

Personalised Interactive Sonification of Musical Performance Data

KatieAnna E. Wolf · Rebecca Fiebrink

Received: date / Accepted: date

Abstract In this article, we describe methods and consequences for giving audience members interactive control over the real-time sonification of performer movement data in electronic music performance. We first briefly describe how to technically implement a musical performance in which each audience member can interactively construct and change their own individual sonification of performers' movements, heard through headphones on a personal WiFi-enabled device, while also maintaining delay-free synchronization between performer movements and sound. Then, we describe two studies we conducted in the context of live musical performances with this technology. These studies have allowed us to examine how providing audience members with the ability to interactively sonify performer actions impacted their experiences, including their perceptions of their own role and engagement with the performance. These studies also allowed us to explore how audience members with different levels of expertise with sonification and sound, and different motivations for interacting, could be supported and influenced by different sonification interfaces. This work contributes to a better understanding of how providing interactive control over sonification may alter listeners' experiences, of how to support everyday people in designing and using bespoke sonifications, and of new possibilities for musical performance and participation.

K. E. Wolf
MeasuringU
Denver, Colorado
E-mail: katieanna.wolf@gmail.com

R. Fiebrink
Department of Computing
Goldsmiths, University of London
E-mail: r.fiebrink@gold.ac.uk

Keywords interactive sonification · music performance · audience interaction · new interfaces for musical expression

1 Introduction

Hermann and Hunt define “interactive sonification” as “the use of sound within a tightly closed human-computer interface where the auditory signal provides information about data under analysis, or about the interaction itself, which is useful for refining the activity” [9]. They “argue that an interactive sonification system is a special kind of virtual musical instrument. It’s unusual in that its acoustic properties and behaviour depend on the data under investigation. Also, it’s played primarily to learn more about the data, rather than for musical expression.”

We propose that there is an underexplored role for interactive sonification within music performance itself, in which a data sonification might be “played” both for musical purposes (e.g., to create and explore expressive sounds) and also to be used and adapted to learn more about various data that may be captured within a performance (or even form the foundation for a performance). In any purely digital, electronic music performance, the music can be understood as being fundamentally driven by data that encodes performer actions and composer intentions—for instance, patterns encoded on a sequencer, or MIDI messages generated by control interfaces. Experimental digital musical instruments can be understood as carefully crafted “sonifications” of data from a wide variety of sources, often including sensors sensing performer gesture, as well as features extracted from real-time audio or even data aggregated from multiple human performers [23, 2]. Many

musical pieces have also featured sonification of real-world data (e.g., social media [3] and scientific [11] data).

In all of these cases, audience members might be interested in or otherwise benefit from deepening their understanding of the underlying data. For instance, deepening an awareness of the data generated by performers of new musical instruments may support an increased understanding of musical structure or performer intention. Furthermore, enabling audience members to interactively modify sonifications could foster a deeper engagement with the sound. Increasing understanding of musical structure, expression, and sound seems especially useful in new music performances in which the instruments and musical structures are unconventional and may be especially opaque to audiences unfamiliar with contemporary practices.

We therefore believe that the use of interactive sonification of musical performance data by audience members is worth exploring as a way to enrich audience members’ experiences of music performances. In this article, we explore mechanisms to provide audience members with *personalised* sonification interfaces. These support tight interaction loops in which audience members (1) dynamically exercise control over sonification of performer movement data and (2) perceive how their actions and performers’ actions together influence the performance sound. Such an approach requires a digital musical instrument (DMI) whose sound can be synthesised in distinct ways for each listener, while ideally still maintaining synchrony between the visual cues of the performer actions and the synthesised audio. Our team have recently designed a DMI that employs gesture prediction to allow latency-free synchronization between the performers’ actions and the sound synthesised on audience members’ personal computing devices [14]. We describe studies on the use of this DMI within two live musical performances. Drawing on log data from audience sonification interfaces as well as surveys of audience members, we examine how and why audience members interacted with the personalised sonification interfaces, and how different approaches to interactive control—exposed by different user interfaces—impacted their experience of the performance. This work leads to a deeper understanding of how interactive sonification can enable new approaches to musical performance and new types of engagement with sound, as well as a deeper understanding of how and why sonification interfaces can be designed for interactive control and exploration by people who are not sonification experts.

