Organised Sound (2025), **30**, 204–214
doi:10.1017/S1355771825000044

CAMBRIDGE UNIVERSITY PRESS

Article

Open Ambisonics Toolkit: A low-cost hardware–software–theory approach to spatial audio

Giuseppe Pisano[1] and PerMagnus Lindborg[2]

[1]Norwegian Academy of Music, Oslo, Norway and [2]SoundLab, School of Creative Media, City University of Hong Kong, Hong Kong SAR, China

Abstract

The increasing number of applications for spatial audio technologies has led to a growing interest in the subject from academic institutions and a more capillary diffusion of techniques and practices to non-institutional contexts, especially independent sound artists. However, the lack of a methodology for learning these technologies motivated our team to develop the Open Ambisonics Toolkit (OAT). Our goal is to promote the diffusion of spatial audio technologies by combining three pedagogical components: a DIY approach to hardware, a selection of open-source software, and a step-by-step introduction to Ambisonics theory through practical applications. The present article focuses on the development of a flexible toolkit and is based in our own practical experience as sound artists and teachers. We describe the process of designing hardware and selecting software components, and report results from objective measurements and listening tests conducted to evaluate different loudspeakers and spatial configurations. To conclude, we discuss future perspectives on the development of tutorials for learning spatial audio with OAT, which we are continually testing in workshop settings with students and independent sound artists.

(Received 15 September 2024; revised 9 December 2024; accepted 24 January 2025; First Published online 15 May 2025)

1. BACKGROUND

Spatial audio is experiencing a surge of interest. Thanks to faster computers, the accessibility of digital audio technologies and the proliferation of virtual reality (VR), more and more artists, content creators, musicians, sound designers and engineers are turning their attention to immersive audio. Dolby Atmos and Ambisonics appear to be tied for the lead in terms of adoption, but while the former is a patented technology used primarily in the film and mainstream music industries, Ambisonics is an open-source, free alternative with a wider range of applications. This is especially true for the video game industry (Horsburgh et al. 2011), thanks to the ability to decode Ambisonics into binaural stereo, which, combined with head-tracking, makes Ambisonics an optimal choice for VR and augmented reality (AR) applications.

Among the first major contributions to the development of Ambisonics there is the work done by Michael Gerzon at Oxford in the 1970s, which for several reasons did not achieve much success initially (Hulme 2018). It survived as a niche technology in a handful of universities, first in the UK and then globally (Malham 2019),[1] which resulted in a range of diversified directions of development. The scene is now changing, thanks to the harmonisation of formats and methodologies that is helping to set workflow standards (Nachbar et al. 2011) and the availability of user-friendly and free tools (IEM 2017; McCormack and Politis 2019). This change in the landscape has made it essential to include Ambisonics in the curriculum of any higher education institution. However, owing to the open-source nature of Ambisonics and its relatively recent entry into consumer applications, educators cannot yet rely on a formalised method for teaching this technology. In addition, the information available online about Ambisonics is often confusing. At the time of writing, the most common types of publicly available information are as follows:

- Tool-specific materials: the instructions, tutorials and documentation that come with specific software packages and hardware usually provide some theoretical contextualisation, although they are not always objective (e.g., no discussion or reference to alternative workflows) (see Blue Ripple Sound 2019).
- Academic papers: this is the core of Ambisonics knowledge and is crucial, as it covers scientific aspects of the technology. However, the presentation is often too advanced for beginners and not always practical for musicians, composers, designers and other end users.
- Independent divulgation: these are mainly tutorials on YouTube, usually on channels that cover sound production topics. Some of this information can be useful, but it is not subject to fact-checking or peer review, which can lead to inaccurate information (Falch 2022).

Our work with the Open Ambisonics Toolkit (OAT) addresses this situation by filtering the existing information for accuracy,

[1]In addition to the University of York we can include the Institute for Electronic Music and Acoustics – Graz; IRCAM – Paris; CNMAT – Berkeley, California; CCRMA – Stanford, California; DXARTS – Seattle, Washington; Virginia Tech – Blacksburg, Virginia.

Corresponding author: PerMagnus Lindborg; Email: pm.lindborg@cityu.edu.hk

Cite this article: Pisano G and Lindborg PM. (2025). Open Ambisonics Toolkit: A low-cost hardware–software–theory approach to spatial audio. *Organised Sound* **30**: 204–214, https://doi.org/10.1017/S1355771825000044

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike licence (https://creativecommons.org/licenses/by-nc-sa/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use.

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

relevance and depth and summarising it on a platform where hardware, software and educational information will be available. At this point in time, 50 years after Gerzon's pioneering work, we believe it is important for music and sound design students to integrate Ambisonics technologies into their skill set. Likewise, creative media students in film, games, AR/VR and so on should understand and be able to use spatial audio at a high level. All media and arts students need exposure to immersive audio. In such contexts, OAT demonstrates the affordances of Ambisonics by providing students with enough theoretical knowledge to understand how spatial audio works, while focusing on practical, hands-on applications that generate ideas and transferable problem-solving skills.

1.1. Context

The development of OAT began at SoundLab, School of Creative Media (SCM), City University of Hong Kong. Despite having nine universities, Hong Kong is rarely mentioned when discussing the development of sonic arts in an international context, perhaps due to the lack of dedicated electroacoustic music institutions which have benefitted local scenes in neighbouring regions (Battier and Fields 2020), and more generally, the democratisation of music technology (Keller et al. 2014). In Hong Kong, the curricula supported by arts and creative industries, and higher education departments, are often broad in nature. The responsibility for creating a context for these disciplines is left to the work of individual faculty members, who are often also active practitioners and usually join the efforts of external organisations to organise events outside the campus. There has been considerable progress in the scene of local sonic arts over the past decade, spear-headed by a generation of practitioners whose characteristics are quite different from the experimental music makers of the past (Ikeshiro et al. 2022). In this context, SCM fully embodies interdisciplinarity in bridging a range of fields, including game studies, human–computer interaction, digital humanities, design, sound, VR, under the headline of 'creative media'. To meet pedagogical needs in relation to spatial and immersive audio, we designed a toolkit that focuses on the most common issues and provides all the necessary information. OAT started as a granted research project with an emphasis on hardware–software development geared at pragmatical solutions for spatial audio. As part of the grant conditions, our approach is practice-oriented and versatile to meet the needs of a wide range of practitioners wishing to use Ambisonics for their applications.

1.2. A hardware–software toolkit for exploring and learning spatial sound

In our experience of teaching spatial sound, it has proven to be fundamental that students gain access to different types of surround sound systems in order to expand their knowledge and understanding of these technologies, and to prepare them for working with immersive media both in theory and in practice. While training students in high-quality facilities using the latest and most advanced products will make them competitive in a professional context, presenting them with highly customisable alternatives can stimulate students' creative imaginations, and allow them to design spatial sound for physical and virtual environments. As a learning experience, this may reveal the underlying functionalities of larger systems, thereby helping to elevate the learner's mindset from that of a consumer or end-user, towards that of a creative media expert.

This approach has been taken by educators such as Gavin Ambrose whose undergraduate course at University of Brighton on sustainable design integrated a module dedicated to building a loudspeaker from found objects and waste (Ambrose 2018), and Otso Lähdeoja who developed an 'orchestra' of found sound objects and loudspeakers, together with strategies for composing with these instruments (Lähdeoja 2016). The second author of this article used a similar DIY method to create an auditory display with 21 small loudspeakers that used metal bowls as enclosures and a low-cost multichannel playback system (Lindborg 2015).

Through the experience of building their own set of loudspeakers, students will learn basic technical skills and develop a deeper understanding of spatial audio and problem-solving skills through experimentation with patches and software solutions that will be provided and explored with students in guided classroom activities. The perception of spatial sound is trained through experiments in room acoustics and psychoacoustics, as well as through critical-evaluative listening to different loudspeaker setups and configurations, and students can be challenged with tasks that require them to apply these skills, such as adapting a sound installation to a given context putting this theoretical knowledge into practice. We believe that this pedagogical emphasis benefits systems thinking, exploration, multi- and interdisciplinarity, and collaboration.

