

Can Realtime Raytraced Audio in Virtual Environments Improve Spatial Reasoning?


GDEV60001 GAMES DEVELOPMENT PROJECT

RYAN MOAKES

PRIMARY SUPERVISOR: KIERAN OSBORNE

SECONDARY SUPERVISOR: PETER COOPER

1 ABSTRACT



Raytraced Audio is undergoing many recent developments in the study and implementation of spatial audio systems for video games. This study furthers the understanding of how different audio systems inside video games interact with users' spatial awareness.

This study focuses on **two** main audio systems that have taken place, pre-existing spatial audio systems that modern game engines perceive as the default sound technology, and less implemented ray-traced audio systems. This study first conducts a literature review to understand and discuss the current literature of spatial audio systems, including traditional game audio pipelines and toolchains, whilst also breaking ground on what audio raytracing is and its fundamentals.

This study then moves on to how spatial awareness could be tested/researched, and how a ray-traced audio framework has been developed to allow for the testing of spatial awareness. The raytraced audio framework consists of a player 'audio listener' and an audio source 'audio emitter' where the player can move around inside the virtual environment to simulate different audio outcomes. This framework also supports spatial audio simulations. These two audio systems have been tested against twenty user participants to gather information on how different spatial audio systems affect user spatial awareness, with a focus on the potential benefits of ray-traced audio.

The findings of this study showcase that there was a **minor benefit** to using raytraced audio systems for improved spatial reasoning, but due to changes in participant hearing group range, a low sample size and considerations with margins of testing error, the potential benefits of raytraced audio to user spatial awareness remain inconclusive until further study on the matter has been achieved.

2 ACKNOWLEDGEMENTS

The completion of this paper would not have been possible without the help of my wonderful friends and family supporting me with the difficulties I face every day. As such, I would like to express my gratitude towards the following people:

***John Moakes and Natalie Moakes** for being the best parents I could ask for.*

***Kieran Osborne** for his mentorship throughout this paper.*

***Louise Ramsden** for always being herself.*

***Cassette Costen** for keeping me company throughout the year.*

***Ryan Rudge** for helping me edit and draft this paper.*

***Tim Hook and Jai Paton** for being amazing all the time.*

***Connor Ratcliff, Lewis Greenwood, Paul Harrison, Harry Hamilton, Frankie Golden and Ian Respeto** for all the encouragement they always provided me with.*

3 CONTENTS

1	Abstract.....	1
2	Acknowledgements.....	2
3	Contents	3
4	Table of Figures	5
5	Introduction.....	6
6	Aims and Objectives	7
7	Literature Review.....	8
	(1.0) Introduction:	8
	(2.0) Discussion of Traditional Game Audio Techniques:	9
	(2.1) Overview of Traditional Game Audio Pipelines:	9
	(2.2) Common Techniques for Spatial Audio Simulation:	9
	(2.3) Sound Propagation Systems:	10
	(2.4) Limitations of Traditional Approaches:	11
	(3.0) Pre-Existing Game Audio Techniques on User Spatial Awareness:	11
	(3.1) Defining Spatial Awareness in Game Audio:	11
	(3.2) Impact of Traditional Techniques on Player Experience:	11
	(4.0) Discussion of Raytracing-Based Game Audio Techniques:	12
	(4.1) Fundamentals of Acoustic Raytracing:	12
	(4.3) Benefits of Raytraced Audio:	13
	(4.4) Challenges and Implementation Limitations:.....	14
	(5.0) Performance Differences Between Traditional Audio Techniques & Raytraced Audio Techniques:.....	14
	(5.1) Computational Performance and Resource Usage:	14
	(5.2) Realtime Feasibility and Latency Considerations	15
	(6.0) Integration of Raytraced Game Audio Techniques into Game Engines:	15
	(6.1) Overview of Current Integration Approaches:	15
	(6.2) Developer Implementation Challenges:	16
	(7.0) Research Gap & Conclusion	16
	(7.1) Gaps in Current Raytraced Audio & Spatial Awareness.....	16
	(7.2) Conclusive Piece	17
8	Research Methodologies	17
	8.1 Research Paradigm	17
	8.2 Testing Framework & Artefact Creation.....	17
	8.3 Materials and Equipment.....	19
9	Results and Findings	22
	9.1 Participant Hearing Range Frequency	22
	9.2 Spatial Audio Tests.....	23

9.3	Raytraced Audio Tests	26
10	Discussion and Analysis	29
10.1	Anomalous Results	29
10.2	Participant Hearing Range Analysis	29
10.3	Analysis Spatial Audio & Raytraced Audio	29
10.4	Criticisms of The Investigation	31
11	Conclusion	32
11.1	Aims and Objectives Checklist	32
11.2	Spatial Reasoning Audio System	33
12	Recommendations	33
12.1	Gaps in Knowledge	33
12.2	Sampling Participants	33
12.3	Raytraced Audio Implementation	34
13	References	35
14	Appendices	38
14.1	User Consent Sheet	38
14.2	User Information Sheet	39
14.3	Average Data Collection Sample	41
14.4	Testing Framework Pre-Test Scenes	42
14.5	Spatial Audio Raw Testing Data	43
14.6	Raytraced Audio Raw Testing Data	43
14.7	Spatial Audio Final Testing Data	44
14.8	Raytraced Audio Final Testing Data	44

4 TABLE OF FIGURES

Figure 1: A general overview of the game audio tool chain.	9
Figure 2: Diagram of a Head Related Transfer Function or HRTF.	10
Figure 3: A “near” linear graph showcasing global illumination rays.....	15
Figure 4: A screenshot of the ray-traced testing environment.....	23
Figure 5: University of Staffordshire (2026) standard games computers	20
Figure 6: Corsair Virtuoso SE (2019) headphones that participants used	21
Figure 7: A screenshot of the main menu of the developed testing framework.....	21
Figure 8: A box plot graph of the user hearing ranges (Hz)	22
Figure 9: A bar chart graph on the average hearing range.....	22
Figure 10: The spatial audio ‘total duration’ results (scatter graph).....	23
Figure 11: The spatial audio ‘total duration’ results (bar graph)	24
Figure 12: The spatial audio ‘duration per emitter’ results (scatter graph).....	24
Figure 13: The spatial audio ‘duration per emitter’ results (bar graph)	24
Figure 14: The spatial audio ‘heading alignment accuracy’ results (scatter graph)...	25
Figure 15: The spatial audio ‘heading alignment accuracy’ results (bar graph).....	25
Figure 16: The raytraced audio ‘total duration’ results (scatter graph).....	26
Figure 17: The raytraced audio ‘total duration’ results (bar graph)	26
Figure 18: The raytraced audio ‘duration per emitter’ results (scatter graph).	27
Figure 19: The raytraced audio ‘duration per emitter’ results (bar graph).....	27
Figure 20: The raytraced audio ‘heading alignment’ results (scatter graph)	28
Figure 21: The raytraced audio ‘heading alignment’ results (bar graph).....	28
Figure 22: The results of the raytraced/spatial users ‘total duration’ values.....	30
Figure 23: The results of the raytraced/spatial users ‘duration per emitter’ values ...	31
Figure 24: The results of the raytraced/spatial users ‘heading alignment’ values	31

5 INTRODUCTION

The project showcases the creation of a realtime raytraced audio simulation artefact that will be used to gather and track accurate audio data of a simulated/virtual environment. One of the goals of this executable will be that it is able to process audio raytraced data in a real-time environment, an environment which can be manipulated by the user to provide different audio outputs that are modified by properties. Some of these properties include the shape of the environment, the surface properties of the environment and the distance between the audio output/audio receiver.

Alongside a literature review, that will not only analyse traditional audio techniques and technologies that are currently used in video game contexts but will also dive into new and emerging audio technologies like raytraced audio, this project will conduct research through user-participated testing. This data will be gathered and analysed with the intention of answering the research question presented below:

Problem: Can Realtime Raytraced Audio in Virtual Environments Improve Spatial Reasoning?

Breaking down this research question into its different segments allows a better understanding of what is being evaluated in this project. Firstly, **‘Realtime Raytraced Audio’** refers to the tracking and simulation of audio “waves/rays” to follow and change their properties in real-time to gather accurate sound data. Gathering the results of the sound/audio raytraced data and using it for real-time headphone/speaker output. This would allow for dynamic and environmental sound effect changes on the fly; for example, a sound designer could make a room bigger to get a smaller echo.

Secondly, **‘Virtual Environments’** refers to the creation and simulation of virtual environments and surfaces to see the effect and changes on sound/audio properties. This would result in different audio output effects, as previously mentioned (echo, etc.), but would also result in the modification of audio through these surface properties. Common surfaces such as wood, glass, stone and other soundproofing surfaces would all have different effects on the output of the audio produced.

Lastly, **‘Improve Spatial Reasoning’** refers to inviting User Participants to test against a set list of test parameters to identify the change that raytracing audio has on spatial reasoning (better object/environment/enemy sense in a space). This is the main data gathering scenario that will facilitate the analysis required to answer the proposed question.

6 AIMS AND OBJECTIVES

This project aims to complete the following objectives to help and further the understanding of ray-traced audio technologies and techniques; the spatial perception of users who utilise these techniques will also be studied and discussed. Testing data will be collected and analysed as part of this process, and a literature review will be conducted to thoroughly understand the topic at hand, which will inform how this paper addresses the questions presented.

Desired Outcomes by the end of the Project:

- A Realtime Sound Simulation Artefact supporting the Research Question Proposed.
- A collection of user-related information on the improvement/changes to spatial reasoning under different audio technologies, such as acoustic raytracing compared to simpler distance-based audio techniques.
- A literature review on the topic of audio-based techniques utilised in games and how raytraced audio can be used for the benefit of improving spatial perception.
- A series of tests performed in the Realtime Sound Simulation Artefact to gather user-related information.
- A discussion on the results and findings of user testing regarding whether spatial perception was improved by one audio technique more than another.
- A conclusion on the analysis to answer the proposed research question.
- A series of recommendations to anyone looking to further the study of ray-traced audio technologies in a general and gaming environment.

Realtime Sound Simulation Artefact – An executable application that allows the real-time raytracing of sound in a simulated environment; this will be used to determine if real-time acoustic raytracing is effective in a game-related context and if the technology can improve users' experience via increased capability for spatial reasoning.

User Related Information – A collection of user-based data and feedback to see the effects that real-time audio raytracing has on spatial reasoning in a virtual environment.

Literature Review – A review of the current state of audio techniques in games and a discussion of the implementation of ray-traced audio.

User Tests on the Sound Simulation Artefact – A series of user-based tests to gather the previously stated user information.

Results and Findings Discussion – A discussion looking at the raw data gathered from the research that has been conducted, and how an interpretation of it can inform future game developers and audio engineers.

Conclusions Discussions – Gathering answers and resolving the research question.

Recommendations for Future Studies – Forming new questions and analysing gaps in the topic.

7 LITERATURE REVIEW

(1.0) INTRODUCTION:

This review stands as a bridging point for those who hope to learn more about game audio programming techniques, whilst also showing interest in the future of game audio programming with the possible introduction of 'realtime raytraced audio' techniques being seen in a wider video game context (Terletskyy and Trushevskyy, 2024). It will also briefly discuss themes of user spatial awareness that are relevant to the higher purpose of this paper.

Modern-day video games have many different interpretations of audio techniques that are used in conjunction with each other to create a realistic and seamless experience for players, whilst also improving players' spatial perception in games which take place in virtual environments. (Cao et al., 2023). Although there is a lot of discussion on the currently used and existing audio techniques (Guy Sombeg, 2024), the analysis of audio techniques, including 'raytraced audio' in a real-time gaming environment, has not seen a significant amount of research. The following review on game-based audio techniques will address this issue by first discussing pre-existing audio techniques, identifying their effects on user spatial perception; then moving on to how 'raytraced audio' could be utilised with brief discussions into game engine integration, performance concerns and benefits/drawbacks, hoping to expand on topics mentioned by Samsurya, 2023 & Terletskyy, 2024.

The review also examine a wide range of the literature available on the current state of audio techniques and technologies used in the current generation of video games, with a deep discussion into the implementation/effects of 'raytraced audio'. Looking at how it has been previously implemented, whilst also looking at how it could be adapted to more modern realtime interpretations.

This review comprises six main sections, with each section having sub-sections to better focus the current discussion at hand. Each section will examine the topic in an in-depth and detailed discussion with the support of academic and peer-reviewed literature. The following table briefly details the contents of each section:

Section Number	Section Description
Section 2.0	Discussion of traditional game audio, including common techniques and their associated limitations.
Section 3.0	Exploration of user spatial awareness within virtual environments and how traditional audio techniques influence this.
Section 4.0	Examination of the fundamentals, implementation methods, and challenges of <i>ray-traced audio</i> in real-time gaming contexts.
Section 5.0	Analysis of performance differences between traditional audio systems and <i>ray-traced audio</i> implementations.
Section 6.0	Discussion on how <i>ray-traced audio</i> techniques might be integrated into existing software frameworks.
Section 7.0	Identification of the current research gap and recommendations for future development.

