PRODUCING IMMERSIVE AND INTERACTIVE AUDIO CONTENT WITH ORCHESTRATED PERSONAL DEVICES

Kristian Hentschel, Jon Francombe, Emma Young, Sonal Tandon, and Danial Haddadi

BBC Research and Development, United Kingdom

ABSTRACT

Access to multiple Internet-connected media devices such as phones, tablets, and smart speakers is now common in domestic environments. Device orchestration makes use of these devices to reproduce or augment media content, enabling novel interactive experiences. BBC R&D's experimental production tool, *Audio Orchestrator*, allows creators to rapidly produce demonstrators to explore this concept with audio content, building on previous work in media synchronisation and object-based audio. In this report, the tool is outlined and evaluated. The production process for four publicly released case study trial productions (with music, radio drama, and sport content) is described and feedback evaluated.

1. INTRODUCTION

Device orchestration is the concept of using multiple synchronised devices to augment or reproduce a media experience. Cieciura et al. (1) found that UK households commonly have access to multiple Internet-connected media devices. Audience members are already used to interacting with them or even consuming different media from multiple sources at the same time (by 'casting' content, linking up wireless speakers for surround sound or multi-room playback, or interacting with social contacts while watching TV). However, beyond early experiments with 'second screen' applications, little use is made of using these powerful devices for facilitating coherent, immersive, and interactive object-based media experiences at home. Orchestration builds on previous ideas for producing and consuming object-based audio and synchronised multi-screen content, offering vast opportunities for distinctive new content experiences, as well as more interactive, personalised, and accessible ways of consuming legacy content.

Francombe et al. (2) demonstrated audio orchestration through a trial production, and Francombe and Hentschel (3) investigated use cases in co-creation workshops. To enable rapid exploration of such use cases, it was necessary to develop a production tool that simplified the process of producing orchestrated audio experiences. This led to the design and development of *Audio Orchestrator*, which has been used to create several public trials with various partners, as well as experimental prototypes made by independent members of the creative community.

In this paper, the production tool and workflow are introduced (Section 2). The tool was evaluated in various ways, including analysing requests for access from the creative community, a survey, and semi-structured interviews (Section 3). Four production case studies are described and evaluated in Section 4. The work is discussed in Section 5.

2. PRODUCTION TOOL

Audio Orchestrator is a tool for building and previewing orchestrated audio experiences from existing audio files. It is available on request from BBC Connected Studio MakerBox, our platform for providing access to production tools and growing a community of practice.

Audio Orchestrator was designed to simplify the process of creating orchestrated audio experiences by giving producers creative control of how their content should adapt to variable multi-device setups, and outputting a configured template application handling pairing, synchronisation, and audio playback. After authoring audio files in a digital audio workstation (DAW), these are imported as audio objects and annotated with metadata (Figure 1). These, together with information about the connected devices, are processed by an algorithm to allocate objects to devices in the prototype application. The allocation algorithm is re-run every time a device joins or drops out, or a control value is changed.

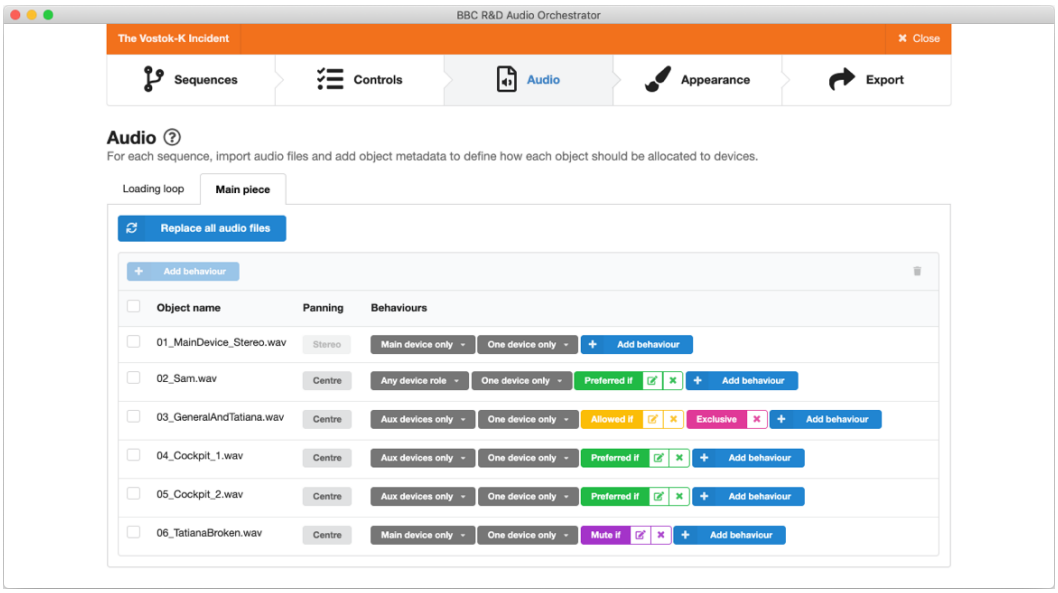

Figure 1 – Screenshot of the *Audio Orchestrator* interface for authoring behaviours