Keeping equipment and software costs low is fundamental as an egalitarian principle that promotes inclusivity, engagement and a sense of student-driven ownership of their learning experience. For this reason, open-source operating systems and software (i.e., Linux, Pure Data, Octave) and microcomputers (Raspberry Pi) were chosen. The overall cost of hardware items is an important factor for how to roll out the solutions in a classroom as part of coursework (for a budget outline, see Supplementary Materials 1). Therefore, the methods and tools we develop are modular; for example, OAT software (Pure Data patchers) can be introduced in audio/music classes focusing on sound mixing, installation and spatialisation. DIY construction of loudspeakers and micro-controllers, with evaluation and potentially advanced applications, might be topics of individual studio project; for example, as an undergraduate thesis work. These course modules or projects can be conducted as individual creative works, or in interdisciplinary workshop groups. The pedagogical project focuses on the development of practice-oriented technical skills and know-how in the field of surround sound design. Hardware and software are integral parts, always in parallel, to create a toolkit that is scalable, easy to maintain and inexpensive.

The components of any toolkit such as OAT would have to be designed in an iterative process that includes listening tests and on-site, live applications. Within the timeframe and scope of the present project, we have conducted three evaluations (reported in the next main section) and three workshops (Lindborg and Pisano 2023; Pisano and Lindborg 2023a, b; Lindborg et al. 2024), with more in the pipeline. At this point, it is important for us to gain feedback on the spatial audio pedagogy and to train users to employ and test OAT materials in their creative work.

1.3. Toolkit modules

The toolkit we have designed aims not only at providing the necessary theoretical information, but also to design, present, and document software and hardware solutions that would allow students to use Ambisonics in a variety of contexts. These could go from designing sound installations at an art exhibition, to

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

monitoring their work in a home studio setting (Prior 2018). The democratisation of music technologies and the emergence of 'bedroom producers' could lead to universities or public-supported research centres being less central to the development of sonic arts (Ikeshiro et al. 2022). We believe that educational institutions need to bridge the gap between professional audio industries and creative media cultures by integrating tools that promote critical thinking.

We have therefore designed OAT as interconnected modules, where each is also useful on its own. The three main modules are: 1) hardware, in which we describe solutions to build inexpensive loudspeaker setups for Ambisonics in a DIY fashion; 2) software, where we present software tools available for Ambisonics work with a specific focus on free/open-source software compatible with Linux and Raspberry Pi; and 3) theory, where we explain Ambisonics in detail and provide tutorials and other forms of documentation. In the next phase of development, we will stimulate practical and pedagogical student projects using OAT to build a database of applications or case studies. Tracking the results of student projects using OAT, both artistic and perceptual listening tests, will accumulate evidence in support of Ambisonics in general and OAT in particular, to demonstrate how this approach can support learning spatial audio.

1.4. Hardware

Working in Ambisonics can be demanding, particularly when it comes to hardware. A common Ambisonics rig in a studio context uses between 8 and 16 loudspeakers, while in concert spaces there are usually between 24 and 48 loudspeakers, or more than 200 for The Cube at Virginia Tech (Lyon et al. 2016). There is a significant upfront cost in equipment, especially acoustically suitable facilities, and equipment such as loudspeakers that are usually high-grade studio monitors, and powerful computers and audio interfaces. It goes without saying that similar systems are beyond the reach of most students and young artists, who must rely on the access to schools and other institutions equipped with such gear to develop their multichannel spatial audio projects. Institutions of this type are few and far between, and most have strict access requirements. A partial remedy for this is the practice of decoding Ambisonics mixes and sound fields into binaural stereo for headphone mixing. However, not only do most earbuds or headphones have mixed acoustic performance (Lim and Lindborg 2013), but, to ensure accurate spatial reproduction in binaural, users should also use an individual Head Related Impulse Response (HRIR), but this is cumbersome and/or costly to produce (Guezenoc and Seguier 2020). Furthermore, binaural techniques do not solve all the problems of implementing spatial audio in public spaces, and while it may be useful in the production phase, other solutions might be required for physical installations; for example, at galleries or exhibition rooms.

For these reasons we imagine that building a small, portable and relatively inexpensive multichannel loudspeaker setup could become an asset for students and artists who want to emancipate their workflow and explore Ambisonics on a shoe-string budget. In what follows, we will describe the design process for OAT hardware.

1.4.1. Loudspeakers

Concerning loudspeakers, we chose in-car systems because of their price point and because most modern in-car speakers are coaxial, which means the tweeter is inside the woofer, so a two-way system can be packed into a single unit. Moreover, the behaviour of coaxial



Figure 1. Loudspeaker driver for cars Kenwood KFC-1066S.



Figure 2. Loudspeaker driver for cars Pioneer TS-F1634R.

loudspeakers in spatial audio application is a closer approximation of an audio point source than their non-coaxial two-ways counterparts, which provides increased directivity (Kessling and Görne 2018). The models we considered suitable for this first phase of experimentation were the Kenwood KFC-1066S 4″ (Figure 1) and the Pioneer TS-F1634R 6.5″ (Figure 2). These two speakers are nominally similar in terms of performance. The main difference is in the frequency response, with the Pioneer claiming to go down to 31 Hz compared with the Kenwood's 45 Hz (see Supplementary Materials 2). However, our choice was also motivated by the fact that they were readily available in a local car stereo parts shop, reflecting a practice of sourcing materials that we wanted to encourage in the students. In our test we will determine which driver performs best, using both frequency response measurements and a perceptual evaluation.

1.4.2. Amplifiers and power supply

For the amplifiers (Figure 3), we decided to use pre-assembled aftermarket digital stereo amplifiers that we found in local electronics shops in Hong Kong. They operate on 12 V and deliver 20 W per channel into 4 ohms, not exceeding the RMS power of the speakers we chose (25 W RMS for the Pioneer and 21 W for the Kenwood). These are D-class amplifiers, which are known to be the cheapest and most energy-efficient class of amplifiers commonly

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press



Figure 3. Generic stereo amplifier for cars.



Figure 4. Power Supply Mean Well LRS-350-12.

used in car audio (Jiang 2017). Concerning the power supply (Figure 4), we are using a Mean Well LRS-350-12, which is standard in electronics development, delivering up to 350 W at 12 V. This choice allows us to comfortably deliver power to four 50 W amplifiers.

1.4.3. CPU: Raspberry Pi

We decided to use Raspberry Pi (Figure 5) to control our system. Running 64x Raspberry OS seemed like a good solution for implementing software tools in Pure Data, Octave and Supercollider. In addition, with the release of RNBO by Cycling '74, it is now possible to run Max patches as standalone applications on Raspberry Pi, further extending the capabilities of this microcomputer.

1.4.4. Audio Interface

There are several options for audio interfaces that work with Raspberry Pi. Typically, most class-compliant USB audio interfaces will work, but not many offer official Linux support, and even fewer have available control software. It is almost impossible to keep track of audio hardware devices and their compatibility with different Linux releases, and most websites listing compatible devices struggle to be up to date (Carneiro 2016; Linuxaudio 2013). One solution for us was to consult dedicated forums and repositories, such as the Reddit subgroup r/linuxaudio, which has an active and helpful community. For our project, we decided to use Audioinjector Octo (Figure 6), a multichannel audio interface designed specifically for Raspberry Pi. It offers six channels of input and eight channels of output in a compact setup. Other low-cost options include the Behringer UMC1820, which we are currently testing.

1.5. Hardware costs

The total cost of building an eight-point-one-speaker system is approximately $1,400, including accessories such as microphone stands to support the speakers, a carrying case and a stool. Costs were kept low by purchasing almost all materials from local vendors (for details, see Supplementary Materials 1), but we believe that an OAT system could cost even less with further optimisation of sourcing.