Table 1: A table containing the section number and a description of each segment of the literature review.

(2.0) DISCUSSION OF TRADITIONAL GAME AUDIO TECHNIQUES:

(2.1) Overview of Traditional Game Audio Pipelines:

Game audio pipelines are a commonly overlooked factor in modern-day game engine implementations, as Lindskev and Riddersholm (2019) discuss. Whilst there may be quite a lot of information on game audio sound design, there is a lack of information on the subject of game audio workflow and pipelines. The simplest and most common game audio pipeline for modern audio implementations includes the standard design and use of audio inside game engines, but also one additional measure known as the middleware stage. (Beauchemin, 2023) (Hooiveld, 2025) (Lindskev and Riddersholm, 2019). The middleware stage is where audio passes through embedded software inside the game engine for processing any additional effects. This software also allows for audio manipulation and gives audio programmers some features and flexibility quickly, instead of requiring lots of prior implementation from previous audio libraries.

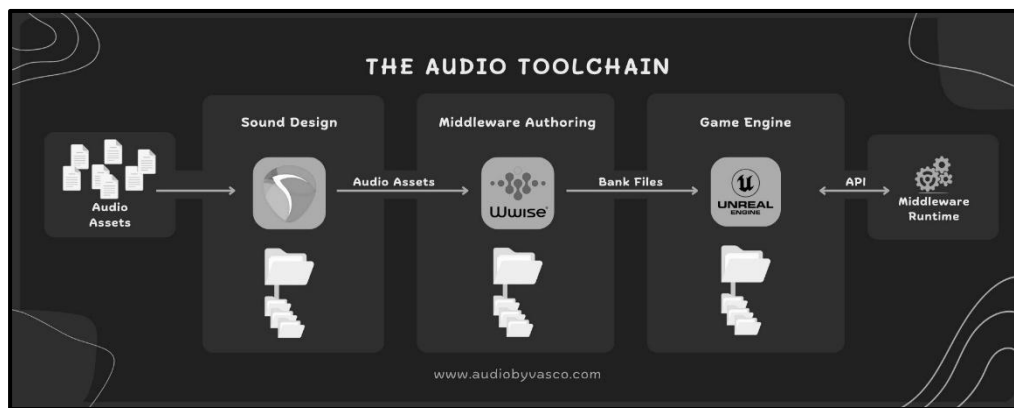


Figure 1: A general overview of the game audio toolchain for Unreal Engine. Showcasing the developer sound design stage, the middleware authoring stage and the final game engine utilisation stage (Hooiveld, 2025).

Some of these middleware stages include software such as Wwise, Csound and FMOD. Industry standard middleware featured in both Unity and Unreal game engines seems to lean more towards the use of Wwise, but Beauchemin (2023), Lindskev and Riddersholm (2019) all argue that the use of the engine's basic sound development toolkit could be enough for any dedicated audio programmer.

(2.2) Common Techniques for Spatial Audio Simulation:

A big component of spatial audio simulation is the use of HRTF's (Head Related Transfer Functions) to convert stereo audio to accurate positional/spatial sound that sounds reasonably realistic to the average user (Zhang et al., 2009). For human sound localisation to occur, three dimensions of hearing need to be accounted for and simulated. These dimensions are range/distance to determine how far the audio source is away from the perspective of the user, a directional vector from the audio source to determine where the audio source is coming from or being played and finally, which ear is hearing the audio source or sometimes referred to as azimuth (Zhang et al., 2009). This is why head-related transfer functions are calculated, to produce a sound solution which correctly corresponds to how an ear would perceive a sound coming from a point in space. Using a pair of HRTF's, a binaural sound can be created, which can be used for accurate simulation of sound inside a space; this

effect, combined with basic sound attenuation, can be used for the very simple approximation of spatial audio inside a space.




Figure 2: Diagram of a Head-Related Transfer Function or HRTF (Zhang et al., 2009).

Another technique for the simulation of spatial audio is the use of the ‘Doppler Effect’. A real-life effect that needs to be recreated to provide the user with more sound-based realism, as further talked about by Beauchemin (2023). The Doppler Effect describes a change in the sound wave frequency between a stationary observer or audio listener and a moving audio source or emitter. This effect is usually seen in the change of pitch of a sound frequency, with common examples including a stationary listener and a car or train. (Gimenez et al., 2007). From the perspective of the user, the audio source is perceived approaching or getting farther away by from them, by the sound frequencies being increased or reduced, changing the resulting spatial sound.

This effect can be calculated using the following **Equation [1]** below and is vital for the simulation of accurate spatial sound.

$$f = \left(\frac{v_m \pm v_r}{v_m \mp v_s} \right) f_0$$

Equation [1]: Calculation of the Doppler effect (not at the speed of light), where f is the user observed frequencies, f_0 is the observed frequency at the point of audio emission, v_m is the speed of the wave through a medium, v_r is the speed of the observer and v_s is the speed of the audio emitter (Davari, 2021) (Beauchemin, 2023).

HRTFs and the Doppler effect are not the only two techniques of simulating realistic spatial audio but are major contributors to the creation of spatial audio.

(2.3) Sound Propagation Systems:

Schissler (2011) discusses how sound propagation can be split into three separated components:

1. Direct Sound

2. **Early Reflections**
3. **Late Reverberation**

Direct sound is the sound waves that are directly transmitted from the audio source/emitter to the audio receiver. Early reflections are the first echo of the sound waves that reach the audio listener. Late reverberation is the last component of the sound wave heard by the audio listener; it comprises ‘higher order’ reflections to produce a decaying sound.

(2.3) Limitations of Traditional Approaches:

Many of these traditional approaches have significant limitations which arise for the simulation of spatial audio in virtual environments. Two main limitations are raised, including audio accuracy and dynamic audio design (Geier et al., 2010). The realism for audio accuracy is lost when cheaply simulating spatial audio, as many factors are not considered, obstacles and geometry are usually left out of audio mixing and produce unrealistic sound in some situations – sound not being dampened by the environment, etc. The lack of dynamic scaling audio systems is another issue; sound design/programmers are commonly left with the implementation of per audio emitter effects, which slows down development time significantly. For example, if an audio emitter is in a cave-like environment, the echo and reverb of the sound should be intensified; instead of an audio emitter understanding this implicitly, the sound designer will have to add implementation effects to maintain realism (Rickard, 2020).

(3.0) PRE-EXISTING GAME AUDIO TECHNIQUES ON USER SPATIAL AWARENESS:

(3.1) Defining Spatial Awareness in Game Audio:

At its core, spatial awareness within the context of virtual spaces is the user’s ability to understand the environment they are presented with and their state within it (Nair et al., 2022), and how they use visual and audio information to reason the relationship between entities in that space. In terms of spatial awareness through audio, several uses aid the player in various ways. For example, the scale and shape of an area can be reasoned through how audio interacts with the environment. In a large chamber area, the user will be capable of reasoning that the area is of a larger scale with well-designed audio through hearing changes in late reverberations, as they will take a longer period to reflect to the user. The increased delay in sound can be used to intuitively determine the approximate scale (Schissler, 2011).

Understanding the wider environment is also key for a user to be spatially aware. Audio feedback relevant to surrounding areas can be a key component in communicating surrounding areas to the player; for example, the sound of rolling waves to communicate to the player that they are near a coastline. This process allows users to cross-reference other cues to reason their location within a larger area, allowing them to know their general position without the need for more specific aids, such as maps. (Nair et al., 2022). Spatial awareness is highly important within games as it directly affects the user’s ability to interact with the environment provided to them effectively and serves as a form of communication between the game and the user, relaying information for them to consider and act upon to complete the gameplay loop.

(3.2) Impact of Traditional Techniques on Player Experience:

Traditional techniques in audio tend to follow the principle of creating the illusion of a realistic audio field by performing simple left-right sound balancing and reducing volume of sounds coming from more distant emitters. This effect is created by

leveraging the human body's real world sound orientation techniques, by detecting changes in how sound waves interact with the head, torso, shoulders and ears, the brain can detect the distance and direction of the sound, thus by playing a sound louder in the left side of the speaker/headphone the user can experience that sound as if originated from that direction. (Paterson, et al., 2010). Through this technique, the user can be more heavily immersed in a scene and are able to more accurately detect the direction and distance of a sound cue and be able to react accordingly to the inferred information this affords them.

Reverberation has a similar effect immersion upon the player - sound cues reflect off of physical surfaces, and as a result, those cues will arrive at the player from different angles with varying degrees of delay. Creating an envelopment around the player helps to convey the relative depth of the scene before them, as well as its shape. Not only is this helpful for reasoning within the space, but it also helps to immerse the player by simulating the acoustics of the environment they are traversing within the virtual space. (Paterson, et al., 2010). Ambient noise also serves a very similar purpose in this regard, although mostly useful for thematic environmental storytelling. Audio ambience helps the player to feel central to the environment, encasing them in the low-level noises expected of any environment.

(4.0) DISCUSSION OF RAYTRACING-BASED GAME AUDIO TECHNIQUES:

(4.1) Fundamentals of Acoustic Raytracing:

Using acoustic raytracing, an audio source within a virtual environment will propagate rays based on the given ray sampling technique; these rays serve to simulate the movement of sound waves throughout a space (Mueller, 1999). When the ray detects a collision with an obstacle, the ray will experience reflection, absorption and diffraction around the obstacle. Firstly, the ray will reflect off the surface, redirecting the ray in line with the law of reflection as seen in **Equation [2]**.

$$\theta_i = \theta_r$$

Equation [2]: The Law of Reflection, stating that the angle of incidence is equal to the angle of reflection (Terletskey and Trushevskyy, 2024).

Secondly, the direction of the reflected ray is calculated through the formula seen in **Equation [3]**.

$$\mathbf{r} = \mathbf{i} - 2(\mathbf{i} \cdot \mathbf{n})\mathbf{n}$$

Equation [3]: Ray Direction calculation where \mathbf{i} is equal to the incident ray direction, and \mathbf{n} is equal to the surface normal (Terletskey and Trushevskyy, 2024).

Additionally, upon contact with the obstacle, the ray will, in most optimised systems, lose 'energy', which is a value generally assigned at propagation, generally assigned a value of one, which will decrease to zero as the ray travels throughout the scene or will stop at its current value when an audio emitter is encountered. (Samsurya, 2023). This value will also decrease on contact with an obstacle based on its material, with a more extreme decrease for materials with higher absorption values.

The sound absorption of the audio ray takes place by using an absorption coefficient – between zero and one – to determine the dampening of the audio ray as seen in **Equation [4]**.

$$I_{\text{reflected}} = (1 - \alpha)I_{\text{incident}}$$

Equation [4]: Intensity of reflected sound calculation, where α is the absorption coefficient (Terletskeyy and Trushevskyy, 2024).

As the ray's energy value decreases, the volume of the played noise from that audio emitter decreases in tandem. This dramatically helps in creating a realistic audio field, as this principle calculates volume based on how far the listener is from the audio source, as well as the properties of the materials that the ray reflected off.

Lastly, rays experience diffraction around obstacles, which is a crucial step to simulate a realistic audio field. As the rays approach a corner of an object, the sound waves will scatter, often bending behind them to simulate a realistic sound wave. (Beauchemin, 2023). The extent of this bend depends primarily on the wavelength; longer wavelengths are less susceptible to this feature, while shorter wavelengths will bend more dramatically. This is calculated in Maekawa's empirical formula, which is widely used in acoustics. (Maekawa, 1968). Maekawa states that the diffraction loss is calculated via the size of an obstacle and the wavelength of the sound as seen in **Equation [5]**.

$$L_d = 10 \log_{10} \left(3 + 20 \left(\frac{\lambda}{D} \right) \right)$$

Equation [5]: Where λ is the sound wavelength, and D is the effective size of an obstacle, which describes the wavelength's ability to diffract in relation to the size of the obstacle (Maekawa, 1968).

In terms of optimisation, and to be computed at all, the system would not allow rays to travel infinitely; this would cause many issues with the system itself if rays were not terminated. (Nvidia, 2018). Thus, optimised acoustic raytracing systems will often include three ray termination methods for the travel length of the ray. Firstly, the energy value. When energy hits zero, the ray should be inaudible from its current endpoint, and if the ray is inaudible, it should be prevented from travelling any further to save computing budget. This can also be defined as the maximum distance the ray can travel. (Samsurya, 2023). Secondly, it is common to include a maximum bounce count - max recursive depth - which not only prevents cases of rays bouncing infinitely in small spaces, which would crash the program, but also helps reduce latency by reducing the number of reflection calculations that must be run in real time. (Samsurya, 2023). Lastly, the sound wave will often terminate when it hits the audio emitter, as its journey to find an audio emitter will have finished, and that audio has been confirmed as heard by an observer. (Lixandru, 2021).