The workflow is structured around five core capabilities, accessed on pages of the tool.

1. **Sequences:** creates the structure of the piece, by linking independent content sections together and setting them to loop, automatically transition to the next sequence, or ask the listener to make a navigation choice.
2. **Controls:** defines input elements (such as multiple-choice questions or sliders for numerical values), for listeners to provide information on each connected device. This information forms part of the device metadata used by the allocation algorithm.
3. **Audio:** shows the imported audio files for each sequence, and lets producers add metadata about the desired allocation behaviour to each object.
4. **Appearance:** provides customisation options including text, images, and colours for the exported prototype, as well as audio settings such as dynamic range compression.
5. **Export**: controls a preview server for immediately trying out the experience with real devices on the local network, and a final web application bundle can be exported.

The object behaviours can be used to: directly link audio object allocations to choices made by the listener (e.g., their preferred language); define rules for constrained random assignments based on other objects or information about the connected devices such as

their joining order or form factor; define if and when objects can move between devices in response to changes; and modify the gain of an object under certain conditions. Objects can be allocated to one or multiple devices under this scheme. Additionally, images and lighting effects at defined times during a sequence may be allocated to devices in the same way.

The exported prototype web application is modelled after familiar music playback apps and provides a user interface for connecting devices (by scanning a QR code or typing a pairing code), controlling playback, making choices to move to the next sequence, and calibrating the latency of connected devices. The session is started on a *main device*; any additional connected devices are referred to as *auxiliary devices*. Synchronisation and communication between devices in a session is facilitated by an open-source cloud-based synchronisation service[1] using the framework described by Hentschel and Francombe (4). Screenshots of the prototype application are shown with the case studies, in Figures 4–7.

3. EVALUATING THE PRODUCTION TOOL

Audio Orchestrator was released on MakerBox in July 2020. The public release was used to learn about the benefits, use cases, and challenges of audio device orchestration, as well as about the usability of the tool. Various methods were used for evaluation. When requesting access, potential users fill in a short form describing their background and intended use of the tool. These responses (309 when this analysis was conducted in October 2020) were analysed using a thematic analysis process to determine the professional backgrounds of users interested in the tool and what kind of projects they were interested in using it for. The tool was then further evaluated through a survey and interviews with four individual producers after they had used it for a few weeks.

Analysis of MakerBox tool requests

The information provided in the access requests by prospective users of *Audio Orchestrator* (309 by October 2020) was grouped using a thematic analysis process with emergent themes. Responses were coded and clustered into a three-level hierarchy, separately for user backgrounds and use cases, in an iterative process. A single response might be reported in multiple themes. The use case themes are shown in Figure 2.

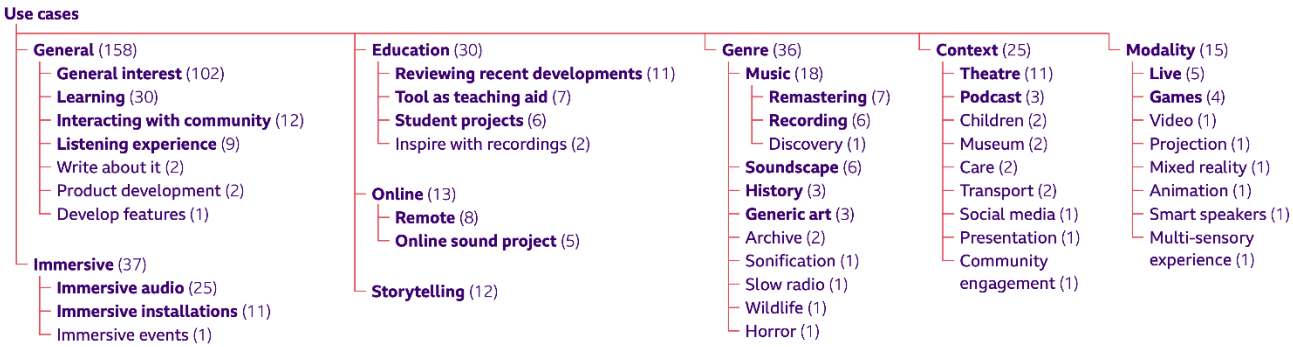

Figure 2 – Hierarchy of use case themes from tool access requests

By far the most frequent theme for the use cases was expressing a **general interest** without elaborating on specific uses. This was followed by **learning**, such as keeping up with developments in audio technology. The **education** cluster included using the tool in