1.6. Software

The software tools for this project were chosen because of compatibility and price. We only included free software, with a strong preference for open-source programs. This was not a strong limitation for the design of our workflow, mainly because most of the Ambisonics tools are developed in academic institutions and are free and well supported. In terms of compatibility, the software we used had to be compatible with the Raspberry OS release of Linux and ARM processors. Note that not everything that runs on Linux will run on a Raspberry Pi. To ensure that all students have easy access to these tools, we also favoured software that could be downloaded and installed directly from the Linux APT or, in the case of Pure Data, from external libraries. A few exceptions to this are a few Pure Data libraries that are not yet available for Raspberry Pi but are expected to be.

Audio in Linux relies on a specific framework that includes Advanced Linux Sound Architecture (ALSA) and a sound server. ALSA provides audio and MIDI functionality to Linux (Alsaproject 2019). Typically, applications communicate with ALSA through a server; the two most popular are PulseAudio (Debian Wiki 2010) and JACK (Jackaudio 2006). While PulseAudio is mostly used by consumer-grade apps, only providing basic routing and volume control, JACK is its professional counterpart, used by DAWs and other music-making applications, providing low-latency functionalities and advanced routing options.

The Ambisonics user community is now a large and very helpful one, but High Order Ambisonics presents some concepts that are quite difficult for a beginner to grasp. We have taken this into account in the selection of the software, trying to balance ease of use with processes that have a pedagogical function and imagining a good common end goal for a hypothetical class. Designing a functional decoder capable of playing back an Ambisonics file seems appropriate, as it is necessary, presents a reasonable challenge, and provides valuable insights and theoretical background for OAT users.

1.6.1. Pure data and libraries

Among open-source programming languages for audio signal processing, both Pure Data and Super Collider provide necessary features to work with Ambisonics on Raspberry Pi (ATK Community 2016; Kjeldgaard 2022). We have chosen Pure Data partly because it is closer to Max, which is what most students at the School of Creative Media are already acquainted with.

We explored three Pure Data libraries dedicated to Ambisonics: iem_ambi (Musil 2000), CEAMMC Pure Data (Poltavsky and Nadzharov 2021) and [vstplugin~] (Zmöelnig 2019). They all

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

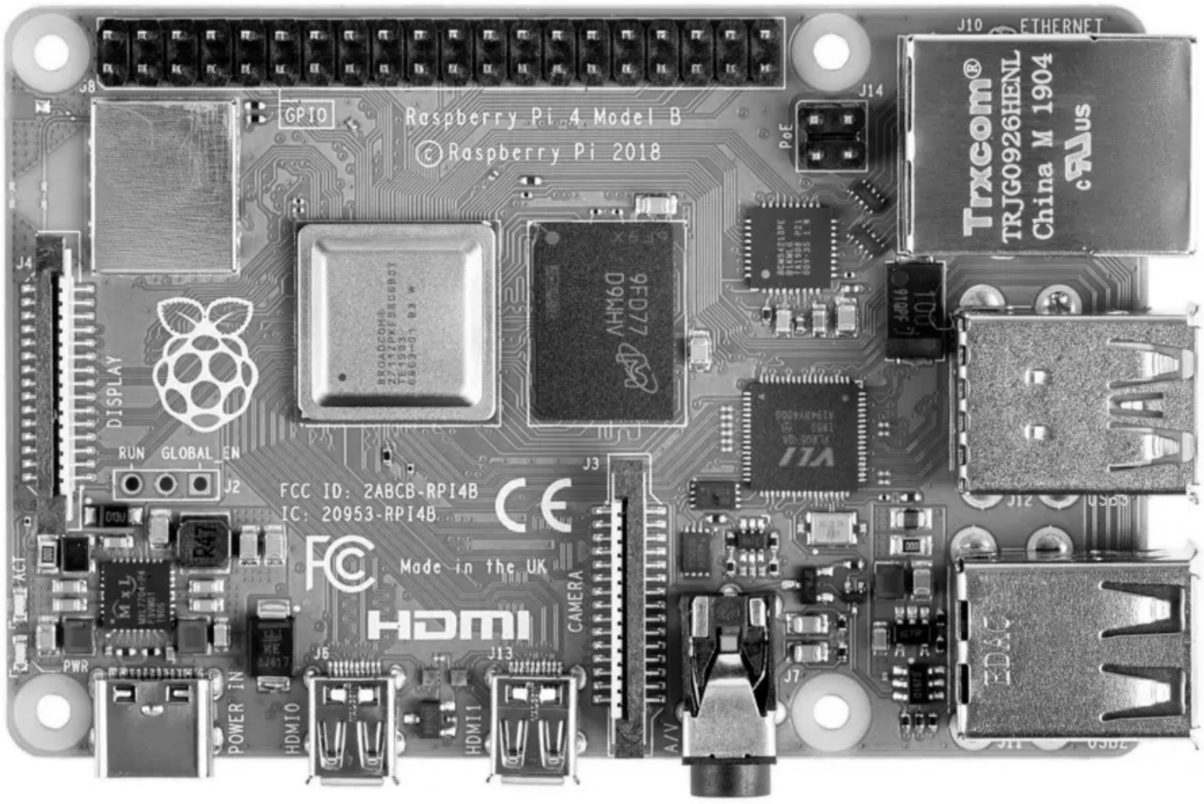


Figure 5. Raspberry Pi.

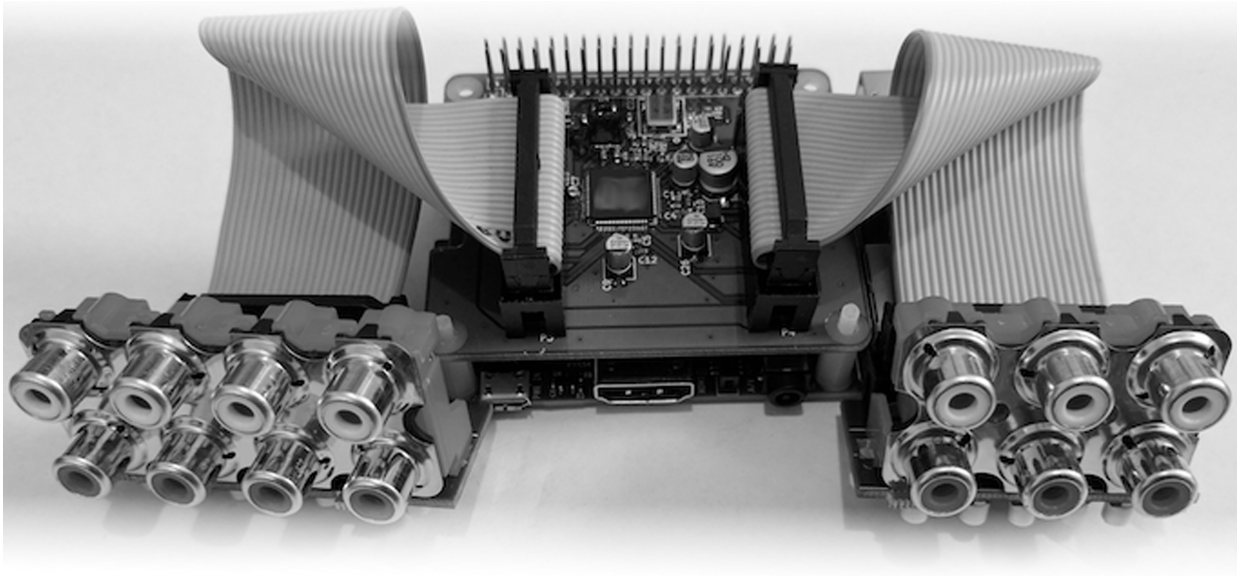


Figure 6. Audioinjector Octo Audio interface.

represent different approaches, each with its pros and cons, but at the time of writing none of them alone is suitable for the project. On the other hand, they can be synergistic to varying degrees, and all three are worth exploring to gain some theoretical insights.