(4.2) Benefits of Raytraced Audio:

One of the purposes of realtime raytraced audio is to simulate the activity and behaviours of sound waves as closely as possible while still working with a real time solution; by propagating rays in every direction the source creates a more realistic field of audio due to the diffraction interaction with objects, scaling volume through walls based on the objects relative size against the wave length of the sound (Terletskeyy and Trushevskyy, 2024). In addition, each ray's interaction with objects based on material properties helps prevent unrealistic audio outcomes without much manual interaction (Samsurya, 2023). Another benefit of ray-traced audio is the automation of the absorption of sound in real time. This removes the need to add filters or change the desired audio file to fit the surrounding materials, therefore streamlining the development of virtual spaces – especially within game environments where audio sources can often be dynamic (Guy Sombeg, 2024).

This helps save memory constraints by preventing the need for multiple versions of the same audio file to be played in different scenarios to keep the player immersed.

Additionally, ray-traced audio is particularly effective in dense scenes where the rays can create a realistic field of audio without the need for manually or algorithmically placed audio zones, which can become extremely complex in areas such as dense woodland, streamlining development as well as general audio realism in these areas. (Samsurya, 2023).

(4.3) Challenges and Implementation Limitations:

Global illumination ray tracing is an inherently high-cost method in terms of computational budget, which can create several issues, including development optimisation trade-offs and technological accessibility issues for consumers (Nvidia, 2018). However, the same cannot be said about audio raytracing. Audio raytracing can not only function by itself but also can run much more efficiently when utilising the techniques that were originally implemented for global illumination raytracing (Terletskyy and Trushevskyy, 2024). Firstly, when using ray tracing in a real time, previously mentioned optimisations to max bounces - max recursive depth - are often a necessary concession. With each bounce, the cost of the audio field increases the computational load, as more updates, checks and calculations are forced to occur for each ray. While this can be mitigated to an extent with ray-caching, other such systems to prevent the need for constant recalculations and updates; these approaches are often useless for moving dynamic audio sources due to the nature of caching. It is also important to note that an infinite ray recursive depth is not feasible with any level of modern hardware and would cause a stack overflow. (Nvidia, 2018).

The aforementioned issues are, however, symptoms of the larger issue of the constraints of real-time computation. Ray tracing is a dynamic system that updates constantly to provide the user with the most realistic audio experience possible; thus, the system must constantly update all ray-traced sources, while concessions can be made, such as finding methods to remove audio emitters that the player is out of range from or restricting diffraction. These methods can indirectly affect the quality and immersion of the system, sometimes defeating the purpose of the audio system in the first place. (Lixandru, 2021).

While global illumination raytracing might raise some accessibility concerns, forcing users to buy expensive and modern hardware like an RTX 3060 graphics card as a recommended minimum (Nvidia, 2018), Audio raytracing does not require the same player investment, due to its much further scaled-back raytracing demands. With all of this considered, raytraced audio scalability and user accessibility are much less of an issue for audio raytracing compared to global illumination raytracing. More comparison on global illumination raytracing and audio raytracing will be discussed further in section (5.2).

(5.0) PERFORMANCE DIFFERENCES BETWEEN TRADITIONAL AUDIO TECHNIQUES & RAYTRACED AUDIO TECHNIQUES:

(5.1) Computational Performance and Resource Usage:

An initial concern for audio raytracing is the real-time performance any modern raytracing technique might face; there is no question that modern global illumination raytracing techniques require multiple optimisations to make real-time raytracing an option for players in the first place. Optimisation techniques such as BVH's (Bounding Volume Hierarchies), ray-path caching and other techniques are all vital systems when raytracing through a scene with global illumination raytracing (Nvidia, 2018). Audio raytracing, however, does not require such systems in place (Schissler, 2011). As such, a far lower computational and resource budget needs to be considered

when implementing real-time ray-traced audio compared to ray-traced global illumination techniques – more on this is discussed in the [following section](#).

(5.2) Real-time Feasibility and Latency Considerations

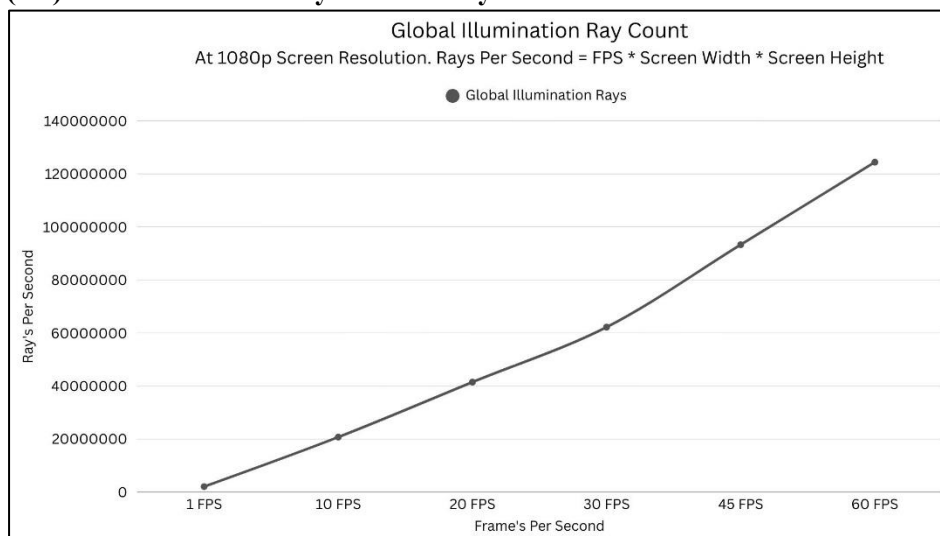


Figure 3: A “near” linear graph showcasing how many GPU-accelerated global illumination rays are necessary for a resolution of 1080p at different frame rates.

Note that in **Figure 3**, the number of rays necessary at 60FPS, 1080p resolution for global illumination techniques is 124416000 RPS (rays per second), the number of rays necessary to be generated and dispatched per second is a 10367900% increase over Terletskyy & Trushevskyy’s recommended 1200 RPS audio ray specification. With these two values considered, it is clear that the use of GPU-accelerated raytracing is unnecessary when performing audio raytracing, and modern CPU architecture is more than capable of handling the small amount of ray generation and audio processing that audio raytracing requires. This is not to say that audio raytracing would not benefit from a GPU accelerated techniques or even other more functional techniques like ray-path caching in a more optimised system.

(6.0) INTEGRATION OF RAYTRACED GAME AUDIO TECHNIQUES INTO GAME ENGINES:

(6.1) Overview of Current Integration Approaches:

Practised techniques within real-time engines are primarily built for ease of use by the developers and to allow efficient sound propagation for simple spatial reasoning in these virtual environments. For example, the current default pipeline within the Unity (2026) game engine involves the insertion of audio data into an ‘audio source component’; the system then uses ‘audio listeners’ to calculate when the user should be able to hear the output of the audio component. When the engine detects the listener is within range of the source, secondary calculations are performed to calculate the intensity of the sound based on their position. The further from the source the listener is, the less intense the sound output is. This falloff is highly customisable using features such as falloff curves and reverb modifiers (Unity, 2026). Additionally, the use of audio programming languages such as Csound is common to apply filters, effects and other useful additions to the output of audio.

Csound specifically is an open-source audio synthesiser and library with a flexible array of tools across both orchestral and real-time audio environments, which is particularly useful for the control over audio outputs that it offers. Engines such as

Unity allow for easy integration of such programs, producing the development of sound quality through techniques such as reverberations and delays that help simulate realistic audio fields. Other engines follow a similar pipeline, such as Unreal Engine (2026). While the process is similar in design, it offers its own in-engine mixers and effects library featuring nodes for similar audio modulation through editable graphs and modifiers (Epic Games, 2026). Notably, no widely used game engine currently supports acoustic ray tracing, although some demonstrations of plugins have been showcased by developers within the communities of these engines. (jungle_Pinpoint, 2025).

(6.2) Developer Implementation Challenges:

Raytracing as a technology has been researched and applied in some fields and industries as far back as the 1970s. The preprocessing stage has historically been colossal in both times to compute and the computational resources to calculate the result. (Allen and Berkley, 1978). Such techniques have been commonly used in architectural processes to understand the results of lighting and acoustics before building occurs. While in this field, the lengthy process of generating a single frame is not much of an issue; within real time environments, use of such systems would lower the performance to completely unacceptable levels, especially considering industry standards of 30-60 frames per second. (Schissler, 2011). The application of this technique has thus remained extremely niche for real-time settings such as video games until recent advancements in computer hardware, keeping both acoustic and global illumination raytracing from being widely used.

Due to audio raytracing in real time being such a recent topic of research, very little prior implementation exists. As mentioned previously, there are no widely used game engine implementations and only niche systems developed by community developers such as @jungle_Pinpoint, 2025. This lack of modern research and development into audio raytracing provides little to work with when creating such a system; while information does exist, the process of constructing the system is much less widely discussed and explored, so designing for the system it will be used in can be a complex process.

Fortunately, however, pre-existing systems of widely used game engines can be repurposed to suit the needs of acoustic ray tracing, as the system of making rays and testing rays is used for other aspects of game development. The system included in Unreal Engine 5 is one such system. This system allows feedback between a start point and an end point to be gathered by detecting collisions and simulating values such as vectors. (Epic Games, 2026). Similarly, Unity uses the physics raycast system, which holds much of the same functionality and can be applied similarly to raytracing. (Unity, 2026). However, using these systems to implement acoustic ray tracing requires developing the system from scratch due to the lack of built-in features to build upon in these engines. (jungle_Pinpoint, 2025).

(7.0) RESEARCH GAP & CONCLUSION

(7.1) Gaps in Current Raytraced Audio & Spatial Awareness

When analysing the current literature on the topic of ray-traced audio in reference to spatial awareness, very little research has been conducted. The raytraced audio topic would benefit from much more on the combination of research into the implementation of raytracing audio and how this audio system affects users' spatial awareness, and other factors that might be relevant for the state of audio technologies for video game players. This paper hopes to act as a starting point for this combination of factors.

(7. 2) Conclusive Piece

In summary, this review aimed to bring together established and emerging game audio techniques within game audio programming, with an emphasis on how realtime raytraced audio may change the player experience and developers' design processes.

This was done by combining the discussion of more traditional game audio techniques, spatial awareness processes, and the broader technical implications of ray-traced audio through critical analysis of the current literature regarding the topic, contributing to the understanding of a variety of different audio technologies and how they can be used in gaming-based contexts.

8 RESEARCH METHODOLOGIES

8.1 RESEARCH PARADIGM

When gathering and researching spatial reasoning /awareness against different audio techniques for this study, mixed research paradigms were used in tandem to achieve all the knowledge required. Using Taherdoost (2022) as a guide, quantitative methods were used. Users participated in a testing framework to gather quantitative data to help see the differences in spatial perception between spatial audio systems and ray-traced audio systems. Users were split into two groups to avoid any learning bias, which is the recommended method for the testing of different systems. (ITU Radiocommunication Assembly, 2001) (ITU-R, 2015) (Taherdoost, 2022).

To understand what the methods for testing spatial reasoning inside a virtual space entailed, Cao et al. (2023) and Vontzalidis et al. (2025) works were analysed for any additional information that might be necessary.

8.2 TESTING FRAMEWORK & ARTEFACT CREATION

Over the course of the project duration, an acoustic raytracing artefact was produced as part of the requirements for the user testing framework that would be used for the user testing/data gathering stage. This audio raytracing system was made with the guidance of Samsurya (2023), Terletskey and Trushevskyy (2024) inside the Unity game engine. Unity has, as previously discussed, contains out-of-the-box raytracing functionality, which meant a lot of the standard raytracing implementation was alleviated from the development process. (Unity, 2026).

An enquiry into the implementation and use of the Csound middleware was also considered for the purposes of creating a testing framework/aspect, as also recommended by Samsurya (2023), Terletskey and Trushevskyy (2024). Csound is an involved audio programming library, and although its integration would have greatly benefited the quality of the finalised raytracing solution, the development time necessary to learn and identify its capabilities would have taken the project duration above the allotted time.

An important note to make for the raytraced audio system developed for this project is that rays permeate the scene from the audio listener instead of the naively expected audio emitter. Although it is noted that ray-traced audio systems can feature one or the other method of ray location generation, opting for the first method mentioned decreases the ray count necessary for simulation, as every audio emitter added to the environment does not need to add more rays, but instead more options for rays to interact with an emitter in the scene.

The finalised raytraced sound solution involved two main audio modifiers that change the user's final heard audio output:

- Volume Modification
- Muffling Modification

The volume modifier was calculated based on how many rays per frame interacted with an audio emitter, and the next calculation verifies the rays' distance travelled before this interaction had taken place. Using these two values, a new volume lerp value can be constructed between zero and one.