We begin by further contextualising our approach to interactive sonification of performance data, and by describing how this approach differs from other approaches to sonification, music composition, and digital musical

instrument design (Section 3). We then propose three research questions about user interaction and experience with personalised interactive performance sonification, and we describe the methodology used to address these questions in the context of two live performances (Section 4). In Sections 5 and 6, we present the results from each performance study, and we discuss additional insights that are suggested by the synthesis and comparison of the two studies in Section 7.

2 Background

In “parameter mapping” approaches to sonification, features of the data are mapped to acoustic attributes [9]. Human interactive control of such a system can include both controlling the data to be sonified and controlling the mappings or mapping-related parameters [13]. Whereas new digital musical instruments can be viewed as sonification systems in which performers interact by controlling the data (e.g., influencing the values of sensors using gestures), the approach we consider here additionally incorporates the ability for audience members to interactively alter the sonification mappings.

Interactive sonification approaches have been applied in a variety of domains, such as the creation of accessible interfaces for people with visual impairments [24], sports training [6], and medical diagnosis [22]. In any domain, giving sonification users control over aspects of the sonification can allow users to adapt the sonification for their own needs and interests. Ground and Hermann discuss the importance of knowing the purpose of the sound in a sonification and the listening mode that is required to achieve that purpose [8]. By giving users control of the aesthetic design of a sonification, we take advantage of their individualized listening modes (which may change over time) and their personal goals for what they want to hear. Furthermore, hearing the sonification change in response to users’ changes to the data and/or sonification puts the user in the centre of a control loop that helps them learn from their interactions and make changes to the sonification accordingly. In the work described in this article, we aim to provide all of these benefits in a musical performance context, by providing audience members with the ability to interactively manipulate parameter mapping sonifications throughout a performance.

A number of past musical performances have incorporated audience involvement and participation. In *Constellation*, Madhaven and Snyder presented a performance in which the audience’s mobile devices acted as speakers [18]. Their aim was to experiment with group music-making and create a “distributed-yet-individualized listening experience”. Many other perfor-

mances have incorporated mobile computing and network technology to enable audience participation (for additional work see [12,17,20]). *Mood Conductor* and *Open Symphony* put audience members in the role of “conductors” of a live musical composition where the client/server web infrastructure allows the audience and performers to exchange creative data [1]. *Tweetdreams* is a performance where the sound and visuals are influenced by audience members’ Tweets [3].

In these and most other DMI performances, the concert experience remains traditional in that all audience members hear the same sonic rendering of the performance, and individual audience members still have relatively small influence over the overall sound of the performance. Further, these performances have not given audience members much (or any) control over the instrumentation or synthesis.

We seek to explore a new approach to live musical performance in which every audience member can interactively influence the sounds of the performers on stage. In this approach, each audience member uses software running on a personal device (laptop, smartphone, tablet, etc.) to interactively control the parameters of a personalised sonification of performer activity. The personal devices use these parameters to render a real-time performance sonification, which audience members listen to through headphones connected to their devices.

In this approach, each audience member creates a personalised musical piece that is unique to *them*. Each person can thus explore the sound space to create a version of the performance that appeals the most to them individually, and this approach potentially affords each audience member new means to understand, engage with, and experience the performance, as they become an active rather than a passive participant.