- Iem_ambi is a library developed by IEM at Kunstuni Graz. It integrates with other libraries developed by IEM and runs perfectly on Raspberry Pi. However, the library appears to be partly outdated, using a very unusual channel order (e.g., SID) (Daniel 2003) that requires conversion, and a decoding method (mode-matching) (Poletti 2000) that is not ideal for irregular loudspeaker setups; however, it is very useful for encoding in Ambisonics, because it operates at the closest level to coding of the three and it requires a fair amount of patching and provides

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

insight into some of the mathematical processes usually hidden in other software packages. This is both a challenge and an opportunity for students.

- CEAMMC (formerly known as HoaLibrary) is a project developed by CIMC between 2012 and 2014 (Sèdes et al. 2014). It was discontinued but is now being incorporated into a larger library at the Moscow Conservatory Centre for Electroacoustic Music. While much easier to use than the Iem_ambi approach, it is currently limited to 2D Ambisonics. CEAMMC includes useful additional objects such as a reverb, several different panners and three decoding methods (basic, in-phase, and Max-rE) (Murillo et al. 2014), allowing for dual-band decoding options (Heller et al. 2008). This library works on Linux, but is not yet implemented on Raspberry Pi, but we found it useful to include in the teaching materials.
- The external object [vstplugin~] loads existing VST plug-ins in a wrapper and is therefore the easiest of the three packages to use. Known plugin suites for Ambisonics such as the IEM plugin suite (IEM 2017) and the ambiX plugin suite (Kronlachner 2011) are available for Linux and can easily be used in the wrapper. However, the specific syntax of the wrapper makes parameter declaration somewhat counterintuitive, making it the least ideal choice from a pedagogical point of view. However, using [vstplugin~] is very convenient for binaural decoding. The [vstplugin~] has recently been released for ARM processors, but it does not show compatibility with all VST plugins, but the project is under constant development on the project's GitHub repository (Ressi 2023).

1.6.2. ADT + AmbDec

A fourth solution we explored is based on the synergy between a matrix calculation software called Ambisonics Decoder Toolbox (ADT) (Heller et al. 2014) and a software called AmbDec (Adriaensen 2009), a Linux-compatible decoder with dual-band decoding, distance compensation and near-field compensation (Heller et al. 2008). ADT comes as an external for GNU Octave or as a FAUST package (Heller 2014) and only computes the coefficient matrix for the decoder given a loudspeaker array. The coefficient matrix for each loudspeaker layout is saved as a.m file. The file is then imported into AmbDec, which performs the audio processing.

Using this selection of software in combination with Pure Data as a modular playback engine for multichannel files and Jack for routing proved to be the most suitable approach for our project. This choice provides OAT with a high degree of flexibility and reliable sonic output. As a pedagogy, opting for a DIY assembly of software components has both positive and negative effects. On the positive side, hands-on experimentation with speaker configurations generated theoretical insight; for example, about spatial audio perception, or the structure of the software system. On the negative side, the workshops have taught us that the process of installing and configuring the OAT software can be finicky and cannot realistically be mastered in a few hours. To lower the threshold of entry, we have written a careful step-by-step manual which we made available to our students together with theoretical materials and example Ambisonics patches.[2]

1.7. Theory

In the final module of our project, we focus on developing a theoretical method to help students navigate the wide range of widely available and highly diverse learning resources on Ambisonics. The process we adopted – which is a work in progress – is to map the available information and filter it according to accuracy, relevance and depth in order to create a coherent and gradual pedagogical path. The aim is also to produce a document in which this information will converge to convey the theory behind Ambisonics, with the specific aim of tuning this knowledge to the work of sound designers and composers, explaining its main concepts and processes, while focusing on its practical applications. We have identified several different topics to cover, such as elements of psychoacoustic and spatial listening and perception, A-to-B format transcoding techniques, encoding techniques, decoder design and signal processing in Ambisonics. The chapters will provide an easy understanding of key concepts and provide references to other articles for more in-depth study, as well as pieces and works organised in curated listening sessions. We expect this friendly introduction to Ambisonics to be the second article output of this project. Further developments may include a blog where prominent users of spatial audio technologies will describe their approaches and techniques.

2. EVALUATION

In the process of designing the DIY multichannel loudspeaker arrays, we carried out three successive evaluations: measurements of loudspeaker responses; listening test of stereo speakers; and listening test of speaker configurations. The first two evaluations compared 'candidate' loudspeaker against reference models, first by objective measurements and then by participants listening to music and voice recordings. This led us to choosing the Kenwood speakers. In the final evaluation, we created three different spatial setups to estimate the quality of the sound fields reproduced across these configurations. We invited MA and PhD students from SCM to participate in the listening tests, also to try out protocols that could be incorporated into OAT as a pedagogical model. The listening tests that we designed are, as a secondary outcome, inspiring the student-driven experimentation in learning settings.

2.1. Loudspeakers evaluation

The first stage of the evaluation process was to measure the response of each loudspeaker to determine which model would be most functional for our project. This preliminary selection was based on the models that were available for purchase in local shops in the Sham Shui Po area of Kowloon, Hong Kong, in the price range of $20 to $35 per pair. The drivers were mounted in custom-built loudspeaker cabinets purchased from the same vendors. The loudspeakers chosen as references were Genelec 8020B and Auratone A2-30. The former is a common and widely respected near-field speaker, relatively close in size to the DIY designs built for this project. The latter is a studio monitor designed to sound like 'real life' loudspeakers (car systems, lo-fi stereo systems) and is often used by sound engineers to cross-check their mixes (for details, see Figure 7 and Supplementary Materials 2).

2.1.1. Room

The materials were set up in a dedicated sound studio at SCM. The studio was a rectangle measuring 2.80 × 2.45 m with an area of approximately 6.8 m^2. The height was 4.08 m, giving a volume of approximately 28 m^3. The studio had some reverberant soundproofing material on the walls. The background noise level was measured to be LeqA (1 min) = 31.5 dB using a calibrated StudioSixDigital iTestMic2 and AudioTools app (Studio Six

[2]https://drive.google.com/drive/folders/1JR9nfK6LIIb7Geu0WqtkCyIUaF_WW9Lt.