The muffling modification to the audio can also be calculated from the values and, in a similar vein, can be used to construct a muffler value to increase or decrease the amount of frequency muffling the audio emitter produces.

Two audio ray sampling techniques were implemented into the raytraced audio solution; these sampling techniques are 'Monte Carlo' sampling and the dubbed 'Golden Ratio' sampling. Using Monte Carlo sampling for ray direction generation was the better method out of the two, as it combined with the audio lerp methods to provide smooth audio whilst the audio listener was moving around. Golden ratio sampling, on the other hand, uses the same ray directions each frame (generated from a golden ratio approximation around the audio listener), which provides much less realistic simulations of the audio when presented to users.

Following the creation of the raytraced sound solution, a testing framework was developed with two sound systems – the raytraced sound system and Unity's ready out-of-the-box spatial audio system. This testing framework was constructed under the written advice of Geier et al. (2010) and ITU Radiocommunication Assembly, (2001).

The testing framework, following previously mentioned guidance, has an audio emitter and audio listener – in other words, the player. The audio emitter is projecting a constant sound into the scene, and - in the raytraced version of the testing framework – the audio listener is projecting out multiple rays each frame to seek the audio emitter. The goal of the user when testing each audio system is find and collect the audio emitter ten times inside the testing framework. To avoid users from just walking around the scene and using their sight to see the yellow audio emitter and collect it, audio emitters are invisible until directly approached by the user. This would mean that they must use their hearing abilities rather than walking around the level.

To view the information sheet that was given to the user participants to educate them on the topic of the project and what they would be testing, refer to [Appendix \[14.2\]](#).

When an audio emitter is collected by the user, it moves to somewhere else in the scene. This position is on a set list and is the same order for each user to avoid any random placement in the variables of the testing. When the testing framework is created, it randomly chooses from an audio system for the user to test to try to reduce any bias that the primary researcher may introduce, as guided by Taherdoost (2022).

Figure 4 showcases the testing environment below:

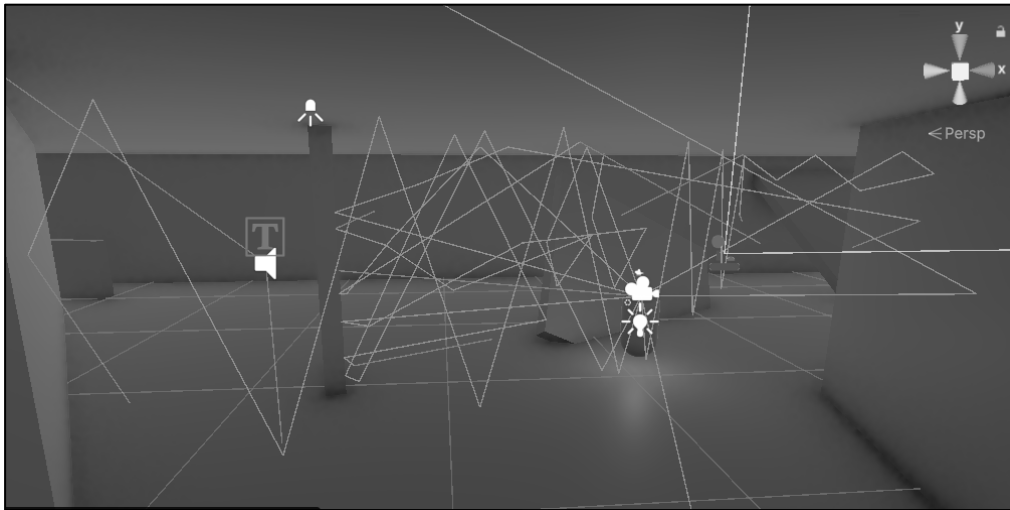


Figure 4: A screenshot of the raytraced testing environment, showcasing rays bouncing around the constructed testing level to try and find the audio emitter. The audio emitter is displayed inside the screenshot with the **T** symbol, whilst the player is shown with the **camera & light** symbol. An important note here is to mention that the number of rays has been reduced from 1200 per second - Terletskeyy & Trushevskyy's recommendation - to 20 per second for visual demonstration purposes.

Before the tests were conducted, users were also instructed to sign a consent form to signify their consent to all expectations and outcomes.

To view the consent sheet that was given to the user participants for them to sign to be able to continue and complete the user testing, refer to [Appendix \[14.1\]](#).

A pre-test was also conducted on users before the main testing to gather additional data, like hearing range and to meet with ethical criteria, such as adjusting volume for comfort levels. This is explained in the [next section](#) when describing the testing framework pre-test further.

Metrics on the users' performance with the different sound systems were written to a log file upon each test completion and then collected for later analysis.

To view the layout and the average user results for a participant – in this case, user number one – refer to [Appendix \[14.3\]](#).

An important note to make here is that, originally, users who would test one audio system would test the next audio system after their first test. The results from each one of these secondary tests were almost always biased to the second sound system, which is clear learnership bias. Because of this, secondary test results from the same user were discarded, and the decision was made to split the users who tested the ray-traced sound system first into one group and the users who tested the spatial audio sound system into another group.

8.3 MATERIALS AND EQUIPMENT

For testing the users, computers at the University of Staffordshire were the platform used; these computers contain the following specifications:

- RTX 3080 GPU
- I7 11700k CPU (8 Cores 16 Threads)
- 32GB DDR4 RAM

- 2TB SSD
- Windows 11 OS

These systems were chosen due to convenient access, as they are situated on the university campus in areas that the testing demographic already frequents, making testing a much faster process. Additionally, the specifications of these systems are by far advanced enough to run the testing framework without issues such as drops in framerate, crashes or other errors; nevertheless the testing framework can run on much lower specifications, which prevents hardware-based issues before they can become a problem. Keeping the system consistent throughout testing also eliminates hardware-based anomalous results.




Figure 5: University of Staffordshire (2026) standard video games development computers that participants were tested on. Featuring a mouse, keyboard and dual-screen monitors. An important note here is that only one monitor was used for the duration of the user testing.

In each testing session, the participants were given a pair of Corsair Virtuoso SE headphones (2019). These headphones have the following specifications as per Corsair (2019):

- Frequency response ~ 20hz – 40,000hz
- Sensitivity ~ 109dB (+/- 3dB)
- Audio ~ 7.1 Surround

While keeping PC specifications prevents issues with the testing framework, keeping the audio device consistent among all users ensures the testing groups are hearing as close to the same audio quality as reasonably possible. Using different headphones can cause issues such as differences in frequency ranges, causing different audio output, differing audio qualities, obscuring sounds that other participants would have heard more clearly (affecting average response time and spatial reasoning), as well as the possibility of hardware faults in the different headsets that could cause various anomalous results.



Figure 6: Corsair Virtuoso SE (2019) headphones that participants used for the duration of the testing. Each participant used the same headphones to help mitigate any hardware differences that using different headphones may cause.

Before testing, the participants were required to complete a pre-test to gather relevant information and calibrate the system's volume. The participants were required to confirm they had completed the relevant consent forms as required by the project's ethics regulation guidelines. These forms were completed in person and by hand before testing was initiated. Next, the participants were required to calibrate the volume of the test to be comfortable for them. Loud volume can damage hearing in the long term, so comfortable hearing levels per person, thus they were asked to calibrate this themselves for optimal safety, comfort and performance.

A key step for the interpretation of the resulting data was the hearing frequency range test, which consisted of playing increasing audio frequencies that the participant was required to pause as soon as they could no longer hear the audio. While due to human reaction time and processing, this number is not exact, it is the closest estimation available and gives a relatively accurate response in terms of the data for the needs of this test. Lastly, the same screen requires a response to the frequency test in Hz, as well as the given tester number, so their data can be anonymously catalogued.



Figure 7: A screenshot of the main menu of the developed user testing framework, which users would start with. They would start the pre-testing process by clicking on

the ‘Pre-Test’ button, which would take them to the following screens, refer to [Appendix \[14.4\]](#).

9 RESULTS AND FINDINGS

9.1 PARTICIPANT HEARING RANGE FREQUENCY

Participant hearing ranges were measured to provide a known metric when used in the analysis of the testing results. A hearing range frequency value was also collected to mitigate any significant user sampling concerns that could occur, with the comparison of hearing ranges between participants being the only user-specific metric necessary when testing audio systems. Because of this age, gender and other known metrics were not considered to be relevant.




Figure 8: A box plot graph of the user hearing ranges (Hz) between the two groups, showing the median, the quartiles and any outliers. It is important to note that group “Spatial Audio” had eleven participants, whilst group “Raytraced Audio” had only nine.

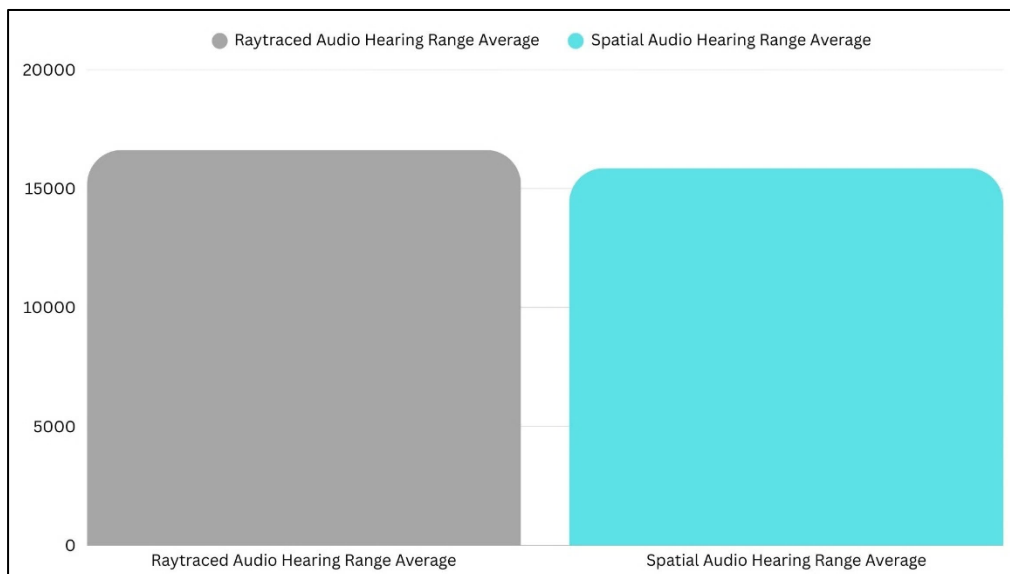


Figure 9: A bar chart graph on the average hearing range difference between the two testing groups. The raytraced audio testing group had an average hearing range of 16,621 Hz, whilst the spatial audio testing group had an average hearing range of 15,855 Hz. Measured in hertz (Hz).

9.2 SPATIAL AUDIO TESTS

Three main testing metrics were gathered from both audio system testing groups when testing the user framework. These are as follows:

- **Total Duration:** for the completion of the test.
- **Duration per emitter average:** for the average duration it took for users to collect each audio emitter.
- **Heading alignment accuracy:** to determine if users were aligned with the audio emitters within a 45°-degree angle.

The results of the spatial audio testing groups are seen below.

For any spatial audio user testing raw data, refer to [Appendix \[14.5\]](#).

The following are some interpretations of the spatial audio data condensed into graph form, featuring bar charts and scatter graphs:

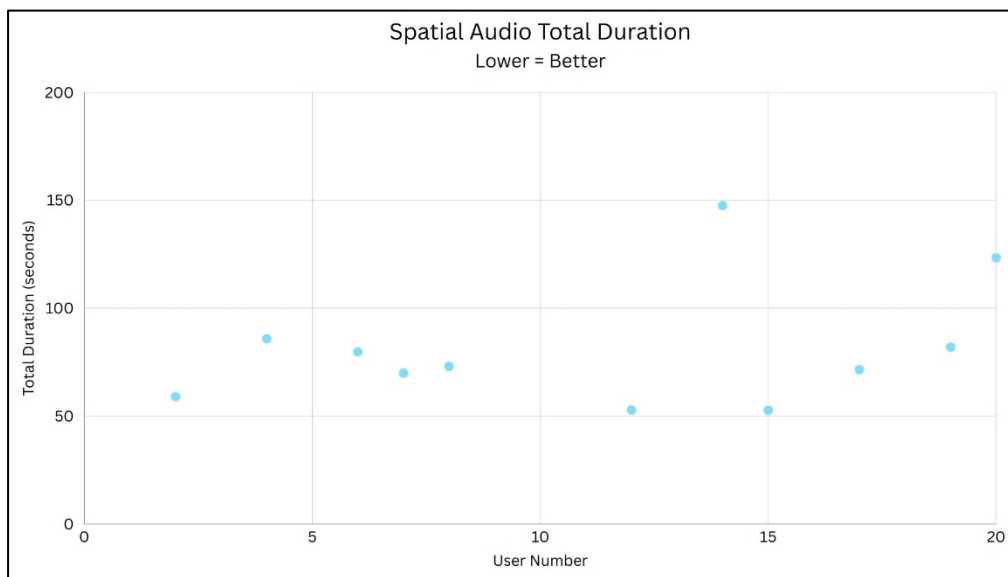


Figure 10: The spatial audio ‘total duration’ results - interpreted in scatter graph format - from users two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in seconds (s).