3 Technologies to Enable Personalised Performance With Interactive Sonification

We can enable this new approach to performance using a digital musical instrument (DMI) that separates the gestural input of the performer (sensed using a gestural controller) from the generation of its sound via digital sound synthesis [19]. Jin et al. have proposed a new class of DMIs called “distributed, synchronized musical instruments” (DSMIs): these allow for sound generation to occur on multiple devices in different locations at once, while still maintaining the synchronization between the input gestures of the performers and the sound output [14]. DSMIs achieve this synchrony by predicting performer movements, such as percussive strikes meant to trigger particular sounds, before

the performers complete those movements. Information about these predicted performer movements can be sent (e.g., over the internet) to all sound generation devices in advance, so that the sound is synthesised at all locations in synchrony with the performer’s completion of the movement.

Our approach to performance personalisation, utilises DSMIs to allow for synchronized sounds to occur on multiple audience devices simultaneously. However, whereas typical DMIs and DSMIs use a single, fixed mapping strategy to define how sound should be controlled by performer movements, our approach allows individual listeners to modify the mapping strategy used for sound generation on their own device. While the creator of a new DMI must design the mapping strategy as part of the instrument, the creator of a new musical work using our approach must decide on which aspects of the sound mapping are fixed and which may be changed (and how) by the listeners.

We have implemented a software system that supports the design of new instruments for interactive personalised performance sonification. As described below, performers’ percussive striking gestures are sensed using the MalLo DSMI [14]. A web application running on audience members’ personal devices provides a user interface for personalised sonification control, and it also synthesizes the personalised audio stream. This system supports the implementation of different mapping strategies and exposure of different types of personalised sonification control. The mapping and control for the two musical performances we studied are described in Sections 5.1 and 6.1.

3.1 Sensing and Predicting Performer Gesture

The sensing and prediction component of our system uses an implementation of the original DSMI called MalLo. MalLo was originally created to provide a way for geographically distributed performers to collaboratively make music over the Internet [14]. It takes advantage of the predictability of percussion instruments [4] by tracking the head of a mallet/stick, fitting a quadratic regression to its height, and computing the zero-crossing time of the resulting parabolic curve to get the strike time prediction. Those predictions are continuously sent over a network to a receiver, each prediction more accurate than the previous as the mallet gets closer to striking. Once a prediction is delivered that is below a specified accuracy threshold, the note is scheduled to be sonified by the receiver. With this approach the strike sound will be heard at both the sending and receiving locations simultaneously, thus allowing computational processes like audio synthesis to

begin 50–60 ms sooner than with a non-predictive instrument (with a timing error of less than 30 ms, which is nearly imperceptible).

When all performers and audience members are co-located in the same physical space, this additional time can still be beneficial for enabling synchronization between performer action and sound synthesized on networked audience devices. For instance, real-time synthesis methods running on devices can compute the audio data in advance, and the system can be robust to delays due to local network congestion.

In our system, MalLo uses a Leap Motion sensor¹ to track the position of the tip of a wooden dowel, which is measured at 120 frames per second. As MalLo uses the velocity of the tip of the dowel to make the strike-time predictions, the performers can simply “air drum” over the sensor, without needing to make contact with any surface. This allows for the gestural controllers to be entirely silent; the only sounds in performance are those synthesized on the audience’s devices.

For each frame captured by the Leap Motion, a message containing the timing prediction for the next strike is sent from the instrument to our web server. Along with the note prediction time, the message contains additional strike information, such as the predicted velocity, and the x- and y- coordinates of the tip of the dowel (parallel to the surface of the Leap Motion).

3.2 Web Application for Interaction and Sonification

Each audience member hears and interactively modifies the sonification of performer actions using a web application. This requires each audience member to have a personal device (e.g., a laptop or mobile phone) registered on the same Wireless Local Area Network (WLAN) as the performers. Each audience member can access the web application by navigating to a URL on their device, and they hear the performance through headphones connected to this device. The clocks of the MalLo instruments and the web applications are synchronized, so that the strike-time predictions will be synthesised in the web applications on the audience devices at the same time as the Leap Motion dowels reach the bottom of their “air-drum” strikes (Figure 1).