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

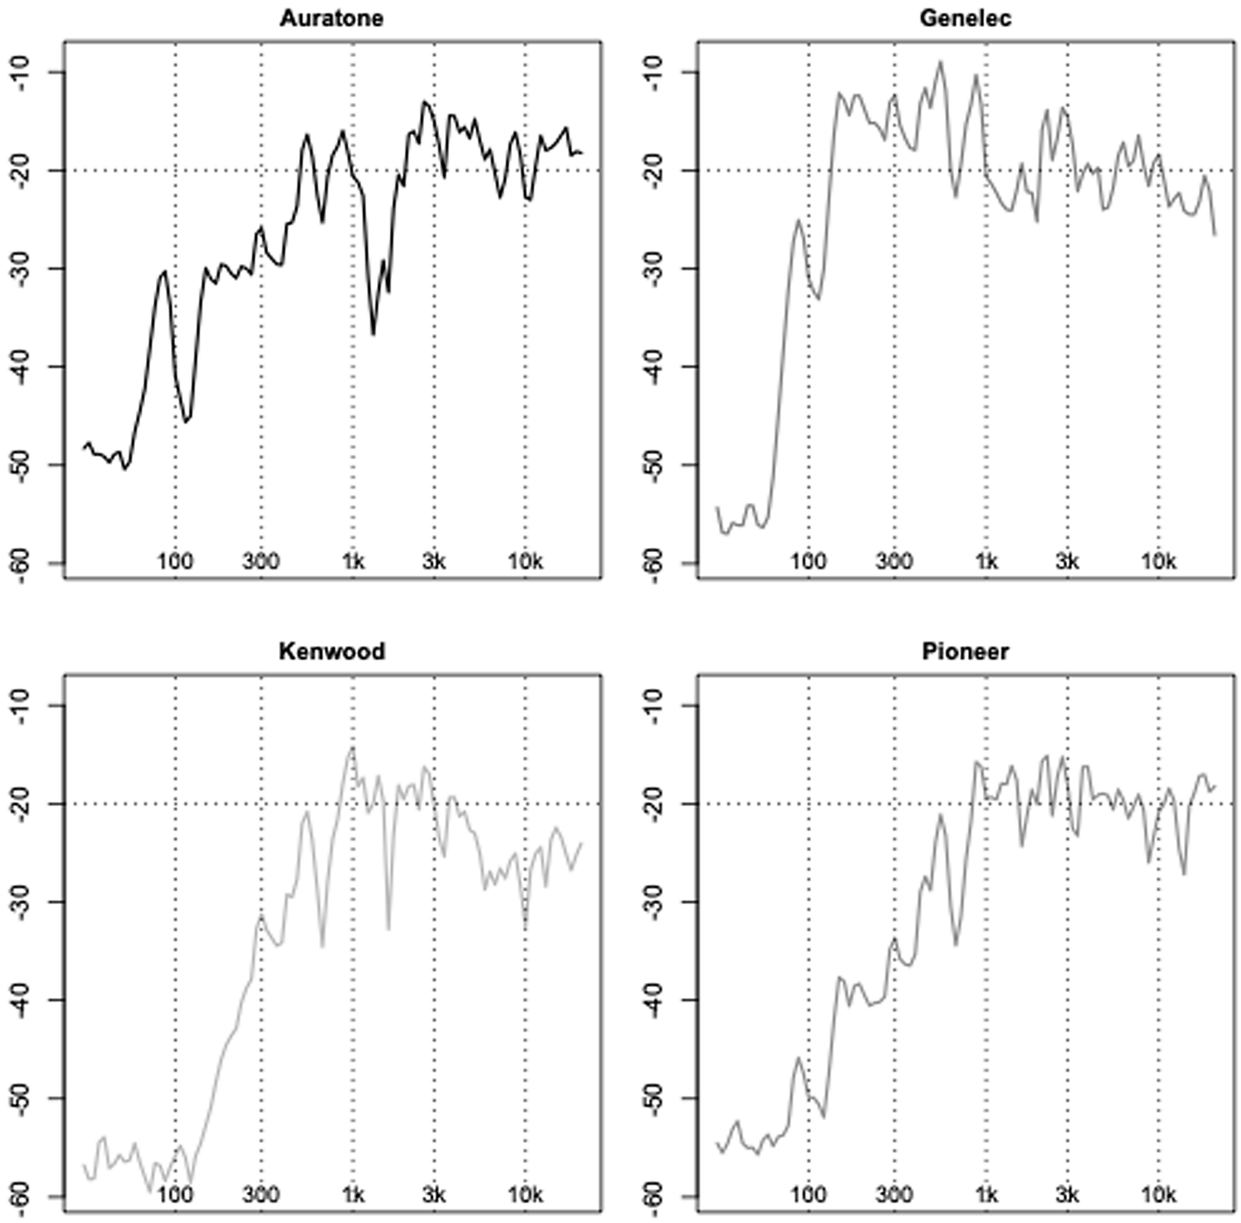


Figure 7. Impulse response plotting of the four loudspeakers tested.

Digital 2022). Room reverberation as measured with the BuzziSpace app was estimated to be 0.20 s (RT60 for frequencies above 250 Hz). Using broadband noise, the amplifiers were adjusted so that the loudspeakers produced the same measured SPL at the listening position. The room was set up like a normal home studio with minimal acoustic treatment, representing the listening conditions in which such systems are most likely to be used.

2.1.2. Measurements

Measurements were taken in accordance with best practice (British Standards Institution 2009). A time-smoothed 'sweep' (40 seconds sine wave at 25 to 20000 Hz, log–frequency) was recorded separately from each loudspeaker. The signals were generated by a custom Pure Data patch and played through a Raspberry Pi mounted on the Audioinjector–Octo audio interface. Two calibrated microphones were tested to record the responses: Behringer ECM8000 and Rode NT4 (with an omni capsule), both connected to a Focusrite Scarlett 2i2. The results from these two mics were very similar, and Behringer (which is the least expensive) is reported here. As can be seen in Figure 7, Genelec had the flattest response; Kenwood had a reasonably flat response from around 400 Hz, while Pioneer and Auratone had a considerable roll-off from around 800 Hz. The response diagrams confirmed our general listening impressions.

2.2. Listening test: stereo pairs

To decide the project's loudspeaker, we conducted a listening test that included four pairs of the above-mentioned measured loudspeakers, in stereo configurations. The purpose was to perceptually evaluate their performance in an ecologically valid setting. The room used for this listening test was the same as for the loudspeaker response measurements.

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

2.2.1. Participants and procedure

Students and faculty members from the researchers' lab were invited to participate in the listening test (N = 8). There were five females and three males, between 22 and 54 years old (median = 28). No one reported any hearing impairment. The evaluation was performed using a subset of the Spatial Audio Quality Inventory (SAQI) (Lindau et al. 2014), including 15 rating scales relevant to stereo configurations. All the participants had previous experience with sound-related topics and technologies in a wide range of practices and at various levels, which guaranteed a level of expertise using the vocabulary of SAQI, which was distributed to them before the listening test. The researchers assisted in explaining any technical terms by clarifying the definitions used in SAQI. The participants were individually seated on a low stool facing the loudspeakers, at angles of ±30° and 1.5 m distance. They were tasked to perceptually evaluate and rank the loudspeaker models along the semantic scales. They could freely switch listening materials within a selection of music pieces (mainly pop music pieces of different styles such as electronic dance music, rock, funk, ballads) and could also at any time switch between the four loudspeaker pairs. Each participant took around 50 to 60 minutes to complete the evaluation.

2.2.2. Results

While some of the detailed findings might be of interest, for the purposes of brevity, we limit ourselves to reporting participant ratings on one general scale, *Degree-of-Liking*. In-depth comments provided by some participants were separately transcribed and showed that preferences depend highly on the type of sound material. The Kenwood speakers had a *Degree-of-Liking* almost as high as the reference Genelecs (which cost many times more), and as an optimal compromise between cost, listening test results, and construction, they were adopted for the next development.

2.3. Listening test: spatial configuration

In the third evaluation, we compared various spatial configurations, each using eight-point speakers and a subwoofer. The purpose was to determine, from a perceptual standpoint, which of the three might be better at reproducing a sound field.

2.3.1. Room

As this part of the experiment required more space, the materials were set up in the Centre for Motion Capture Experiments (MOCAP) studio at the School of Creative Media. The studio is a shoebox-shaped room measuring approximately 12 × 5 m with a ceiling height of approximately 4 m, giving a volume of 240 m^3. The room has some reverberant soundproofing materials on the walls, mainly black curtains and a moquette floor, resulting in a medium dry reverberation characteristic (RT60 was 0.65 s, estimated as before). The loudspeaker configurations occupied a smaller part of the room and were installed so that the average diameter of the setup was 3.0 m. This ensured a reasonable loudspeaker density. It can be considered as a model for a production or mixing studio with minimal acoustic treatment, or a relatively small exhibition space.

2.3.2. Participants

The participants (N = 8) were from the researchers' laboratory and a workshop in spatial audio held concurrently. They were aged between 24 and 54 years (median = 29), and four were female. Some but not all had participated in the previous listening test, and none reported any hearing impairment.

2.3.3. Evaluation

The evaluation was carried out using a slightly different subset of SAQI that included semantic scales relevant to spatial audio, that is, primarily related to the perception of sound source image and spatial characteristics. We selected 19 rating scales in several categories: *Timbre* (Tone colour), *Geometry* (Distance, Depth, Width, Height, Localizability, Spatial disintegration), *Time behaviour and Dynamics* (Echoes, Sequence of events, Loudness, Dynamic range), *General* (Artefact, Clarity, Voice intelligibility, Naturalness, Presence), *Degree-of-Liking*, and finally, two additional scales labelled *Immersivity* and *Source distribution*. As before, all participants had previous experience with spatial sound and related technologies at various levels and were sufficiently familiar with the relevant terminology. Nevertheless, each participant was individually briefed to ensure that they had thoroughly understood and integrated the terms used in the SAQI, and any doubts that arose were explained.