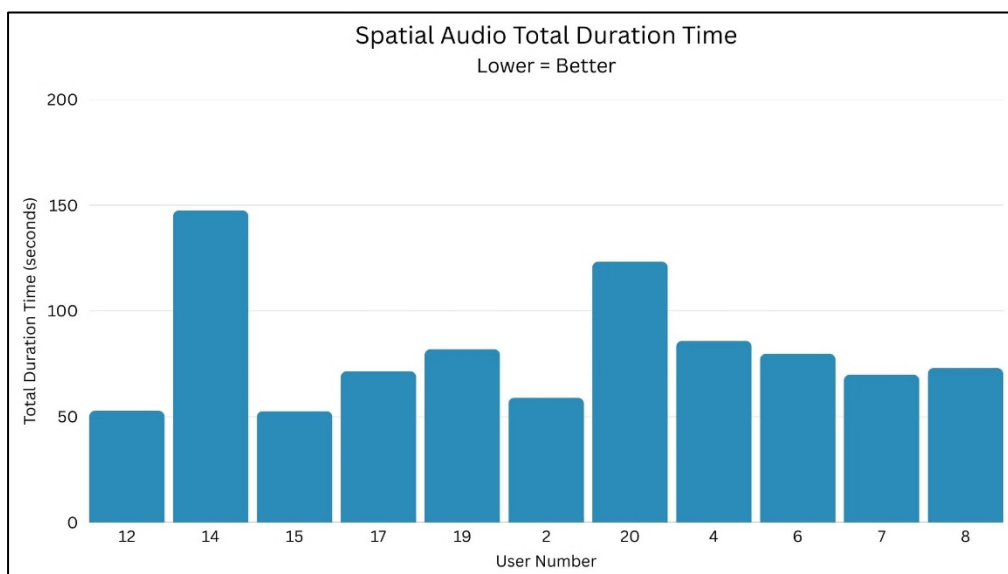


Figure 11: The spatial audio 'total duration' results - interpreted in bar graph format - from users two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in seconds (s).




Figure 12: The spatial audio 'duration per emitter average' results - interpreted in scatter graph format - from users, two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in seconds (s).

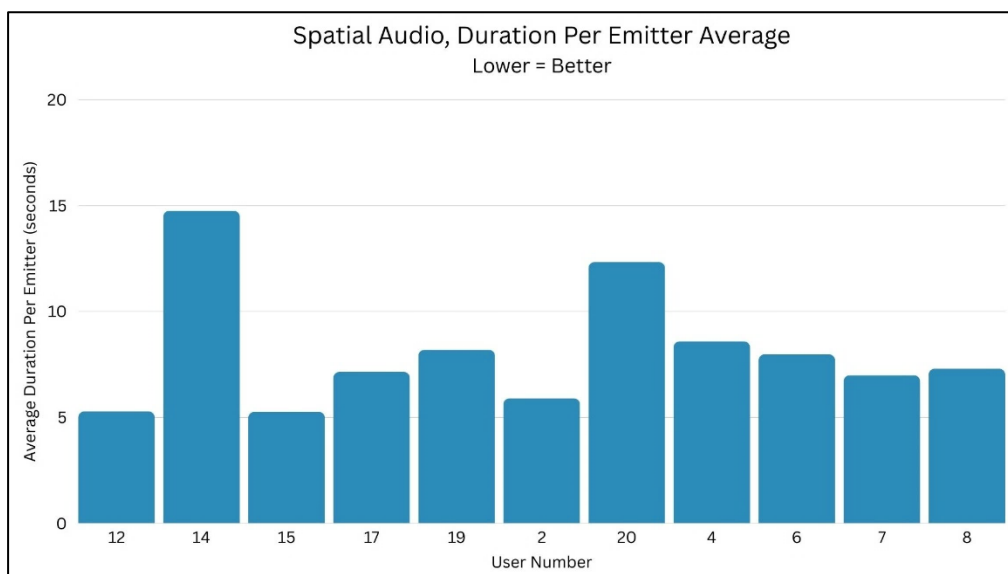


Figure 13: The spatial audio 'duration per emitter average' results - interpreted in bar graph format - from users, two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in seconds (s).

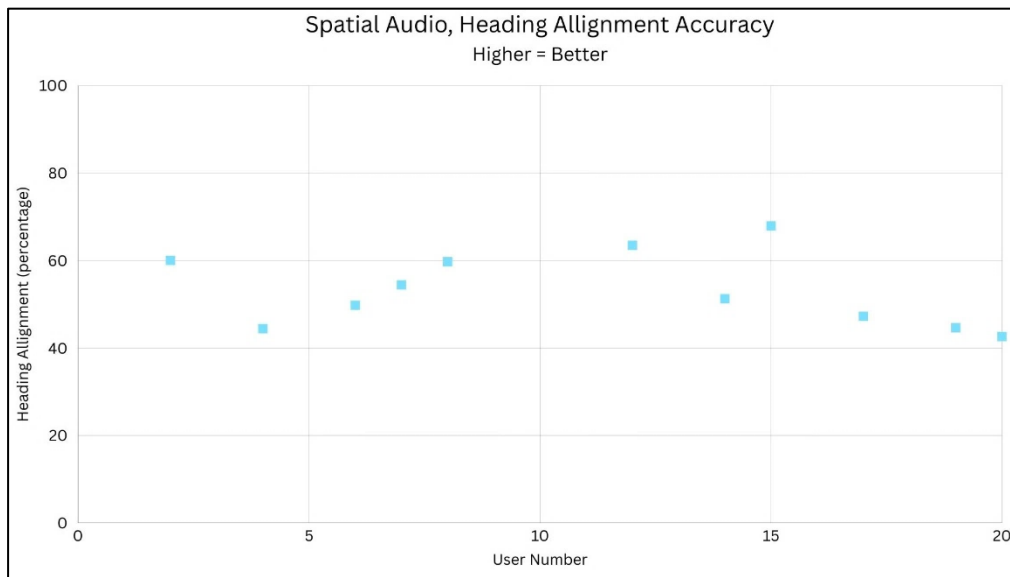


Figure 14: The spatial audio 'heading alignment accuracy' results - interpreted in scatter graph format - from users two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in percentage alignment with the audio emitter out of one hundred (%), within a 45°-degree angle.

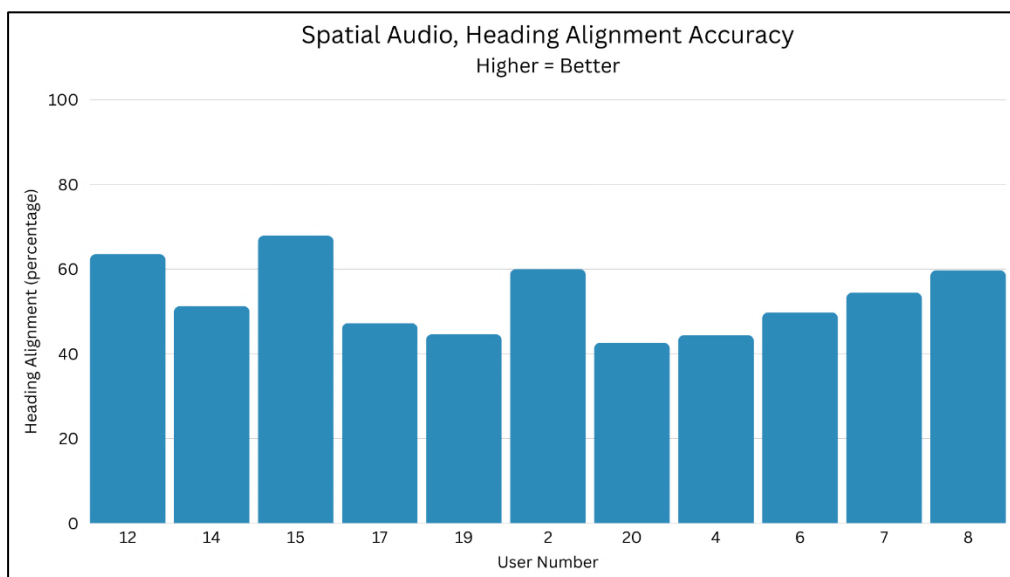


Figure 15: The spatial audio 'heading alignment accuracy' results - interpreted in bar graph format - from users two, four, six, seven, eight, twelve, fourteen, fifteen, seventeen, nineteen and twenty. Measured in percentage alignment with the audio emitter out of one hundred (%), within a 45°-degree angle.

9.3 RAYTRACED AUDIO TESTS

The same metrics used for the spatial audio testing group were also gathered and compiled from the raytraced audio testing group; they can also be seen below, interpreted as a bar and scatter graph format. No other differences apart from the audio system were changed between testing the two groups.

For any raytraced audio user testing raw data, refer to [Appendix \[14.6\]](#).

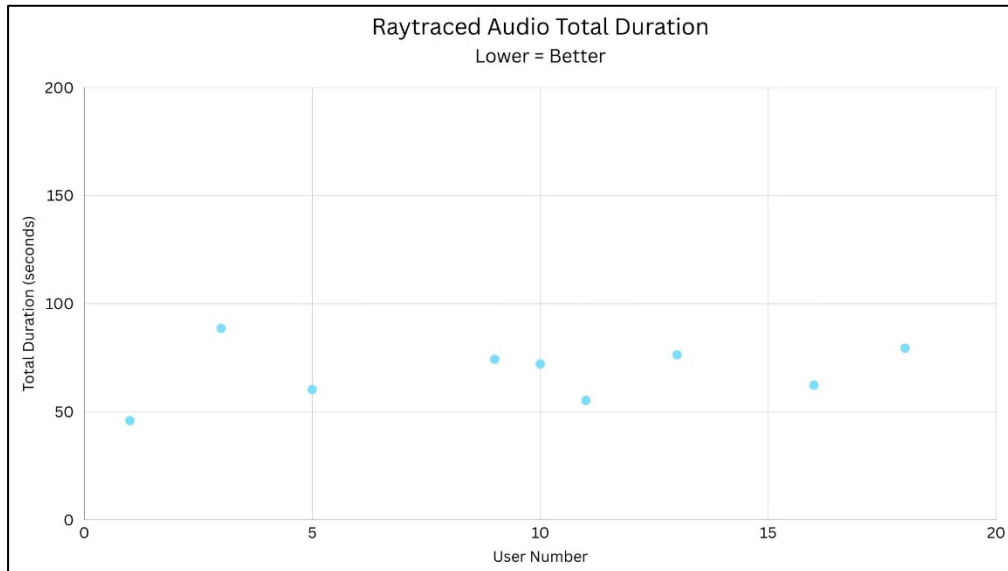


Figure 16: The raytraced audio ‘total duration’ results - interpreted in scatter graph format - from users one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in seconds (s).

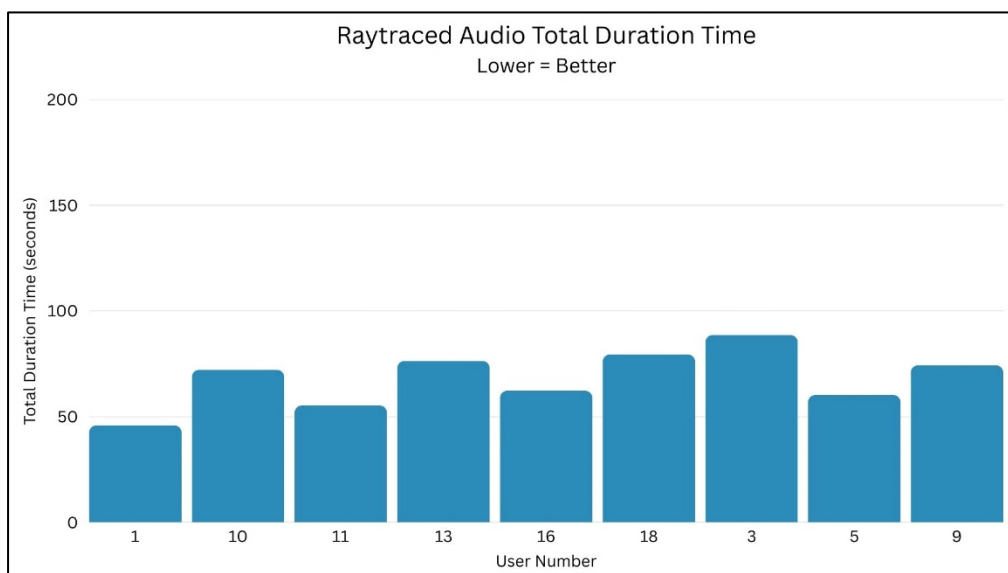


Figure 17: The raytraced audio ‘total duration’ results - interpreted in bar graph format - from users one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in seconds (s).