[1] https://github.com/2-IMMERSE/cloud-sync

courses **reviewing recent developments** in audio technology, using it as a **teaching aid** to illustrate technical or musical concepts, and **student projects**. Most creative production uses fit into the **immersive** theme, but many individual **contexts**, combinations with other **modalities**, and specific **genres** were also mentioned by a small number of users each.

A variety of user backgrounds were found, primarily from the **creative industries** (audio production, music, publishing, theatre), **education** (teachers, students), or **research** (academic and commercial). There were also many **hobbyists** expressing a non-professional interest in the tool, a few **BBC** staff members, participants in a creative talent development scheme, and some individual responses from completely different fields.

User experience

A short questionnaire was sent to all MakerBox users who had requested access to *Audio Orchestrator* to collect qualitative and quantitative data about their experience of using the tool. Eleven responses were received, but one respondent was not able to use the tool (because of the lack of a Windows-compatible version at the time) so did not answer most of the questions. Questions covered use of *Audio Orchestrator*, the prototype application, and overall feedback.

The responses were mainly very positive. The application received a median rating of 5 out of 5. The features of *Audio Orchestrator* were generally well understood, but the hardest features to use and understand were the controls and various behaviour types (see Figure 3). When asked about missing features, respondents suggested that more customisation of the prototype application, easier discovery and pairing of devices, video and image authoring/playback, and loudness calibration would all be valuable. One participant suggested that the technology is useful but that an entire redesign of the interface is required. All features of the prototype application were considered easy or very easy to use, and the general look and feel of the prototype application was given a median rating of 4 out of 5. Feature requests for the prototype application included reducing audio glitches, improving the layout, and more components for navigating the experience.

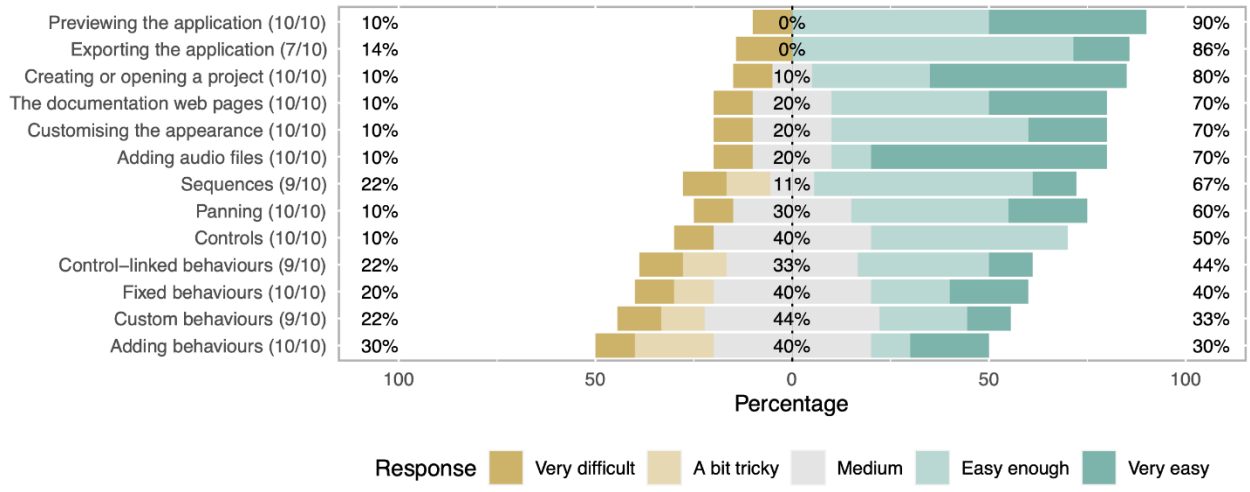

Figure 3 – Survey responses for ease of use of various *Audio Orchestrator* features