2.3.4. Materials

The sound material used for this part of the experiment included field recordings, excerpts from acousmatic music pieces, and vocal material. The field recordings came from the first author's personal practice and were all recorded using a Zylia ZM-1 microphone; they included both indoor and outdoor spaces in both natural and urban contexts. The acousmatic music excerpts included excerpts from pieces by Brona Martin, Ernst Van der Loo and Giuseppe Pisano. The vocal material included songs and speech material commissioned for OAT, kindly provided by Ernst Van der Loo and Mariam Gviniashvili. The speech was recorded using a Zylia ZM-1 microphone, and the song material was a mono recording encoded in Ambisonics 3D using the IEM Room Encoder. All files were supplied in third-order Ambisonics format, to ensure that playback was possible via third- or lower-order decoders. The duration of all excerpts was between 1 and 2 minutes, and the duration of the entire playlist was approximately 12 minutes. The sounds were played in pseudo-random order to counteract sequence bias.

2.3.5. Configurations

The following three loudspeaker configurations were constructed:

- OCTO: a ring of eight equidistant speakers. Sounds were played back on this system using a 2D (azimuthal) third-order decoder.
- HEMI: a hemisphere with two rings of five and three equidistant speakers, respectively. The first ring was at ear level and the second ring at 45° elevation. Sounds were played back with a 3D second-order decoder.
- CUBE: a skewed cube composed by two rings of four loudspeakers each, horizontally offset by 45°. The two rings were located at 45° above and below ear level. For this configuration, sound playback was through a 3D first-order decoder.

These configurations required clear trade-offs in decoder design. For example, OCTO presents the maximum possible loudspeaker density at ear level, allowing third-order decoding, but no height. On the other hand, CUBE distributes the sound more evenly across the surface of a sphere but offers lower spatial resolution and presents no speakers at ear level. HEMI is a

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

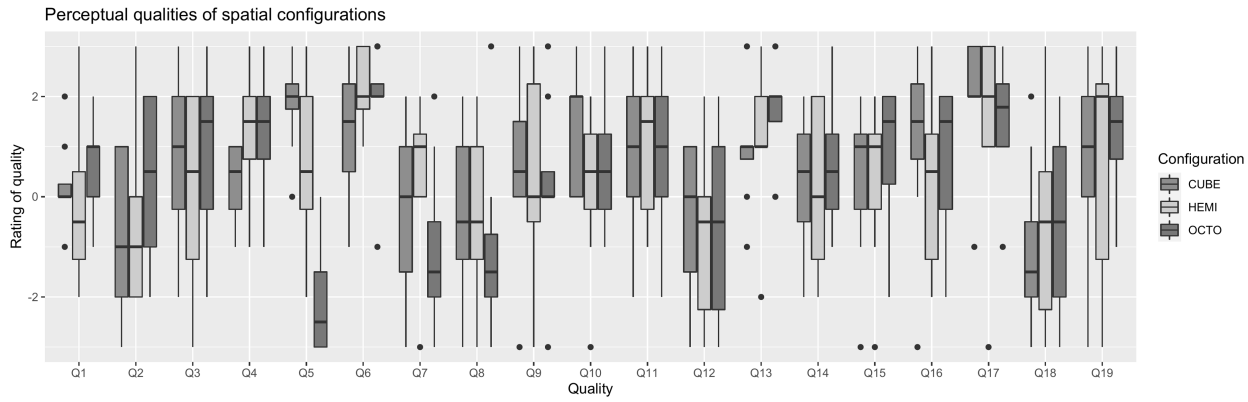


Figure 8. Perceptual qualities of spatial configurations.

compromise between the two, with relatively good spatial resolution at ear level and a larger sweet spot size than CUBE due to the higher-order decoder. However, it only covers the upper hemisphere of the Ambisonics sound field.

2.3.6. Decoders

AmbDec was used as our Ambisonics decoder. The filter matrices were created using the Ambisonics Decoder Toolbox on GNU Octave. Pure Data was used as a multichannel player and gain matching module for the loudspeaker signals. Routing was done in Jack.

2.3.7. Listening test

The test took place over two days to allow the researchers to reconfigure the setups; each day one group of participants would listen to the three configurations in a different order, to counteract sequence bias. The first day participants listened to the three configurations in an OCTO–HEMI–CUBE order, the second day participants listened in the reverse order, CUBE–HEMI–OCTO. This was a functional compromise that allowed us to rearrange the loudspeakers as little as possible, hence improving the consistency of the experiment. Measurement for the physical placing of loudspeakers was done with a laser meter and a protractor. For each configuration, the appropriate filter matrices and the decoders were set.

2.4. Results

The present analysis focuses on how the primary variable of interest, *Degree of Liking*, relates to ratings on the other 18 scales of SAQI (Lindau et al. 2014). As can be seen in Figure 8, the mean *Degree of Liking* was highest for the HEMI configuration, followed by OCTO and CUBE. Across all scales, the interrater agreement was moderate (standardised alpha = 0.64). It was higher for the OCTO configuration (0.75) than for CUBE (0.64) or HEMI (0.43), which might indicate that participants found it harder to evaluate spatial audio quality when height was an added factor. We conducted exploratory factor analysis on the 18 scales (excluding *Degree of Liking*). The BIC criterion in a Very Simple Structure analysis (Revelle and Rocklin 1979) indicated that three factors might be optimal. An ordinary least squares (OLS) factor analysis with promax rotation yielded three latent factors which together explained 45% of the variability in the data.

- The first factor (17%) loaded onto *Voice intelligibility* (loading = 0.96), *Tone colour* (0.71) *Naturalness* (0.64) and *Source distribution* (−0.52). It was labelled *Colouration*.
- The second factor (16%), labelled *Timing*, loaded onto *Sequence of Events* (0.92), *Echoes* (0.75), *Artefacts* (0.70) and *Distance* (−0.57).
- The third factor (12%) loaded onto *Presence* (0.73), *Depth* (0.66) and *Immersivity* (0.55), and was labelled *Presence*. In each latent factor, the underlying scales are listed in order of strength.

We then investigated correlations between *Degree of Liking* on the one hand, and each of the three latent factors on the other. All were significant at the alpha = 0.05 level: *Colouration* (Pearson's $r = 0.42$, $p = 0.042$), *Timing* ($r = 0.43$, $p = 0.038$) and *Presence* ($r = 0.70$, $p = 0.00013$). To explore the factors further, we created linear regression models using stepwise reduction. A model with two predictors, *Timing* and *Presence*, explained 58% (R-squared adjusted for number of predictors) of variability in *Degree of Liking*. Model assumptions were met, as indicated by adequate linearity of fit (from inspecting a QQplot), and equal variance and distribution normality of residuals (by inspecting a scatter plot, histogram and passing Shapiro-Wilks's test). Despite the limited size of our group of participants, these results from the spatial listening test indicate that it was possible to predict (from SAQI ratings) the most general and overall perceptual characteristic, the *Degree of Liking*, for the three configurations that we built and tested. They also indicate a potential simplification of the protocol, approximately halving the number of scales, which might be useful for any future research that needs to evaluate the overall liking for a spatial audio setup. Further research might extend this line of inquiry with more participants and other loudspeaker configurations. We are also keen to compare the performance of OAT, a small DIY system, against larger loudspeaker rigs with high-performing equipment. While we expect blind-listening tests to favour more expensive systems, it will be valuable to understand which aspects of OAT should be improved in the next iteration of the project, while keeping costs under strict control.