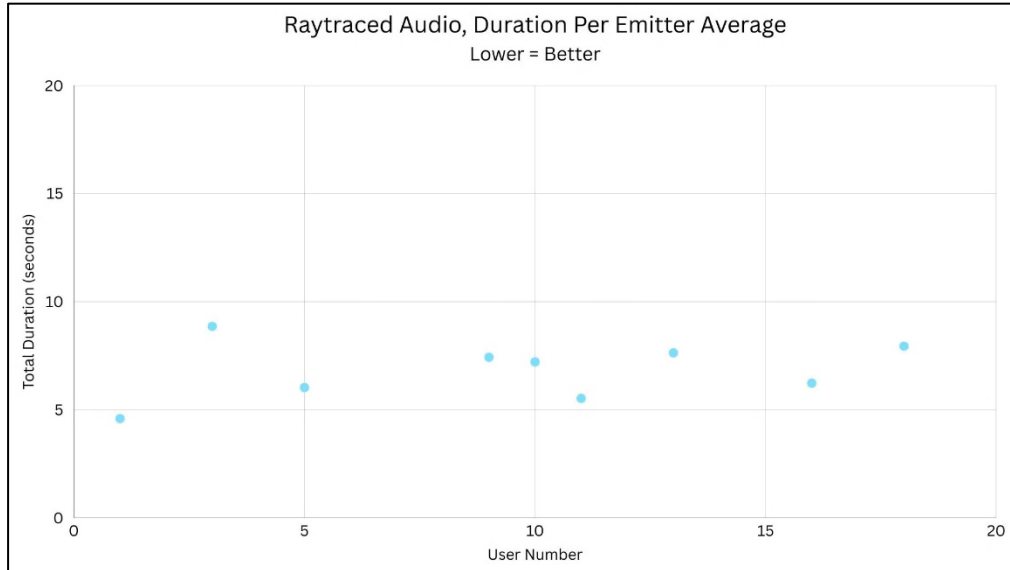


Figure 18: The raytraced audio ‘duration per emitter average’ results - interpreted in scatter graph format - from users one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in seconds (s).




Figure 19: The raytraced audio ‘duration per emitter average’ results - interpreted in bar graph format - from users, one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in seconds (s).

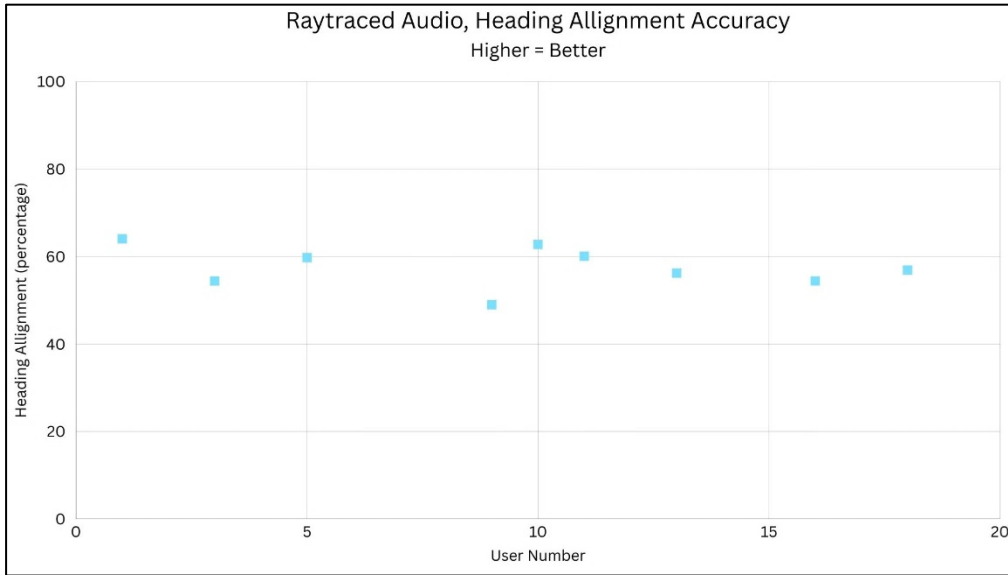


Figure 20: The raytraced audio ‘heading alignment accuracy’ results - interpreted in scatter graph format - from users one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in percentage alignment with the audio emitter out of one hundred (%), within a 45°-degree angle.

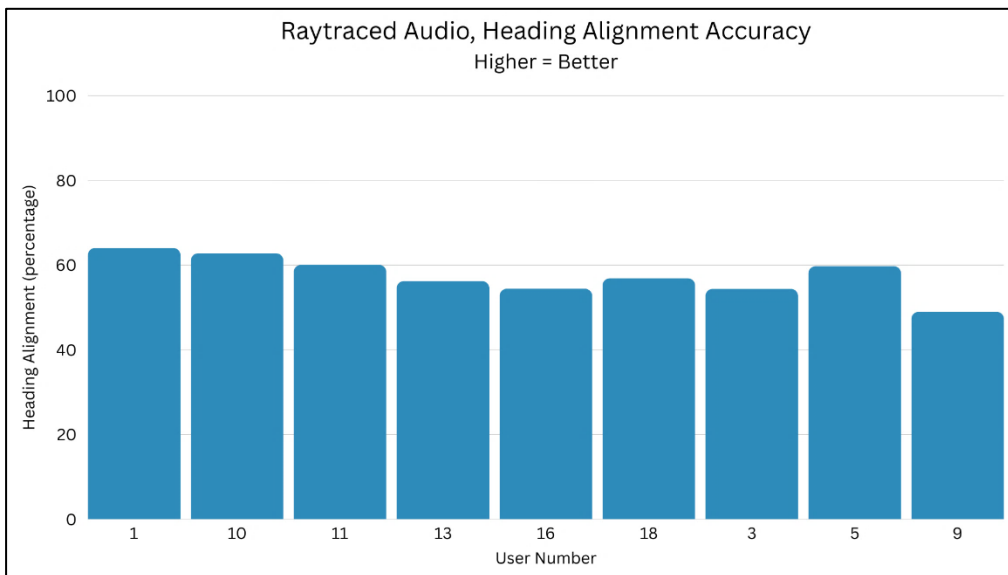


Figure 21: The raytraced audio ‘heading alignment accuracy’ results - interpreted in bar graph format - from users, one, three, five, nine, ten, eleven, thirteen, sixteen and eighteen. Measured in percentage alignment with the audio emitter out of one hundred (%), within a 45°-degree angle.

10 DISCUSSION AND ANALYSIS

10.1 ANOMALOUS RESULTS

On inspection of the total collected results of the twenty user participants, it was identified that two out of the twenty results were classed as anomalous and were far out of the bounds of acceptability that the other participants scored within. To keep the tests rational, these results were removed from the spatial audio testing group to stop the group from averaging higher than they realistically should be. For clarity, the user numbers that were removed were users numbered fourteen and twenty, who scored significantly higher than their fellow participants.

The average total duration time to complete the spatial audio test was 69.60 seconds, whilst these users scored 147.53 seconds and 123.30 seconds respectively. These results will still be present in any further graphs on the matter, but will be removed from the final conclusive results, hearing range, total duration, per emitter duration average and heading alignment accuracy. The removal of these participants' data from the study also balances the uneven weighting of the testing groups from eleven spatial audio participants to nine; this now puts the current weighting of groups at nine users per group.

10.2 PARTICIPANT HEARING RANGE ANALYSIS

The total participants' hearing range average among the group of eighteen (see previous section where two users' anomalous results were removed) was 16158.5Hz, which is in the average range for human hearing (Operating Eurovision and Euroradio, 2023). With this number in mind, it is safe to assume that the rest of the user testing group did not have any anomalous hearing range results that might sway any further testing results. Breaking this down further, the spatial audio user testing group had a hearing range average of 15695.56Hz, whilst the raytraced audio group had a hearing range average of 16621.44Hz. This meant that the ray-traced audio group had a slight hearing range increase of 5.89%, which will be important to consider when concluding the following data.

10.3 ANALYSIS SPATIAL AUDIO & RAYTRACED AUDIO

This next section will cross-examine the results gathered from the eighteen remaining user data points and plot all remaining data on scatter graphs seen below. This way, similar data can be presented at the same time for easier comparison. Participants who tested the raytraced audio system are marked with a red colour, whilst participants who tested the spatial audio system are marked with a blue colour. As previously mentioned, each audio system test is the same as each other, same starting location, same positional order of audio emitters and the same level design. The only difference between the tests is the audio system which the two groups interacted with.

For finalised spatial user testing data, with averages and user numbers fourteen and twenty removed, refer to [Appendix \[14.7\]](#).

For finalised raytraced user testing data, with averages, refer to [Appendix \[14.8\]](#).

The graph shown in **Figure 22** displays the total duration time of the testing framework for both sets of groups. This graph shows a neutral correlation between the data points, excluding the outlier results. The average total duration time for both groups was 68.95 seconds, with the average time for the spatial audio group being 69.60 seconds, whilst the average time for the raytraced audio group was 68.30

seconds. This shows a 1.86% decrease in the total duration timing of the tests when using a ray-traced audio solution for users over a spatial audio solution. This is a minor improvement over the spatial audio system, but it is within the margin of error, and when accounting for the raytraced groups' improvement in average hearing range does not represent a statistically significant result.




Figure 22: The combined results of the raytraced/spatial users 'total duration' values recorded from the testing framework. Users who tested the raytraced audio system are marked in red, whilst users who tested the spatial audio system are marked in blue. **User numbers fourteen and twenty, coloured in purple here to showcase anomalous results, have been disregarded from the final average for this metric.**

The graph shown in **Figure 23** shows the duration per emitter average for the two groups and shows another neutral correlation between them. The average total for the duration per emitter time for both groups was 6.89 seconds, with the average time for the spatial audio group being 6.96 seconds, whilst the average time for the raytraced audio group was 6.83 seconds. Like the graph shown in **Figure 22**, another decrease of 1.86% is seen when users tested a ray-traced audio solution over a spatial audio solution. This can be yet another case within the margin of error or justified by the ray-traced groups' improved hearing range over the spatial audio group.

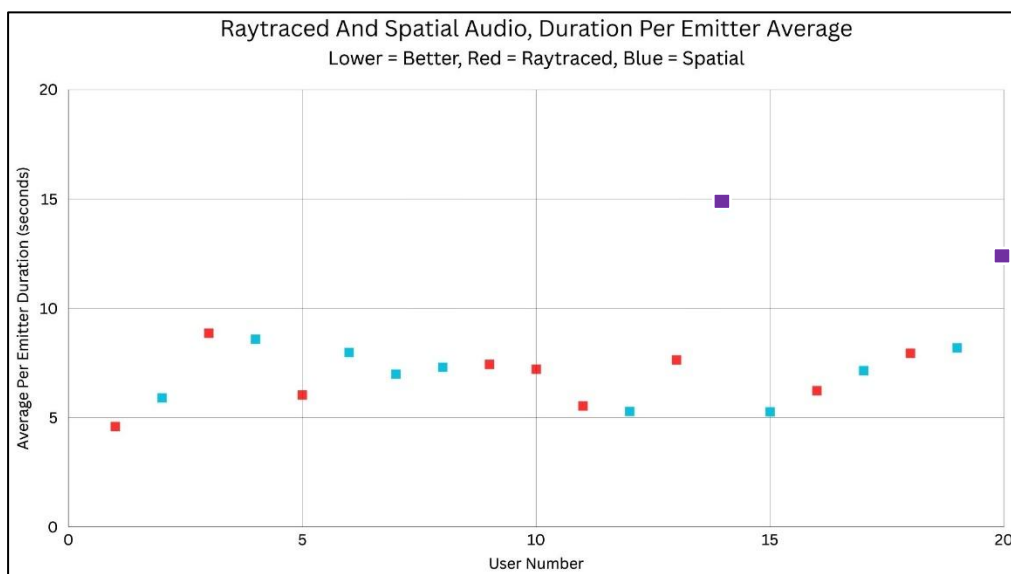


Figure 23: The combined results of the raytraced/spatial users 'duration per emitter average values recorded from the testing framework. Users who tested the raytraced audio system are marked in red, whilst users who tested the spatial audio system are marked in blue. **User numbers fourteen and twenty, coloured in purple here to showcase anomalous results, have been disregarded from the final average for this metric.**

The graph shown in **Figure 24** showcases the combined heading alignment results of both groups testing the raytraced/spatial sound system framework. This graph shows another neutral correlation with both groups, with the combined group average heading alignment being 56.07%. Looking at this result further, the spatial audio groups' average alignment heading was 54.63%, whereas the raytraced audio groups' average was 57.51%, which is a unique result compared to the other testing metrics gathered between the two groups. A 2.88% increase is seen when users tested with the raytraced audio solution over the spatial audio solution; this result is outside the margin of error (+- 2%) and could be an indication that raytraced audio provided better user spatial awareness. It is still important to remember that the ray-traced audio group did have a better hearing range, which could be a slight factor in this improvement.