In addition to the survey, semi-structured interviews based around a topic guide were conducted with four audio professionals who had used the tool. Each interview was led by one researcher, whilst a second took notes. Full summaries of the interviews are beyond the scope of this paper, but there were some commonalities in the responses. The

feedback received was generally positive, with users finding it easy to install the tool and make a simple project. The tool was often used to repurpose existing works (such as multi-channel installations or object-based audio demonstrators). As in the survey, it was notable that the behaviours were hard to use and understand. Audio and/or visual monitoring were requested to make it easier to check that behaviours were having the desired effect. Some simple concepts, such as allocating a group of objects to as many different devices as possible, were hard to achieve with the existing behaviours. Conceptually, participants tended to think about the connected devices as if they were a fixed multi-loudspeaker installation, allocating one channel to one device. It took some familiarity with the tool to consider how controls and behaviours could be used to enable interactivity and/or to ensure the produced experience was adaptable.

Summary

Analysis of the MakerBox requests suggested that there are many uses of orchestration across the creative industries, including various forms of arts and media as well as in education. Many requesters were simply interested in finding out about some new technology and staying up to date with developments in the field, but others had specific use cases in mind around education and immersive experiences in a range of contexts, modalities, and genres. Results from a small survey suggest that *Audio Orchestrator* is reasonably easy to use, but there are challenges in the more intricate use of behaviours to determine how audio objects are allocated to devices.

4. PRODUCTION CASE STUDIES

Following the release of *Audio Orchestrator*, four orchestrated trials were produced and released to the public on BBC Taster. The trials were designed to investigate different orchestration use cases (radio drama, music, and sport) from an audience perspective and to explore how orchestrated production can fit with existing workflows. The case studies in this section describe the production process and give some insight into the challenges faced by creative teams when producing orchestrated content.

Decameron Nights (immersive audio drama)

Decameron Nights is a radio drama consisting of a series of audio stories that was commissioned through Culture in Quarantine and BBC Arts. The episode *I'm Alright Jack* was refactored as an orchestrated experience in collaboration with performance group 1927. The release of the orchestrated version (Figure 4) on BBC Taster was timed to coincide with the broadcast of the original stereo version on BBC Radio 3 in August 2020.

Figure 4 – Screenshots from 1927's *Decameron Nights*

Production process

Remote production support was given to 1927's creative team to introduce *Audio Orchestrator's* core functions and to help them realise their vision for the orchestrated piece. Sequences were defined to handle three distinct phases within the experience (*Onboarding*, *Main Content*, and *Outro*); these had varying requirements for the allocation of objects and sequence settings. The *Main Content* was the existing content produced for the radio broadcast. Additional content was produced to bookend the experience.

The narrator recorded a voice-over for the *Onboarding* section, with script and delivery to match the aesthetic of the piece, which resulted in a set of connection instructions that flowed seamlessly with the main content. During this section, no explicit instructions were given with respect to auxiliary device placement. Instead, suggestions were made to convey the notion that devices should be spread around the listener (*"Place it somewhere in the room / On Grandma's antique coffee table perhaps? / Or on Grandma's lap / On the slanted shelf behind your head / In the dog's bed"*). The *Main Content* was made up of 58 audio objects containing narration, music, and sound effects for each story within the episode. The producers created a lean back experience, emulating the feeling of *"settling down to listen to an intimate storytelling session"* so when a listener finished connecting their available auxiliary devices and pressed *Begin*, no further interaction was required. The *Outro* sequence wrapped up the experience and provided a themed call to action to listeners, where they were encouraged to complete a feedback survey.

Production challenges

Support was given in authoring more complex behaviours to enable creative ideas for the allocation of audio objects between the main and auxiliary devices. For example, at certain points in the story *Fat Cat*, the team described how their 'cat munch' sound should *"jump between all the available connected devices"*. To achieve this effect, working outside of *Audio Orchestrator* in a DAW, the original mono stem for that sound effect was duplicated across multiple tracks and edited to occur at different points in time. These tracks were then exported as stems and imported into *Audio Orchestrator* as eight separate objects. A custom behaviour was authored to ensure that no two sequential sounds were heard from the same device, achieving an overall effect of both temporal and spatial separation, through use of the following *Preferred if* conditions:

Preferred if: [Objects allocated to same device] [not] [any of] [adjacent object numbers]

When testing the experience (e.g., previewing changes during production and testing user scenarios with various device setups) the production team found that connecting multiple auxiliary devices to assess each change was significantly more laborious and time-consuming than simply pressing play on a DAW to preview a traditional production. This feedback informed subsequent work undertaken to explore how monitoring functionality could be implemented to aid the production process.