3. DISCUSSION

One observation we have made in our project is that focusing on Ambisonics decoding as the main objective for the module seems a sensible choice, because it is one of the most theoretically complex questions in Ambisonics, as well as the most necessary, and it touches on a wide range of relevant aspects – such as formats, orders and irregular arrays – that are therefore addressed with

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press

concrete evidence, thus solving a good number of potential problems that might arise in the design of a listening space. In a pedagogical perspective, learning how to design a decoder is useful in many circumstances, whether playing back recorded sound fields or encoded Ambisonics sound files, and whether designing a listening space in the context of a home studio, a gallery or a public space.

Regarding the benefits of a low-cost DIY approach, there is a balance to be struck between the quality of the output and the cost of the purchase, as well as the time and labour involved. Time and money do not have the same value for those who lack money but can invest time (and vice versa), and there are different contexts in which this is the case. In this project we have found a solution that provides a satisfactory sonic result at a low cost. However, there are many possible compromises that can be made based on different needs and budgets, and different users can conduct their own personal experiments to find what's most useful in their applications.

An important consideration that emerged from the last evaluation is that different spatial configurations performed differently depending on the sound material. This material dependency seems to be particularly relevant when comparing our hemispherical and circular configurations. While the cuboid (CUBE configuration) performed worse than the others, likely because the sweet spot size is noticeably smaller in this setup, the other two gave different results depending on the audio content. This leads us to the conclusion that spatial configurations cannot be defined as simply better or worse. While it is sensible to always lean towards solutions offering the highest possible resolution – especially in smaller systems – it is important to consider the sounds intended for diffusion and the overall purpose of the array when deciding what configuration to employ; for example, whether the installation is designed to host one visitor at the time or several, whether it is a soundscape composition exhibited in a gallery or a concert-style event, and so forth. At this point, we have tried OAT in setups for up to ~5 people in exhibition-style Ambisonics setups, as well as ad hoc spaces (e.g., Lindborg 2024). This shows why it is necessary to include such tests in a pedagogical pathway, as they can teach students to approach the topic in a non-dogmatic way and can create engagement within a lecture. We noticed this during the two workshops we held so far, where rearranging and calibrating an array, and designing a new decoder resulted in a practical way to develop hands-on experience, breaking the frontality of the presentation by directly involving the participants.

During the development of OAT, we came across the project of a research group in Austria called OTTOsonics (Mitterhuber 2022). Their project shares with ours the goal of democratising spatial audio, making it something that everyone can experience in a range of contexts and environments (whether formal or not), and was therefore very inspiring. On the other hand, their technological approach leans more towards the idea of designing and assembling loudspeakers using 3D printing, as well as laser cutting aluminium to make cases for amplifiers and power supplies. This results in a more professional product, but also requires access to tools and facilities that are not yet widely available. This was a reason for us to continue with an approach based on sourcing and evaluating low-cost consumer grade resources that our end user could easily find in any electronics shop. Furthermore, it would be feasible to augment OAT with many of the OTTOsonics templates and resources available on their website. We believe that the main advantage of OAT presents minimal economical and technical obstacles and offers students and independent sound artists an entry point to exploring spatial audio. OAT allows user to develop their knowledge and skills with Ambisonics as well as to other formats of multichannel sound.

Since the design, construction and evaluation described in this text, we have held two workshops: a one-day event at the International Community for Auditory Display (Lindborg and Pisano 2023b) conference at Linköping University, Sweden, and a three-day event at SoundLab (Pisano and Lindborg 2024) and a presentation at UbiMus Conference (Lindborg et al. 2024). Demonstrating the system to people, explaining and discussing OAT, and of course listening to sound in different configurations, will provide us with invaluable information for the pedagogical part of the project. Current plans include a longer workshop for a group of local Hong Kong sound artists and continued pedagogical integration with courses at the School of Creative Media. Students working in small groups will pursue individual research or creative projects, such as electroacoustic music composition, auditory perception evaluations, or VR/AR application. Future potential developments could include the integration of Ambisonics and Wave Field Synthesis, making cutting-edge spatial audio technology affordable and accessible.

Going forward, we expect to address the remaining aspects of the project in two distinct phases. In the first phase, we will create a short manual for composers and sound designers interested in spatial audio. The process of collecting articles, organizing information and developing a substantive yet accessible discourse is time-consuming and iterative, but we are confident we can complete it in a relatively short period of time. The second phase will test the components in context. With adjustments, we can organize different formats, including semester courses, several-day workshops and short classes. We will test the material in various settings in order to gather as much feedback as possible and to improve the quality of the output. The results will be summarised in a review article.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/S1355771825000044.

Acknowledgements. This research project was funded by Teaching Start-up Grant #6000788 and ethical review of test procedures covered by Start-up Grant #7200671 from City University of Hong Kong. The authors would like to thank: Dirk Stromberg, LaSalle College of the Arts, Singapore, for valuable advice on configuring Jack and other software on Linux; Frank Ekeberg for his suggestion on using AmbDec; SCM undergraduate student Cyrus Leung for assistance with acquiring equipment; Ernst Van der Loo and Mariam Gviniashvili for recorded voice materials; and the SoundTeam graduate students for participating in listening tests.

Materials relative to the development of the OAT can be found at these links:

- The GitHub repository of the Open Ambisonics Toolkit: https://github.com/SonoSofisms/OAT/tree/main.
- The website of SCM's SoundLab: http://soundlab.scm.cityu.edu.hk/projects/open-ambisonics-toolkit/.

References

Adriaensen, F. 2009. Ambdec Manual. https://kokkinizita.linuxaudio.org/linuxaudio/downloads/ambdec-manual.pdf (accessed 14 September 2024).

Alsaproject. 2019. ALSA SoundCard Matrix. www.alsa-project.org/wiki/Main_Page (accessed 14 September 2024).

Ambrose, G. 2018. Sustainable Loudspeaker Design: Building a Working Loudspeaker from Found Objects and Waste. *Artefact*. https://research.brighton.ac.uk/en/publications/sustainable-loudspeaker-design-building-a-working-loudspeaker-fro.

ATK Community. 2016. ATK for SuperCollider. www.ambisonictoolkit.net/documentation/supercollider/ (accessed 14 September 2024).

Battier, M. and Fields, K. (eds.) 2020. *Electroacoustic Music in East Asia*. New York: Routledge.

