtual interface for musical expression. In *14th International Symposium on Computer Music Multidisciplinary Research*. Springer, 1010–1021.
[18] Tatsuyoshi Kaneko, Hiroyuki Tarumi, Keiya Kataoka, Yuki Kubochi, Daiki Yamashita, Tomoki Nakai, and Ryota Yamaguchi. 2018. Supporting the sense of unity between remote audiences in vr-based remote live music support system ksa2. In *2018 IEEE international conference on artificial intelligence and virtual reality (AIVR)*. IEEE, 124–127.
[19] Takashi Kasuya, Manabu Tsukada, Yu Komohara, Shigeki Takasaka, Takuhiro Mizuno, Yoshitaka Nomura, Yuta Ueda, and Hiroshi Esaki. 2019. LiVRation: Remote VR live platform with interactive 3D audio-visual service. In *2019 IEEE Games, Entertainment, Media Conference (GEM)*. IEEE, 1–7.
[20] Jaron Lanier. 2017. *Dawn of the new everything: Encounters with reality and virtual reality*. Henry Holt and Company.
[21] Eito Murakami, John Burnett, and Ge Wang. 2024. RayTone: A Node-based Audiovisual Sequencing Environment. In *Proceedings of the International Conference on New Interfaces for Musical Expression*. 627–631.
[22] Kelsey E Onderdijk, Lies Bouckaert, Edith Van Dyck, and Pieter-Jan Maes. 2023. Concert experiences in virtual reality environments. *Virtual Reality* 27, 3 (2023), 2383–2396.
[23] Jieun Park, Youjin Choi, and Kyung Myun Lee. 2024. Research Trends in Virtual Reality Music Concert Technology: A Systematic Literature Review. *IEEE Transactions on Visualization and Computer Graphics* (2024).
[24] Tabitha C Peck, Sofia Seinfeld, Salvatore M Aglioti, and Mel Slater. 2013. Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and cognition* 22, 3 (2013), 779–787.
[25] Simon Pfafff, Olav Lervik, Reto Spoerri, Eleonora Berra, Margarete Jahrmann, and Martin Neukom. 2018. *Games in concert: collaborative music making in virtual reality*. 1–2.
[26] Stefania Serafin, Cumhur Erkut, Juraj Kojs, Niels C Nilsson, and Rolf Nordahl. 2016. Virtual reality musical instruments: State of the art, design principles, and future directions. *Computer Music Journal* 40, 3 (2016), 22–40.
[27] Dan Trueman. 2007. Why a laptop orchestra? *Organised Sound* 12, 2 (2007), 171–179.
[28] Luca Turchet, Rob Hamilton, and Anil Çamci. 2021. Music in extended realities. *IEEE Access* 9 (2021), 15810–15832.
[29] Ge Wang. 2018. *Artful Design: Technology in Search of the Sublime, A MusiComic Manifesto*. Stanford University Press.
[30] Ge Wang, Nicholas J Bryan, Jieun Oh, and Robert Hamilton. 2009. Stanford laptop orchestra (slork). In *Proceedings of the International Computer Music Conference*.
[31] Ge Wang, Perry R Cook, and Spencer Salazar. 2015. Chuck: A strongly timed computer music language. *Computer Music Journal* 39, 4 (2015), 10–29.
[32] Matthew Wright and Adrian Freed. 1997. Open SoundControl: A new protocol for communicating with sound synthesizers. In *ICMC*.
[33] Nick Yee and Jeremy Bailenson. 2007. The Proteus effect: The effect of transformed self-representation on behavior. *Human communication research* 33, 3 (2007), 271–290.
[34] Anıl Çamcı and Rob Hamilton. 2020. Audio-first VR: New perspectives on musical experiences in virtual environments. *Journal of New Music Research* 49, 1 (2020), 1–7.