Additional development requirements

In addition to orchestrating the audio elements of this experience, the team wanted to include several still and moving images that would be displayed on the main device screen at the relevant points in the story. Functionality for timed image authoring was not supported in the first version of *Audio Orchestrator;* therefore, some development support was required to separately add images to the exported application.

Pick A Part (interactive classical music)

Pick A Part (Figure 5) was our first orchestrated piece with music as the primary content, and was produced in collaboration with the BBC Philharmonic. With the orchestra's musicians unable to play together due to Covid-19 restrictions, they began recording their individual parts separately, a production technique well-suited to the object-based audio requirements of orchestration. The experience allowed the listener to choose which musical part to hear on each connected device. They could connect multiple devices and position the instruments around the room, to simulate sitting in the middle of an ensemble.

Production process

Individual stems for four musical pieces were supplied by the BBC Philharmonic and in *Audio Orchestrator*, separate sequences were defined to accommodate each piece (with their associated instrument parts as individual audio objects). User controls specific to each sequence were added to enable the listener to switch between the different instruments on each connected device. Radio button controls were used so that only one instrument could be selected at any time on each device. This was set up by adding three behaviours to each object: (i) an object was *Preferred if* the control was set to the matching instrument; (ii) an object was *Prohibited if* the control was not set to the matching instrument on auxiliary devices (to force it back into the main device); and (iii) a *Change management* behaviour was used to allow objects to move to devices in which they were more preferred, but otherwise stay where they were previously allocated.

As in *Decameron Nights,* bespoke onboarding content was produced, designed to introduce the experience; provide instructions for connecting devices; and encourage users to connect multiple devices before navigating to the main content.

Additional development requirements

Music requires very accurate synchronisation between objects. The synchronisation framework is accurate to within 10–20 ms, but there is additional output latency caused by software and hardware on the reproduction device that cannot be corrected by the cloud-based synchronisation system. This varies greatly between devices, sometimes reaching ~500 ms. To mitigate this, a manual calibration stage was developed, enabling listeners to correct the time-alignment for their devices. The calibration system works by playing the same speech content from the main device and an auxiliary device, and asking the user to adjust a slider until the devices are synchronised. Calibration mode can be enabled as an option in *Audio Orchestrator*.

Additional features implemented to extend the application exported from *Audio Orchestrator* include: (i) displaying an image of the selected instrument on each device; (ii) displaying a representation of all instrument selections on the main device; (iii) ensuring that no two devices could select the same instrument at the same time; and (iv) showing a list of available pieces with titles and images.

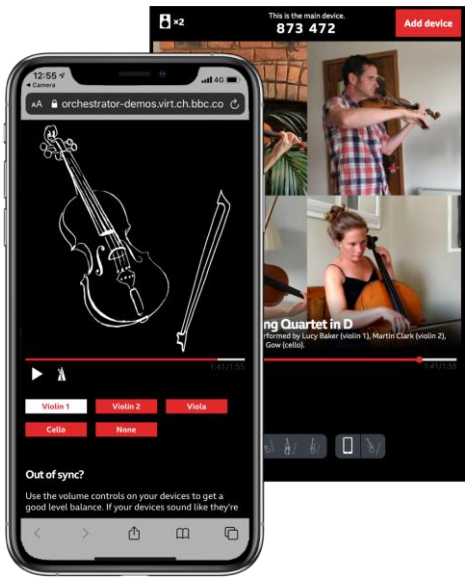

Figure 5 – Screenshots from *Pick A Part*

Monster (immersive horror drama)

Monster was an immersive audio drama in the horror genre that was conceived and developed specifically for orchestrated delivery. It was produced in collaboration with BBC Wales Innovation and BBC Writersroom Wales. *Monster* was initially planned as a physical installation, but restrictions arising from the Covid-19 pandemic prompted the shift to developing an online at-home experience. The pilot was launched on BBC Taster in the run up to Halloween 2020. As a horror story, *Monster* was designed to be experienced in the dark and incorporated timed lighting effects and images that were displayed on the main and auxiliary devices. Screenshots are shown in Figure 6.

Production process

Dialogue by the voice actors was recorded remotely and audio production followed the traditional radio drama workflow, with sound design and the final mix completed in a DAW. The multi-track audio was exported as 67 separate objects and brought into *Audio Orchestrator*. Orchestration of the audio objects and authoring of lighting effects and image sequences were completed in parallel by different producers in separate *Audio Orchestrator* project files, which were then combined by manually editing the project files.