Figure 24: The combined results of the raytraced/spatial users 'heading alignment accuracy values recorded from the testing framework. Users who tested the raytraced audio system are marked in red, whilst users who tested the spatial audio system are marked in blue. **Anomalous results have not been marked here as they follow the correct trend shown in the graph but have been disregarded in final averages for this metric as well.**

10.4 CRITICISMS OF THE INVESTIGATION

The metrics used to measure the participants' performance during testing were somewhat inadequate; the use of the individual's alignment heading relative to the audio source was a misinformed source of information. This is due to general behaviour in virtual spaces; most users will focus the camera on navigating their surroundings, such as pathing around obstacles and panning the camera to understand their surroundings. Alignment heading would only have been an accurate metric in an environment that does not require navigating obstacles, instead moving in a straight line from point A to point B.

Implementation of the framework itself could also have been heavily improved, for instance, due to strict testing deadlines, audio calculations had to be defaulted to Unity's built-in systems, rather than tailored for the ray tracing system. The same can also be said for the exclusion of Csound, which would have greatly increased audio quality and realism to assist in providing adequate spatial reasoning for the participant.

11 CONCLUSION

11.1 AIMS AND OBJECTIVES CHECKLIST

To assess the adequacy of this project, the initial aims and objectives will be questioned to provide a reasonable checklist:

Aim: A Realtime Sound Simulation Artefact supporting the Research Question Proposed.

Result: A real-time ray-traced/spatial audio testing framework was developed.

Aim: A collection of user-related information on the improvement/changes aim to spatial reasoning under different audio technologies, such as acoustic raytracing, compared to simpler distance-based audio techniques.

Result: User information was gathered.

Aim: A literature review on the topic of audio-based techniques utilised in video games and how raytraced audio can be used for the benefit of improving spatial perception.

Result: A literature review was conducted.

Aim: A series of tests performed in the Realtime Sound Simulation Artefact to gather user-related information.

Result: Tests were performed to collect this information.

Aim: A discussion on the results and findings of user testing regarding whether spatial perception was improved by one audio technique more than another.

Result: A discussion on the results of the study was conducted, and final audio system conclusions were discussed in the [next section](#).

Aim: A conclusion on the analysis to answer the proposed research question.

Result: Conclusion has taken place.

Aim: A series of recommendations to anyone looking to further the study of ray-traced audio technologies in a general and gaming environment.

Result: Recommendations for future study have been provided in the [final section](#) of this document.

11.2 SPATIAL REASONING AUDIO SYSTEM

The results of this study are inclusive, but some results lean towards a ray-traced audio system providing users with very minor improvements to spatial awareness over more traditional audio systems, such as spatial audio. There is no full indication of improvement due to the general margins of error set in place for adjusting hearing range differences. Even if both margins of error and hearing range differences between the two testing groups were discarded, the general implementation of the ray-traced audio testing framework could misalign with other developments of the same system.

For more conclusive results, a much wider user sample size would need to be tested against with many varieties of ray-traced audio implementations.

12 RECOMMENDATIONS

12.1 GAPS IN KNOWLEDGE

As mentioned previously, there is a general lack of step-by-step guides for the implementation of acoustic ray tracing within real-time virtual spaces. While methods used in other industries, such as architecture, are well documented, they are not as relevant in real time contexts due to their vastly different approach and drawbacks, such as preprocessing. This being said, some research into the field is documented and the systems are described/tested by Terletsky and Trushevsky (2024), and Samsurya (2023) however implementation itself is not discussed to highly replicable standard, with only certain techniques discussed used to enhance ray tracing audio, this creates the need to develop largely off described principles from scratch rather than follow a proven robust method.

Many of the systems described in these papers also use the Csound library, which can be used to considerably enhance audio quality and realism through its vast array of features. Some of this implementation is described by Terletsky & Trushevsky (2024); however, Csound is an expansive system that could not be implemented due to the constraint of time for the project. However, given further development, this library would be instrumental in improving the testing framework through enhanced audio control and additional features such as reverberations and delay.

12.2 SAMPLING PARTICIPANTS

The sample group for testing consisted of 20 participants, including two removed due to anomalous results because of external factors in testing. Of the included, nine represented spatial audio testing and nine represented testing raytraced audio. This sample size is comparatively rather small, although sufficient; more data could have revealed more accurate correlations and data, allowing the conclusions drawn to be more heavily supported and concise.

Targeting a specific age demographic would also streamline results. Hearing frequency range is generally tied to age, thus by enlisting a small age range (19 to 21) (ITU-R, 2015) The number of anomalous results can be minimised. Anomalous results are already somewhat minimised during current testing, as individuals with pre-existing conditions directly related to difficulty hearing were not invited into the testing sample as per the ethics declaration. A relatively similar amount of experience in playing video games should also be screened before testing in the future, as well, to prevent results reflecting an individual's ability to navigate a virtual space rather than their ability to reason audio within it.

12.3 RAYTRACED AUDIO IMPLEMENTATION

The raytraced audio implementation seen throughout this study – and the same one that users were tested against – could have benefited from many various improvements as discussed in the [literature review section](#) of this paper. The following improvements can be made to the ray-traced audio system used are the following:

- Using audio middleware such as Csound or Wwise for enhanced audio processing.
- Implementations of environmental reverberation and echo.
- Raytraced material interaction against object materials in the scene.

13 REFERENCES

Allen, J. and Berkley, D. (1978). *Image method for efficiently simulating small-room acoustics*. Murray Hill, New Jersey: Acoustics Research Department, Bell Laboratories.

Beauchemin, S. (2023). *Audio Debugging Tools and Techniques*. In *Game Audio Programming 4*. 4th edn. CRC Press.

Cao, Z., Magalhães, E. and Bernardes, G. (2023). *Sound Design Impacts User Experience and Attention in Serious Games*. Springer Science and Business Media Deutschland GmbH.

Corsair (2019) *VIRTUOSO RGB WIRELESS SE High-Fidelity Gaming Headset*. Available at: <https://www.corsair.com/uk/en/p/gaming-headsets/ca-9011180-eu/virtuoso-rgb-wireless-se-high-fidelity-gaming-headset-gunmetal-eu-ca-9011180-eu>.

Davari, N.V.C. (2021). *Efficient 2D sound propagation in video games*.

Epic Games (2026) *Audio in Unreal Engine*. Available at: <https://dev.epicgames.com/documentation/en-us/unreal-engine/audio-in-unreal-engine>.

Epic Games (2026) *Using a Single Line Trace (Raycast) by Channel*. Available at: <https://dev.epicgames.com/documentation/en-us/unreal-engine/using-a-single-line-trace-raycast-by-channel-in-unreal-engine>.

Geier, M., Ahrens, J. and Spors, S. (2010). *Object-based audio reproduction and the audio scene description format*.

Geier, M., Ahrens, J. and Spors, S. (2012). *The SoundScape Renderer: A unified spatial audio reproduction framework for arbitrary rendering methods*. Available at: www.aes.org.

Gimenez, M., Vidaurre, A., Riera, J. and Monsoriu, J. (2007). *Visualising the Doppler Effect*.

Guy Sombeg (2024) *Game Audio Programming 4*.

Hacihabiboglu, H. D.S.E., C.Z., J.J.a.S. (2017). *Perceptual spatial audio recording, simulation, and rendering: An overview of spatial-audio techniques based on psychoacoustics*.

Hooiveld, V. (2025). *Understanding the Game Audio Pipeline*. Available at: <https://www.audiobyvasco.com/blog/understanding-the-game-audio-pipeline>.

ITU Radiocommunication Assembly (2001) *Method for the subjective assessment of intermediate quality level of coding systems*.

ITU-R (2015) *Methods for the subjective assessment of small impairments in audio systems BS Series Broadcasting service (sound)*. Available at: <http://www.itu.int/ITU-R/go/patents/en>.

jungle_Pinpoint (2025) *Pinpoint - Audio Tracing*. Available at: <https://forums.unrealengine.com/t/pinpoint-audio-tracing/2645588>.

Krokstad, A., Str, S. and Sbrsdal, S. (1968) *CALCULATING THE ACOUSTICAL ROOM RESPONSE BY THE USE OF A RAY TRACING TECHNIQUE*.

- Lindskiv, J. and Riddersholm, M. (2019) *Game Audio Pipelines*. Bachelor's project. Sonic College - UCSYD Haderslev.
- Lixandru, A.B.G.S.B.A. (2021). *Acoustic Modelling for Indoor Spaces Using the Ray-Tracing Method*. Springer International Publishing.
- Maekawa, Z. (1968) *Noise Reduction by Screens, Applied Acoustics*.
- Mueller, W.U.F. (1999) *A scalable system for 3D audio ray tracing*. IEEE.
- Nair, V. et al. (2022) *Uncovering Visually Impaired Gamers' Preferences for Spatial Awareness Tools Within Video Games*. Association for Computing Machinery, Inc.
- NVIDIA (2018). *Introduction to NVIDIA RTX and DirectX Ray Tracing*. Available at: <https://developer.nvidia.com/blog/introduction-nvidia-rtx-directx-ray-tracing/>.
- NVIDIA (2018) *Ray Tracing*. Available at: <https://developer.nvidia.com/discover/ray-tracing>.
- Operating Eurovision and Euroradio (2023) *R 128 LOUDNESS NORMALISATION AND PERMITTED MAXIMUM LEVEL OF AUDIO SIGNALS Status: EBU Recommendation*.
- Paterson, N. et al. (2010) *Spatial Audio and Reverberation in Augmented Reality Game Sound Design*. Dublin: Trinity College.
- Rickard, F. (2020). *Engaging gameplay for audio games*. PhD Thesis. Norrköping, Sweden: Linköping University.
- Röber, N., Andres, S. and Masuch, M. (2006). *HRTF simulations through acoustic raytracing*. Halle (Saale), Germany: Universitäts- und Landesbibliothek Sachsen-Anhalt.
- Samsurya, B. (2023) *Realtime Audio Raytracing and Occlusion in Csound and Unity*.
- Savioja, L. and Svensson, U.P. (2015). *Overview of geometrical room acoustic modelling techniques*. Acoustical Society of America (ASA).
- Schissler, C.M.D. (2011). *Gsound: Interactive sound propagation for games*. Audio Engineering Society.
- Taherdoost, H. (2022). What are the different research approaches? Comprehensive review of qualitative, quantitative and mixed method research, their applications, types and limitations. *Journal of Management Science & Engineering Research* pp.55-63.
- Taylor, M. et al. (n.d.) *Guided Multiview Ray Tracing for Fast Auralization*. Available at: <http://gamma.cs.unc.edu/Sound/Guided/>.
- Terletskyy, O. and Trushevskyy, V. (2024) *RAY TRACING BASED APPROACH FOR THE SOUND PROPAGATION SIMULATION*.
- Tsuchiya, T. and Kagawa, Y. (1991) *A Basic Study on Nonlinear Sound Propagation by Finite Element Simulation*.
- Unity (2026) *Physics.Raycast*. Available at: <https://docs.unity3d.com/6000.3/Documentation/ScriptReference/Physics.Raycast.html>.

Unity (2026). *Introduction to the Audio Source component*. Available at: <https://docs.unity3d.com/6000.3/Documentation/Manual/AudioSource-overview.html>.

Venkatesan, H., Durand, E. and Bouillot, N. (2021) *LIVE RAY TRACING AND AURALIZATION OF 3D AUDIO SCENES WITH VARAYSÉmile*.

Weinzierl, S., Lepa, S. and Ackermann, D. (2018). *A measuring instrument for the auditory perception of rooms: The Room Acoustical Quality Inventory (RAQI)*. Acoustical Society of America (ASA).


Zhang, M., Zhang, W., Kennedy, R.A. and Abhayapala, T.D. (2009) 'HRTF measurement on KEMAR manikin.' Canberra

14 APPENDICES

14.1 USER CONSENT SHEET

The user consent sheet was signed by the testing participants for the purpose of the study to answer the research question. Users would sign this after reading the information on the information sheet, then they would play the testing application.

Page 1/1:



RESEARCH PROJECT CONSENT FORM

Title of Project: Can Realtime Raytraced Audio in Virtual Environments Improve Spatial Reasoning?