This web application uses the Web Audio JavaScript API [21] for sound synthesis and the jQueryMobile² responsive web framework for the user interface for interactive control of sonification. The local server hosts the web application and distributes the predicted strike information from the performers to the audience over the

WLAN. The design of the sonification control interfaces in the web application were motivated by our research questions, which we discuss in the next section.

Our system also includes a web application for performers, which functions similarly to the audience application, but without interactive sonification control. (Because the MalLo DSMIs produce no sound directly, the performers must also wear headphones connected to a device that allows them to hear themselves and each other.) Our system also includes a simple user interface for a human “conductor,” which is used to send control messages to other interfaces (e.g., to enable or disable interactive control by audience interfaces).

4 Research Questions and Study Design

This work aims to investigate these following research questions:

1. **How does personalised control over performer sonification impact the audience’s experience of a performance?** Integrating audience interaction into live performance is still unconventional, and its consequences are not widely studied. As far as we know, allowing audience members to edit the mapping between performer instruments and sound has never been done before. Do audience members find this interaction to be valuable? Or engaging, expressive, informative, confusing, stressful? Etc.
2. **What is the effect of the sonification interface design on audience experience?** Although some past work has investigated how to create interfaces that allow non-experts to design their own sonifications (e.g., [5,7]), most such work has focused on supporting end-user designers in non-creative domains such as science, education, and the creation of interfaces for people with visual impairments. An interface for interactive performance sonification by audiences should be immediately usable by people without any specialised expertise, and desired criteria might include not only ease of use but also support for aesthetically pleasing output, self-expression, learning, more deeply understanding the performance, etc. We would like to explore how different interface designs might support such criteria.
3. **What factors motivate audience members to interact with the personalised sonification?** Are audience members in a music performance context interested in engaging with control over sonification as a means of musical expression? As a means to learn about sonification? As a means to more deeply understand the performers? Or out of boredom or novelty? Understanding people’s motivations

¹ <https://www.leapmotion.com/>

² <https://jquerymobile.com/>

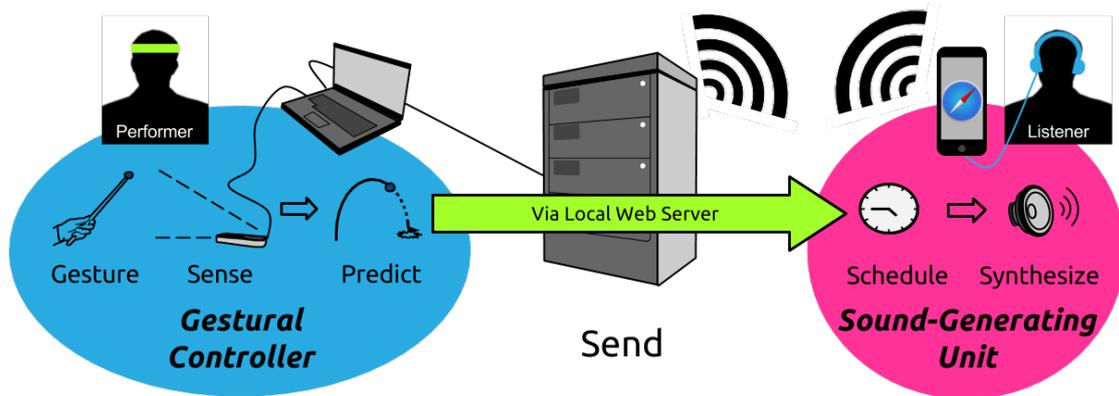


Fig. 1: In our performance setup, the MalLo gestural controller is used to predict the strike time of an “air drumming” action as the performer moves a wooden dowel over a Leap Motion sensor. This prediction is sent over a Wireless Local Area Network and is scheduled by each audience member’s device to synthesize the notes when the dowel reaches the lowest point of its trajectory.

can help us understand how to best build user interfaces to support audience interactive sonification, can help us understand the commonalities and differences between audience members and users of interactive sonification in other domains, and can inform composers and instrument builders designing new musical works.