Monster diverged significantly from the other orchestrated radio drama, *Decameron Nights*, in two ways. Firstly, listeners were required to connect at least two auxiliary devices and were given explicit instructions on where to position these first two connected devices. Secondly, the requirements for synchronised images differed: as well as displaying sequences of images on the main device, separate timed images and lighting effects were displayed on each auxiliary device independently. Lighting effects (flashing or pulsing) were implemented as colour fills overlaying the image, with animated opacity.

Production challenges

The concept of orchestration can be difficult to grasp with most writers and producers familiar with developing content for non-interactive, linear formats. A haunted house themed technical demo was built to showcase some of the features of orchestration. This proved helpful in demonstrating to the production team how one-off triggered sound effects and an interactive phone call could work.

There was concern that users might fail to ensure their connected devices were at a sufficient level to experience a balanced mix between the different audio objects. There is a risk that if a device's output level is too low, or the device is in silent mode, some crucial narrative elements, sound effects, or ambience might be missing or less prominent than the audio producer intended. In *Monster* (and the other trials), this risk was addressed by (i) applying dynamic range compression to boost the volume of sounds sent to auxiliary devices, and (ii) providing explicit voiced instructions to ensure device volumes were raised and not set to silent mode. Results from the evaluation suggest that unbalanced volumes is the most commonly reported problem across all the experiences.

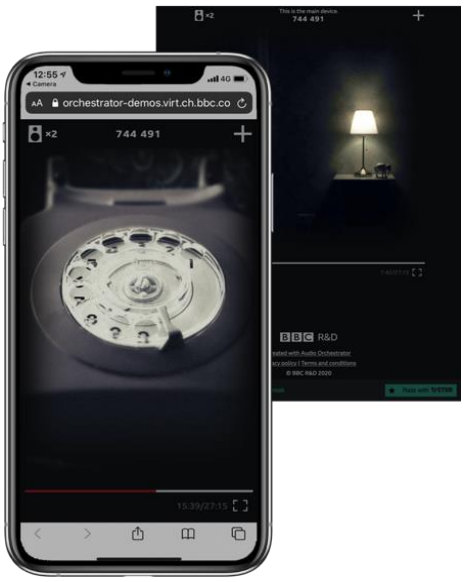

Figure 6 – Screenshots from *Monster*

Six Nations rugby (augmented sports coverage)

An orchestrated version of rugby highlights coverage was created in collaboration with BBC Wales Innovation and the BBC Wales Sport programme *Scrum V,* using content from a 2021 *Six Nations* match. On the main device, viewers could toggle commentary on and off. On auxiliary devices, viewers could switch between additional crowd noise or the referee's microphone feed. The two different interfaces are shown in Figure 7.

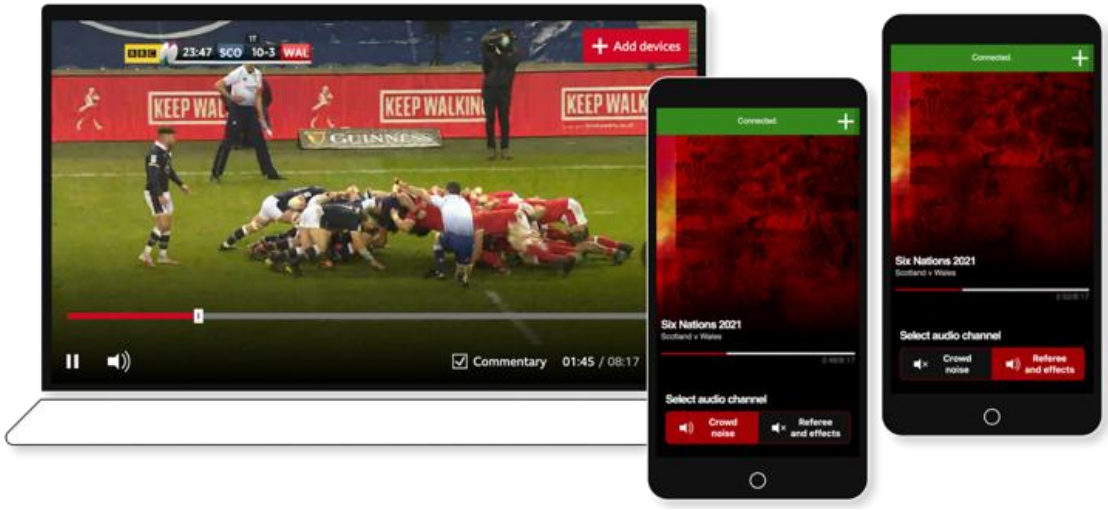

Figure 7 – Screenshots from *Six Nations* rugby

This trial was delivered in a different way to those described above, as the main device needed to play video synchronised to the auxiliary device audio playback. To achieve this, the main device ran the BBC's Standard Media Player (SMP) with a synchronisation plugin based on that developed for the experimental group viewing service *BBC Together*[2]. The plugin was extended to interface with the orchestration library and to handle multi-channel audio. Auxiliary device content was authored in *Audio Orchestrator*.