Blue Ripple Sound. 2019. 3D Audio. www.blueripplesound.com/3d_audio (accessed 14 September 2024).
British Standards Institution. 2007. Acoustics – Definitions of Basic, *BSi ISO/TR 25417:2007*.
British Standards Institution. 2009. Acoustics – Measurements of Room Acoustic Parameters, *BSi ISO 3382-1:2009*.
Carneiro, M. 2016. Multichannel Audio Interface for Linux. https://discourse.ardour.org/t/multichannel-audio-interface-for-linux/89706 (accessed 14 September 2024).
Daniel, J. 2003. Spatial Sound Encoding Including Near Field Effect: Introducing Distance Coding Filters and a Viable, New Ambisonic Format. *23rd AES Conference*, Copenhagen, 13.
Debian Wiki. 2010. Pulseaudio. wiki.debian.org/PulseAudio/ (accessed 14 September 2024).
Guezenoc, C. and Seguier, R. 2020. Hrtf Individualization: A Survey. *arXiv preprint* arXiv: 2003.06183. https://doi.org/10.48550/arXiv.2003.06183.
Heller, A. 2014. Ambisonics Decoder Toolbox. *Bitbucket.org*. https://bitbucket.org/ambidecodertoolbox/adt/src/master/ (accessed 14 September 2024).
Heller, A. J. and Benjamin, E. M. 2014. The Ambisonic Decoder Toolbox: Extensions for Partial-Coverage Loudspeaker Arrays. *Linux Audio Conference*, ZKM/Center for Art and Media, 3 May, 1–9.
Heller, A., Lee, R. and Benjamin, E. 2008. Is My Decoder Ambisonic? *Audio Engineering Society Convention 125*. Audio Engineering Society.
Horsburgh, A. J., McAlpine, K. B. and Clark, D. F. 2011. A Perspective on the Adoption of Ambisonics for Games. Paper presented at the *Audio Engineering Society: 41st International Conference: Audio for Games*, London.
IEM. 2017. IEM Plug-in Suite. *plugins.iem.at*. https://plugins.iem.at (accessed 14 September 2024).
Ikeshiro, R., Charrieras, D. and Lindborg, P. M. 2022. Charting the Scene (s) of Sonic Arts in Hong Kong. *Organised Sound 27*(3): 369–78.
Jackaudio. 2006. JACK Audio Connection Kit. https://jackaudio.org/ (accessed 14 September 2024).
Jiang, X. 2017. Fundamentals of Audio Class D Amplifier Design: A Review of Schemes and Architectures.*IEEE Solid-State Circuits Magazine 9*, 14–25.
Keller, D., Lazzarini, V. and Pimenta, M. S. (eds.) 2014. *Ubiquitous Music*. Berlin: Springer.
Kessling, P. and Görne, T. 2018. Studio for Immersive Media Research and Production: Immersive Audio Lab at Haw Hamburg. *Audio Engineering Society 145th Convention*, New York.
Kjeldgaard, M. 2022. Notes for Setting Up a Raspberry Pi 4 for Audio Work. *madskjeldgaard.dk*. https://madskjeldgaard.dk/posts/raspi4-notes/ (accessed 14 September 2024).
Kronlachner, M. 2011. AmbiX – Ambisonics Plugin Suite. *matthiaskronlachner.-com* www.matthiaskronlachner.com/?p=2015 (accessed 14 September 2024).
Lähdeoja, O. 2016. Composing for an Orchestra of Sonic Objects: The Shakeousmonium. *Proc ICMC*, University of Michigan Library.
Lim, M. J. Y. and Lindborg, P. M. 2013. How Much Does Quality Cost? Listening to Music with Earphones on Buses and Trains. *Proc 3rd Int. Conf. on Music Emotion* (ICME3), Jyväskylä, Finland, 1–6.
Lindau, A., Erbes, V., Lepa, S., Maempel, H.-J., Brinkman, F. and Weinzierl, S. 2014. A Spatial Audio Quality Inventory (SAQI). *Acta Acustica united w Acustica 100*(5): 984–94.
Lindborg, P. M. 2015. *LW24*. Frontal Auditory Display with 21 Channels Data Sonification of Four Decades of Predicted Weather in South-East Asia (2015-55). Custom-built loudspeakers and large sculptural grid structure. Commissioned by the National Gallery of Singapore and ArtsFission for the inaugural exhibition.
Lindborg, P. M. 2021a. Feeling Loki's Pain: Designing and Evaluating a DIY 3D Auditory Display. *Proc ICAD*, Virtual Conference. https://doi.org/10.21785/icad2021.004
Lindborg, P. M. 2021b. Loki's Pain, *Proc NIME*, Shanghai, China.
Lindborg, P. M. 2024. *Stairway to Helheim*. Eight-Channel Site-Specific, Physical Loudspeaker Installation: A Sonification of 138 years of Weather Data from Hong Kong. Juried selection for *International Computer Music Conference*, Hanyang University Museum, Seoul, Korea, 7–13 July.
Lindborg, P. M. and Pisano, G. 2023. Workshop with the Open Ambisonics Toolkit. *The 28th International Conference on Auditory Display (ICAD 2023): Sonification for the Masses*, Norrköping, Sweden.
Lindborg P. M., Pisano, G. and Yu, K. 2024. Exploring the Open Ambisonics Toolkit: Workshop. *Proceedings of 14th International Symposium of Ubiquitous Music (UbiMus)*, University of Saint Joseph, Macao, China.
Linuxaudio. 2013. Linux Hardware Support. *linuxaudio.org*. wiki.linuxaudio.org/wiki/hardware (accessed 14 September 2024.)
Lyon, E., Caulkins, T., Blount, D., Ico Bukvic, I., Nichols, C., Roan, M. and Upthegrove, T. 2016. Genesis of the Cube: The Design and Deployment of an hdla-Based Performance and Research Facility. *Computer Music Journal 40*(4): 62–78.
Malham, D. G. 2019. The Early Years of Ambisonics at York. Paper presented at the *Audio Engineering Society Conference: 2019 AES International Conference on Immersive and Interactive Audio*.
McCormack, L. and Politis, A. 2019. SPARTA and COMPASS: Real-time Implementations of Linear and parametric Spatial Audio Reproduction and processing Methods. *Audio Engineering Society Conference: 2019 AES International Conference on Immersive and Interactive Audio*. Audio Engineering Society.
Mitterhuber, M. 2022. OTTOSonics. *alterbauhof.at*. www.alterbauhof.at/ottosonics/ (accessed 14 September 2024).
Murillo, D., Fazi, F. and Shin, M. 2014. Evaluation of Ambisonics Decoding Methods with Experimental Measurements. *Proc EAA Joint Symposium on Auralization and Ambisonics*, Berlin, Germany, 33–40.
Musil, T. 2000. iem_ambi Pd-objects for Ambisonics Sound Spatialization. https://git.iem.at/pd/iem_ambi (accessed 14 September 2024).
Nachbar, C., Zotter, F., Deleflie, E. and Sontacchi, A. 2011. Ambix-a Suggested Ambisonics Format. *Proc Ambisonics Symposium*, Lexington.
Pisano, G. and Lindborg, P. M. 2023a. Introducing the Open Ambisonics Toolkit. *Proc ICAD*, Linköping, Sweden.
Pisano, G. and Lindborg, P. M. 2023b. Workshop with the Open Ambisonics Toolkit. https://soundlab.scm.cityu.edu.hk/2023/08/07/workshop-with-open-Ambisonics-toolkit/ (accessed 14 September 2024).
Poletti, M. A. 2000. A Unified Theory of Horizontal Holographic Sound Systems. *J. Audio Eng. Soc. 48*: 1155–82.
Poltavsky, S. and Nadzharov, A. 2021. PureData CEAMMC Web Documentation, Category: spat. https://ceammc.github.io/pd-help/help-en/category_spat.html (accessed 14 September 2024).
Prior, N. 2018. New Amateurs Revisited: Popular Music, Digital Technology, and the Fate of Cultural Production. In L. Grindstaff, M. L. Ming-Cheng and J. R. Hall (eds.) *Routledge Handbook of Cultural Sociology*. Abingdon: Routledge, 340–8.
Ressi, C. 2023. vstplugin~ for ARM (Raspberry Pi). https://git.iem.at/pd/vstplugin/-/issues/172 (accessed 14 September 2024).
Revelle, W. and Rocklin, T. 1979. Very Simple Structure: An Alternative Procedure for Estimating the Optimal Number of Interpretable Factors. *Multivariate Behavioral Research*, *14*: 403–14. https://personality-project.org/revelle/publications/vss.pdf (accessed 14 September 2024).
Sèdes, A., Guillot, P. and Paris, E. 2014. The Hoa Library, Review and Prospects. *Proc International Computer Music Conference Joint with Sound and Music Computing*, Athens, Greece.
Studio Six Digital. 2022. Home Page. https://studiosixdigital.com/ (accessed 14 September 2024).
Zmöelnig, J. M. 2019. vstplugin~. https://git.iem.at/pd/vstplugin (accessed 14 September 2024).

Videography

Falch, A. 2022. Ambisonics Spatial Audio with Reaper and IEM. *YouTube*. www.youtube.com/watch?v=CejOR83X6lM (accessed 14 September 2024).
Hulme, D. 2018. Into the Soundfield – Michael Gerzon and Ambisonics at Oxford. *YouTube*. www.youtube.com/watch?v=X23hZNoSkUs (accessed 14 September 2024).

https://doi.org/10.1017/S1355771825000044 Published online by Cambridge University Press