SOUNDDRIFT : A LIDAR-BASED SPATIAL AUDIO EXPERIENCE

CREATING PHYSICALLY-ACCURATE SOUND ENVIRONMENTS IN VR

CSE 493V - Virtual Reality Systems
Sai Sunku & Harshitha Rebala

Introduction

While VR often prioritizes visual fidelity, spatial audio is equally vital for immersion. Traditional spatial audio methods rely on approximations or manual modeling, limiting realism. The challenge lies in accurately capturing how sound interacts with real-world environments.

Motivation & Related Work

Our research builds upon spatial audio in VR, LiDAR environmental scanning, and acoustic simulation. We leverage the Meta XR Audio SDK's spatial audio capabilities while extending its functionality through environmental data.
Key Technologies We Build On:

- Apple iOS LiDAR sensors and ARKit's scene reconstruction
- Meta XR Audio SDK's spatial audio rendering
- Geometric acoustic methods and ray-tracing techniques
- Previous work in real-time 3D mapping and LiDAR reconstruction

For sound propagation modeling, we adapt simplified versions of architectural acoustic prediction systems like ODEON and CATT-Acoustic for real-time VR applications.

Our Approach

- **Two-Part System:**
  - iOS application captures and processes LiDAR data using ARKit
  - Meta Quest VR application renders the acoustic environment
- **Mesh Processing Pipeline:**
  - Implements outlier removal to eliminate noise
  - Uses voxel grid filtering to reduce data complexity
  - Applies adaptive downsampling to balance detail and performance
  - Performs post-processing (spike removal, Laplacian smoothing)
- **Real-World to VR Data Transfer:**
  - Transfers processed environmental data via local network connection
  - Utilizes Meta XR Audio SDK's raytraced acoustics
  - Simulates realistic sound propagation based on actual room properties
- **User Experience:**
  - Allows placement of virtual sound sources within scanned environments
  - Creates physically accurate sound reflections and reverberations
  - Delivers significantly more realistic audio than generic VR presets