Production process

An introductory video was produced as a lead-in to the highlights, featuring the *Scrum V* title sequence and music with voiced instructions for connecting devices supplied by the programme's presenter. These video and audio stems were combined with the highlights content and accommodated within a single *Audio Orchestrator* sequence.

For the auxiliary devices, a single checkbox control was defined and behaviours were authored allowing viewers to switch between the crowd noise and referee microphone objects. The referee feed was only allowed on a single device at any time if multiple auxiliary devices were connected. Extra development was used in the prototype application to change the appearance of the selection control, and to update its status if the referee microphone was selected on another device. Rather than switching between allocated objects, the control was also modified so that both objects played simultaneously, and changing the control set the gain of the object that was not required to –60 dB. The *Gain adjustment if* behaviour developed to achieve this has subsequently been added to a release of *Audio Orchestrator*. This allows near instantaneous switching

[2] https://www.bbc.co.uk/taster/pilots/bbc-together

between objects (as opposed to switching object allocations, where there is a small delay whilst the new object is buffered) at a cost of increased delivery bandwidth.

Production challenges

As current production workflows for sport have been refined for traditional broadcasting and are not specifically geared towards object-based production, there were challenges in assembling content. The provided stems were not well labelled or aligned. The highlights package did not have all the required stems, so it was necessary to rebuild the audio files from the full match content (using a hybrid process of manually identifying the necessary clips and automatically selecting time-aligned segments using cross-correlation).

Taster evaluation of trial productions

The BBC Taster platform facilitates evaluation of pilot experiences through a short questionnaire that viewers are prompted to complete following the experience. For the four experiences described above, similar questionnaires were used. The results are summarised in Table 1. There were some differences in the exact wording of the questions and possible answers; these have been combined for the comparison below. Percentages do not always sum to 100%, as respondents could skip questions or respond “not sure”.

	Decameron	***Pick A Part***	***Monster***	***Six Nations***
Unique tries (days available)	2586 (93)	1969 (219)	11700 (171)	502 (50)
Mean / mode star rating out of 5 (number of ratings)	4.38 / 5 (164)	3.93 / 5 (72)	3.37 / 5 (274)	3.64 / 5 (22)
Would definitely or would like to use this tech again	90%	74%	74%	60%
Preferred orchestrated version / Preferred normal listening	25% / 3%	45% / 6%	65% / 4%	39% / 11%
Percentage using just the main device / +1 device / +2 or more	10% / 34% / 52%	14% / 36% / 37%	- / - / 76%	32% / 37% 26%
What was best about the experience? (Modal response)	Sound effects (43%)	Increased immersion (45%)	Devices as speakers (42%)	Referee mic choice (44%)
How could we improve it/what was worst? (Modal response)	Unbalanced volumes (28%)	Unbalanced volumes (17%), something else (17%)	Unbalanced volumes (22%)	Better sync (22%)

Table 1 – Summary of BBC Taster survey responses for all pilots

The experiences were viewed over 16,500 times in total, and the results are positive. Mean star ratings vary from 3.37 to 4.38, with mode ratings of 5/5 for all experiences. Between 60–90% of respondents would like to use orchestration technology again. The results suggest that connecting devices was common, with 64% or more of respondents across all experiences reporting that they connected one or more devices. It was rare for respondents to prefer normal listening to orchestrated listening (3-11%). When asked about the best part of the experiences, there was generally strong agreement within experiences (modes between 42% and 45%), but different selections for each of the experiences. There was less agreement within experiences about the worst aspects

(modes between 17% and 28%), but across the experiences unbalanced volumes between devices appears to be the main issue.

Summary

In the sections above, four orchestrated audio case studies that used *Audio Orchestrator* were described. They were produced in collaboration with internal and external partners, and all revealed various production challenges and requirements for bespoke development. Many of the features requested have been added to subsequent releases of *Audio Orchestrator*. Positive results from the evaluation suggest that there is audience benefit in such multi-device experiences.