Researcher: Ryan Moakes

I have read and understood the information sheet. Yes No

I have been given the opportunity to ask questions, and I have had any questions answered satisfactorily. Yes No

I understand that my participation in this study is entirely voluntary and that I can withdraw at any time without having to give an explanation. Yes No

I am unaware of any pre-existing hearing conditions that may exclude me from participating in the test. Yes No

I consent to play and test a sound simulation application. Yes No

I consent to wear headphones for the duration of the application playtest. Yes No

I understand that the playtest will be screen captured. Yes No

I understand that the playtesting application records data such as mouse movement, general test timings and other data that can be used for research purposes. Yes No

I consent that data collected could be analysed and used to inform the project and understand that all data will be presented anonymously. Yes No

I agree that data will only be used for this project although the data may also be audited for quality control purposes. Yes No

All data will be stored safely on a password protected computer (electronic data), or locked away securely (hard copies of data) until the project finishes at which point it will be destroyed. Yes No


I understand that I can withdraw my data from the project up to the point where the data has been analysed without having to give an explanation. Yes No

I hereby give consent to take part in this study. Yes No

14.2 USER INFORMATION SHEET

This is the user information sheet, given to the users before signing the consent sheet, to gain the appropriate information about the study before committing to completing the testing framework.

Page 1/3:



V1.1 15/10/25

INFORMATION SHEET FOR PARTICIPANTS

Project Reference Number: 23015485

Title of study
Can Realtime Raytraced Audio in Virtual Environments Improve Spatial Reasoning?

Invitation Paragraph

I would like to invite you to participate in this research project which forms part of my *Level 6 Game Project*. Before you decide whether you want to take part, it is important for you to understand why the research is being done and what your participation will involve. Please take time to read the following information carefully and discuss it with others if you wish. Ask me if there is anything that is not clear or if you would like more information.

What is the purpose of the study?

I am collecting data on this project to determine if taking a new approach to audio simulation techniques such as "Raytraced Audio" can be beneficial to the spatial reasoning of users who play games over older more in-accurate audio technologies seen previously in games. This will then determine that "Raytraced Audio" has more benefits than just more accurate data gathering.

Why have I been invited to take part?

You have been invited to take part as your insights and feedback would be valuable to wards the goal of the project.

What will happen if I take part?

You will be asked to play and participate in a few application-based tests to determine the difference between audio technologies. You will be asked to wear headphones and try to accomplish certain audio/sound activities based in a virtual environment. The application and screen will record your input and gather information useful to the study. After the application-based test has been concluded you may be asked to fill in some qualitative based questions on the experience.

Do I have to take part?

Participation is completely voluntary. You should only take part if you want to and choosing not to take part will not disadvantage you in anyway. Once you have read the information sheet, please contact us if you have any questions that will help you make a decision about taking part. If you decide to take part we will ask you to sign a consent form and you will be given a copy of this consent form to keep.

1



V1.1 15/10/25

What are the possible risks of taking part?

There should be no possible risks of taking part.

What are the possible benefits of taking part?

There are no tangible benefits from taking part.

Data handling and confidentiality

Your data will be processed in accordance with the data protection law and will comply with the General Data Protection Regulation 2016 (GDPR). All data will be stored on a university owned OneDrive account that is password protected. Additionally the data will be destroyed after the project has finished.

Data Protection Statement

The data controller for this project will be Staffordshire University. The University will process your personal data for the purpose of the research outlined above. The legal basis for processing your personal data for research purposes under the data protection law is a 'task in the public interest' You can provide your consent for the use of your personal data in this study by completing the consent form that has been provided to you.

You have the right to access information held about you. Your right of access can be exercised in accordance with the General Data Protection Regulation. You also have other rights including rights of correction, erasure, objection, and data portability. Questions, comments and requests about your personal data can also be sent to the Staffordshire University Data Protection Officer. If you wish to lodge a complaint with the Information Commissioner's Office, please visit www.ico.org.uk.

What if I change my mind about taking part?

You are free to withdraw at any point of the study, without having to give a reason. Withdrawing from the study will not affect you in any way. You can withdraw your data from the study up until the data has been analysed after which withdrawal of your data will no longer be possible due to it being anonymised and decoupled from you after the analysis has taken part.


If you choose to withdraw from the study we will not retain any information that you have provided us as a part of this study.

What will happen to the results of the study?

The data will be analysed and then used to inform the project.

Who should I contact for further information?

If you have any questions or require more information about this study, please contact m015485n@student.staffs.ac.uk

Page 3/3:

V1.1 15/10/25

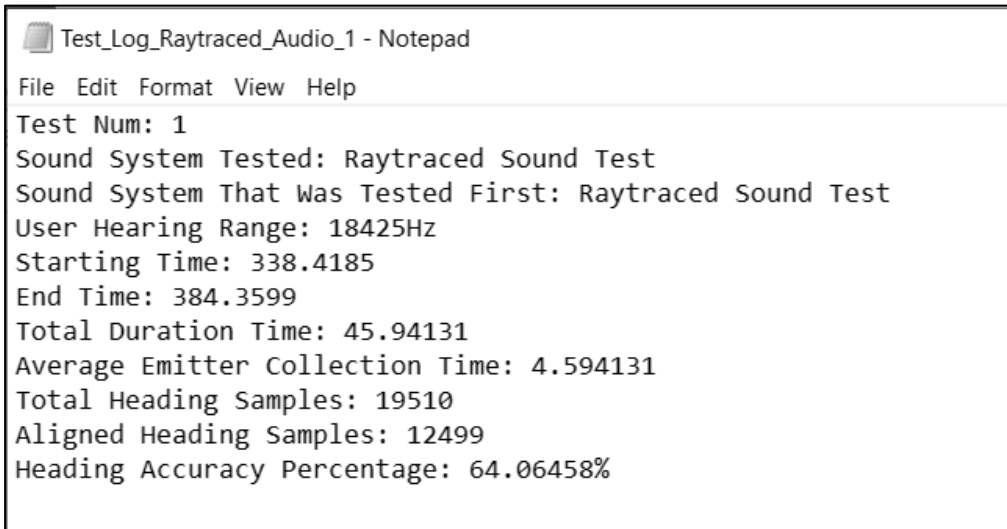
What if I have further questions, or if something goes wrong?

If this study has harmed you in any way or if you wish to make a complaint about the conduct of the study you can contact the study supervisor or the Chair of the Staffordshire University Ethics Committee for further advice and information:

Thank you for reading this information sheet and for considering taking part in this research.

14.3 AVERAGE DATA COLLECTION SAMPLE

This is the testing data format and the regular results from an average user, User Number One in this case.

Page 1/1:

Test_Log_Raytraced_Audio_1 - Notepad

File Edit Format View Help

Test Num: 1
Sound System Tested: Raytraced Sound Test
Sound System That Was Tested First: Raytraced Sound Test
User Hearing Range: 18425Hz
Starting Time: 338.4185
End Time: 384.3599
Total Duration Time: 45.94131
Average Emitter Collection Time: 4.594131
Total Heading Samples: 19510
Aligned Heading Samples: 12499
Heading Accuracy Percentage: 64.06458%

14.4 TESTING FRAMEWORK PRE-TEST SCENES

These are the testing framework pre-test scene components to make sure the audio volume is adjusted, that users have signed the consent sheet & read the information sheet, and the hearing range is tested and captured inside the user output log.

Scene 1/3:

Information and Consent Stage

Please Read the Infomation Sheet. Tick the Box Below Once Read and Understood.

If the infomation sheet has not been provided then click the button below.

Infomation Sheet Read

Please Read and Sign The Consent Sheet. Tick the Box Below Once Read and Signed.

If the consent sheet has not been provided then click the button below, sign it digitally and then return it.

Consent Sheet Signed and Returned

Scene 2/3:

Volume Adjustment Stage

Please put on headphones and adjust the volume with the slider below, then tick the checkboxes once the volume of the audio is at a comfortable level.

Headphones On

Volume Adjusted

Scene 3/3:

User Information Stage

Please Enter Your Provided User Number Below, or ask for it now.

Please Click the Button Below and perform a quick hearing range test.

Then enter your hearing range (Hz) into the box below

14.5 SPATIAL AUDIO RAW TESTING DATA

This is the raw data gathered and compiled into a spreadsheet from the spatial audio testing group.

Spreadsheet 1/1:

	A	B	C	D	E	F	G	H	I	J
1	TestNum	SoundSyst	HearingRa	StartTime	EndTime	TotalDurat	AvgEmitter	Heading	Accuracy	Percent
2	12	Spatial Sou	16745	604.961	657.78	52.8188	5.28188	63.5156		
3	14	Spatial Sou	15479	1746.51	1894.04	147.533	14.7533	51.2724		
4	15	Spatial Sou	14692	2582.01	2634.68	52.6645	5.26645	67.9015		
5	17	Spatial Sou	16325	3002.01	3073.5	71.4849	7.14849	47.2575		
6	19	Spatial Sou	17806	3969.12	4051.06	81.9348	8.19348	44.6377		
7	2	Spatial Sou	11633	286.969	345.98	59.0107	5.90107	60.0219		
8	20	Spatial Sou	17666	4554.64	4677.94	123.3	12.33	42.6245		
9	4	Spatial Sou	15213	302.115	387.98	85.8647	8.58647	44.4324		
10	6	Spatial Sou	16988	129.377	209.14	79.7632	7.97632	49.789		
11	7	Spatial Sou	15374	173.791	243.68	69.8885	6.98885	54.4621		
12	8	Spatial Sou	16484	86.0327	159.06	73.0272	7.30272	59.7241		

14.6 RAYTRACED AUDIO RAW TESTING DATA

This is the raw data gathered and compiled into a spreadsheet from the ray-traced audio testing group.

Spreadsheet 1/1:

14	TestNum	SoundSyst	HearingRa	StartTime	EndTime	TotalDurat	AvgEmitter	Heading	Accuracy	Percent
15	1	Raytraced	18425	338.418	384.36	45.9413	4.59413	64.0646		
16	10	Raytraced	19074	88.8542	161.02	72.1658	7.21658	62.7648		
17	11	Raytraced	16893	240.833	296.14	55.3074	5.53074	60.0838		
18	13	Raytraced	17380	963.537	1039.9	76.3624	7.63624	56.2433		
19	16	Raytraced	17252	2809.5	2871.82	62.3186	6.23186	54.4377		
20	18	Raytraced	14752	3454.48	3533.92	79.4368	7.94368	56.9079		
21	3	Raytraced	14132	510.786	599.4	88.6135	8.86135	54.4179		
22	5	Raytraced	15785	565.629	625.94	60.3107	6.03107	59.7607		
23	9	Raytraced	15900	500.17	574.48	74.3097	7.43097	48.9885		

14.7 SPATIAL AUDIO FINAL TESTING DATA

This is the user spatial audio groups testing data with averages included, and the user numbers fourteen and twenty's anomalous testing results were removed from the final discussion considerations.

Spreadsheet 1/1:

TestNum	SoundSys	HearingRe	StartTime	EndTime	TotalDura	AvgEmittet	HeadingAccuracy	Percent
12	Spatial So	16745	604.961	657.78	52.8188	5.28188	63.5156	
14	Spatial So	Removed	1746.51	1894.04	Removed	Removed	Removed	
15	Spatial So	14692	2582.01	2634.68	52.6645	5.26645	67.9015	
17	Spatial So	16325	3002.01	3073.5	71.4849	7.14849	47.2575	
19	Spatial So	17806	3969.12	4051.06	81.9348	8.19348	44.6377	
2	Spatial So	11633	286.969	345.98	59.0107	5.90107	60.0219	
20	Spatial So	Removed	4554.64	4677.94	Removed	Removed	Removed	
4	Spatial So	15213	302.115	387.98	85.8647	8.58647	44.4324	
6	Spatial So	16988	129.377	209.14	79.7632	7.97632	49.789	
7	Spatial So	15374	173.791	243.68	69.8885	6.98885	54.4621	
8	Spatial So	16484	86.0327	159.06	73.0272	7.30272	59.7241	
		15695.56			69.60637	6.960637	54.63798	

14.8 RAYTRACED AUDIO FINAL TESTING DATA

This is the user raytraced audio groups testing data with averages included and no user results removed. All users testing the ray-traced audio system were within the same testing margins as each other.

Spreadsheet 1/1:

TestNum	SoundSys	HearingRe	StartTime	EndTime	TotalDura	AvgEmittet	HeadingAccuracy	Percent
1	Raytraced	18425	338.418	384.36	45.9413	4.59413	64.0646	
10	Raytraced	19074	88.8542	161.02	72.1658	7.21658	62.7648	
11	Raytraced	16893	240.833	296.14	55.3074	5.53074	60.0838	
13	Raytraced	17380	963.537	1039.9	76.3624	7.63624	56.2433	
16	Raytraced	17252	2809.5	2871.82	62.3186	6.23186	54.4377	
18	Raytraced	14752	3454.48	3533.92	79.4368	7.94368	56.9079	
3	Raytraced	14132	510.786	599.4	88.6135	8.86135	54.4179	
5	Raytraced	15785	565.629	625.94	60.3107	6.03107	59.7607	
9	Raytraced	15900	500.17	574.48	74.3097	7.43097	48.9885	
		16621.44			68.30736	6.830736	57.5188	