Pipeline & Process

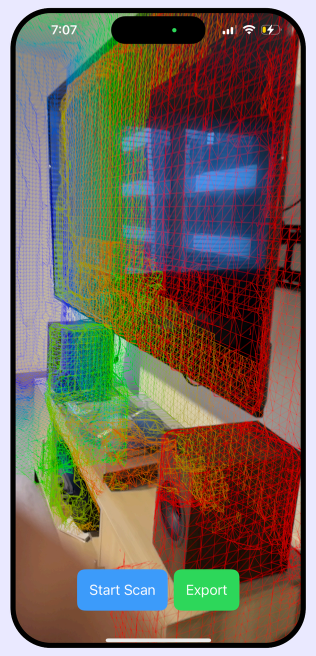

1 Use LiDAR Sensor on your iOS device to capture your enviroment

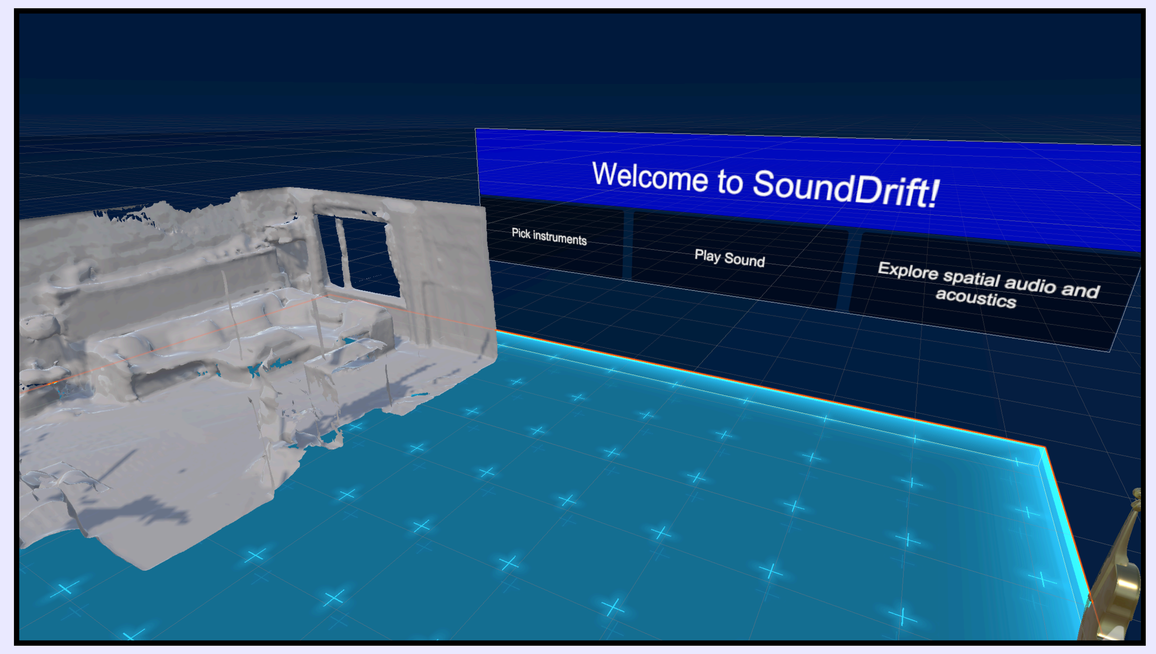

2 LiDAR Scan is processed and rendered in the Meta Quest Headset

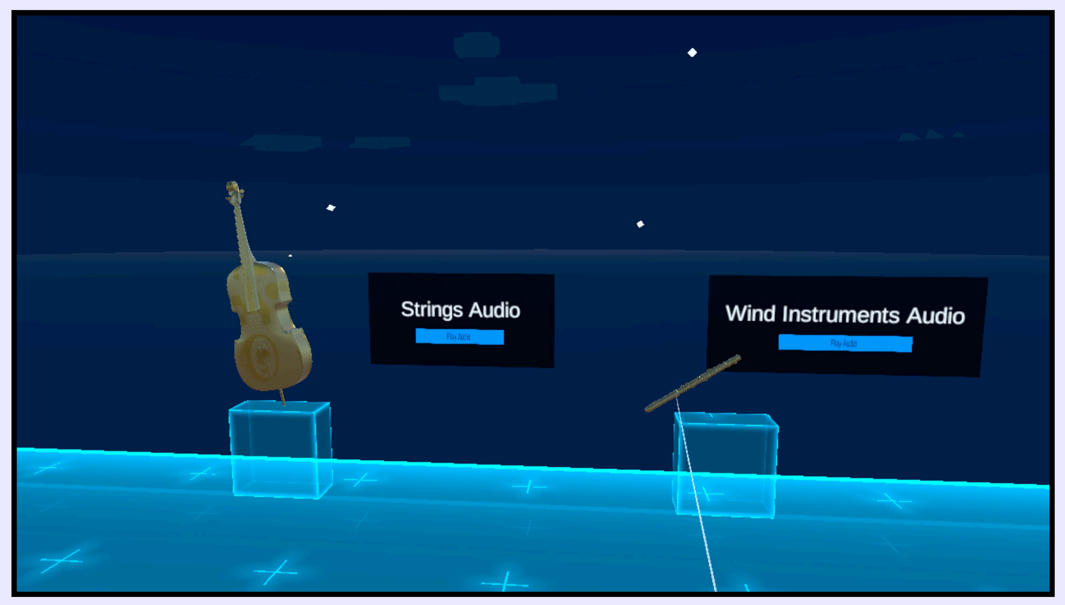

3 Interact with instruments and their audio within LiDAR-scanned environments

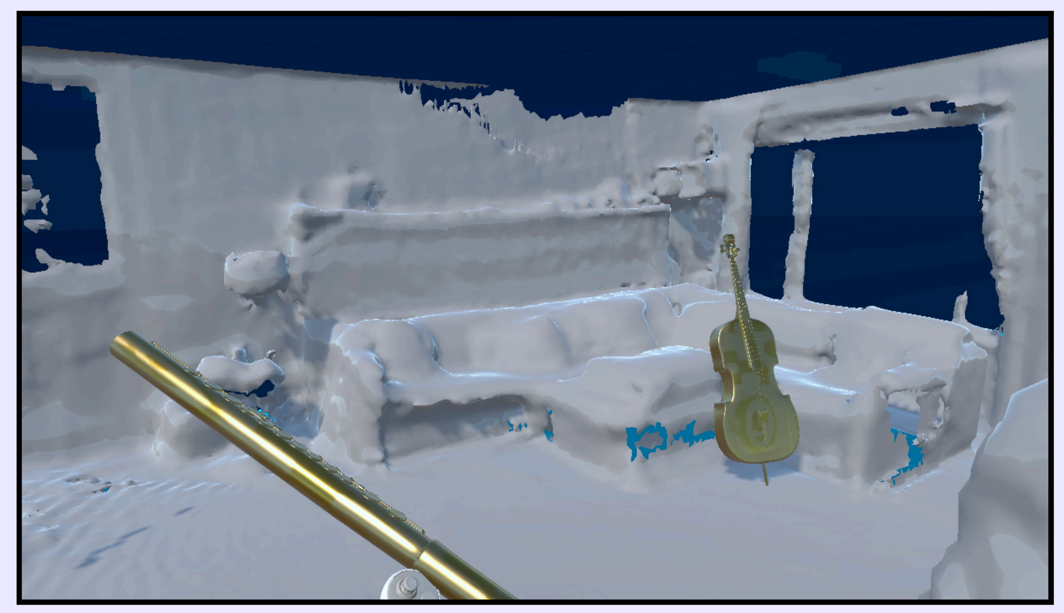

4 Listen to audio that simulates the real-world by accounting for distance, material reflections, and room acoustic

Results

- Our post-processing pipeline effectively reduced artifacts and improved acoustic simulation quality, though processing time (several seconds on iPhone 14 Pro) remains a limitation for real-time applications.
- User studies revealed that participants could accurately distinguish sound source distances, room sizes, and sound directions based solely on acoustic cues, confirming the system's perceptual accuracy.
- Key benefits include environment-specific acoustics that reflect the unique properties of physical spaces, automatic setup without manual calibration, and material-aware processing for nuanced sound environments.
- Limitations persist in material classification accuracy, preservation of small acoustic features, and real-time updates for moving objects.
- Future work will focus on enhancing real-time capabilities and improving material classification accuracy.

References

- Apple Inc. (2024). "Scene Reconstruction." developer.apple.com
- Meta. (2024). "Meta XR Audio SDK Documentation." developers.meta.com
- Thakur, A., Rajalakshmi, P. (2024). "Real-time 3D Map Generation Using Multi-channel LiDAR." IEEE I2CT.
- Tachella, J. et al. (2019). "Real-time 3D reconstruction from single-photon lidar." Nature Communications.
- Savioja, L., Svensson, U.P. (2015). "Overview of geometrical room acoustic modeling." JASA.
- Schissler, C., Manocha, D. (2017). "Interactive sound propagation for multi-source scenes." ACM TOG.
- Valem Tutorials. (2024). YouTube Channel. [8] Unity forums on ARKit integration.

Special thanks to Professor Douglas Lanman for helping us!