5. DISCUSSION AND FUTURE WORK

A prototyping tool for creating synchronised multi-device audio experiences was described. The evaluation showed that creators can independently (or with limited support) make interesting new prototypes with the tool and generally responded positively to it. However, several areas are still hard to understand and there is room for workflow improvement. The audio object behaviours are particularly complex and can feel more like programming than audio editing or mixing. Producers find it hard to predict what will happen in certain situations without connecting devices and trying it out. Audio professionals are familiar with existing content delivery models and therefore think about designing content for fixed rather than variable loudspeaker layouts. They may consider a few different setups, but find it hard to consider a dynamic setup where loudspeakers of different types may join or drop out during the session. A monitoring system for *Audio Orchestrator*, providing visual and auditory previews of the object allocations for a range of presets or custom setups, is currently in development to address these issues by allowing users to experiment with different settings more easily.

Producers tend to use pre-existing content to test orchestrated audio. But when converting existing content (like *Decameron Nights* and the *Six Nations* rugby) into an orchestrated experience, it was still necessary to design and produce extra content to signpost the viewer through the experience (e.g., explaining the technology and device connection).

It is clear from the trial productions that producers see the potential of interactive multi-device experiences using audio but also other modalities. Being able to add timed images and lighting effects in *Monster* made for a much more immersive experience, and these features have since been added to the general release of *Audio Orchestrator*. The *Six Nations* pilot shows that orchestration works well with synchronised video content.

Other features where the pilots motivated developments that have subsequently been included in *Audio Orchestrator* are: support for stereo audio files; timed images and lighting effects; fading out of audio objects when they are removed from a device; a calibration mode to compensate for device output latency; and the option to allow play/pause control from auxiliary devices. Based on community feedback, a conditional gain adjustment behaviour and an operator for easily selecting groups of devices by their joining order have been added, and a Windows version of the tool released.

There is always a trade-off between what is possible in a general-purpose production tool, and custom software written for a particular production; for example, the interactive phone call and one-off triggered sound effects also considered for *Monster* production made for a compelling demonstration but could not easily be integrated with the tool. It is evident that to get the maximum benefits from orchestrated audio technology, it is beneficial to be able

to make bespoke modifications to the template prototype application; however, this is not currently supported in *Audio Orchestrator*.

The pilot productions were well regarded (as shown in the questionnaire responses), but being technology demonstrations working within the limits of the web browser, several technical issues remain. Firstly, the application has no control or knowledge of a device's loudspeaker type and volume settings, and therefore manual calibration of each device is required to achieve a balanced mix. Dynamic range compression can be added on auxiliary devices in an attempt to address this, and providing explicit voiced instructions was often used to check that device output volumes were matched. Secondly, there were occasional rendering artefacts with some types of audio content and performance issues where many streams had to be rendered simultaneously on one device. A method of manual latency compensation was offered to combat synchronisation issues. Ideally this device-specific offset would be detected automatically (but this value is not currently exposed by the browser APIs or operating systems).

Future work

Future work could include developing a live production and distribution system for live content, and reviewing the benefits of orchestration for accessibility needs such as personalised mixes and audio descriptions required by some users. Next generation audio codecs for object-based audio delivery might prove useful for overcoming some of the performance and quality issues, while device manufacturer support might be required to improve the calibration experience and to deliver orchestrated content on a wider range of consumer devices e.g., smart speakers and smart TVs.

ACKNOWLEDGEMENTS

Audio Orchestrator was developed in the BBC R&D Audio Team. We've received excellent support from many colleagues in other areas of R&D (Connected Studio MakerBox and Taster, the Product Engineering Team, and Communications), as well as our partners on the trial productions described above (1927, BBC Philharmonic, BBC Writersroom Wales, and BBC Wales Sport and Innovation). Particular thanks to Catherine Robinson for her work on *Monster* and the *Six Nations*. Finally, many thanks to those who have helped us evaluate the pilots and the production tool.

REFERENCES

1. Cieciura, C, Mason, R, Coleman, P, and Paradis, M, 2018. Survey of Media Device Ownership, Media Service Usage, and Group Media Consumption in UK Households. Audio Engineering Society Convention 145. eBrief 456.

2. Francombe, J, Woodcock, J, Hughes, RJ, Hentschel, K, Whitmore, E, and Churnside, A, 2018. Producing audio drama content for an array of orchestrated personal devices. Audio Engineering Society Convention 145. eBrief 461.

3. Hentschel, K, and Francombe, J, 2020. Exploring audio device orchestration in workshops with audio professionals. Audio Engineering Society Convention 148. eBrief 571.

4. Hentschel, K, and Francombe, J, 2019. Framework for web delivery of immersive audio experiences using device orchestration. TVX 2019. Adjunct Proceedings.