To investigate these questions, we have conducted studies using two live musical performances. The performances featured work by different composers, for audiences with different types of prior experience with sonification. In both performances, we collected log data about audience members’ interactions with the sonification interfaces. We also disseminated post-performance surveys to audience members. These surveys included questions investigating the stated research questions (described above), questions about audience members’ demographics and prior experience with relevant topics (sonification design, sound engineering design, music composition and performance, and human-computer interaction), and open-ended questions about audience members’ experiences of the performance. A complete list of survey questions appears in Table 1.

Next, we describe how the performances and user sonification interfaces were designed in order to address the above research questions.

4.1 Q1: How does personalised control over performer sonification affect the audience’s experience?

To address this question, we designed each performance so that it had two movements. In Movement 1, the audience simply listens to the performers using sonifica-

tion parameters pre-programmed by the composer. In Movement 2, each audience member is able to interactively control the sonification. In both cases, audience members listen to the performance through headphones, synthesised by the web application running on their personal device. In order to minimise musical variation between the movements, performers were instructed to play the same structure of notes for both movements.

Movement 1 was played before Movement 2 in both performances. The order of these two conditions was not varied for two reasons: First, the experience of listening to a performance through headphones connected to a personal device is quite unfamiliar, so by playing Movement 1 first, we hoped to put listeners at ease before asking them to additionally interact with an unfamiliar interface. Second, randomising movement order among participants (i.e., where half of audience members listened passively while half interacted with the sonification interface) seemed likely to be distracting to participants in the Movement 1 condition. Such a distraction was found in a similar concert setting in which only some audience members were allowed to control the stereo panning of the guitar on stage coming out of the speakers by moving their smart phones from left to right [12]. Furthermore, by having all audience members experience the piece in the same way, we aimed to create a unified social experience for the audience, even though they were listening to the performance through their own devices and headphones.

The post-survey questionnaires asked audience members to make comparisons between the two movements (Table 1, questions 2c, 3a, and 3b) and to rate their agreement with statements related to the performance and their interaction with the interfaces (question 4).

- 1a.** Rate your level of expertise in the following areas using the scale “I know nothing,” “I know a little”, “I know a lot”, “I am an expert”:
- Sonication Design
 - Sound Engineering/Design
 - Music Composition/Performance
 - Human-Computer Interaction
- 1b.** ******Select all that apply to you:
- I am a graduate student
 - I am an undergraduate student
 - I am a student of the MUS/COS 314: Electronic and Computer Music
 - I am university faculty or staff
- 1c.** ******If you are a student, faculty, or staff, with which department(s) are you affiliated?
- 2a.** Which movement did you prefer? (Select one)
- Movement 1
 - Movement 2
- 2b.** Why did you prefer that movement?
- 2c.** Describe how your experience was different between Movement 1 and Movement 2.
- 3a.** Which interaction mode did you prefer? (Select one of the following):
- Movement 1 - No Interaction
 - Interaction Mode 1
 - Interaction Mode 2
 - Interaction Mode 3
- 3b.** Why did you prefer that interaction mode?
- 4.** Rate your agreement with the following statements (1 - Strongly Disagree, 2 - Somewhat Disagree, 3 - Neutral, 4 - Somewhat Agree, 5 - Strongly Agree):
- I understood how to use the interface
 - The controls were satisfying to use
 - The interface allowed me to creatively express myself
 - The interface allowed me to create sounds I enjoyed
 - The interface allowed me to create unique sounds
 - I felt more engaged in Movement 2 than in Movement 1
 - I wanted to share my designs with others
 - I would attend another performance of MalLo March
 - ******The technical difficulties affected my enjoyment of the piece
- 5.** Explain what motivated you to make changes to your design throughout the performance?
- 6.** What aspects of the control were most useful?
- 7.** What was your favorite part about the performance?
- 8.** Do you have any other feedback for the [insert performance name] Performance?
- 9.** ******During the NumX Performance, did you share a device with others?
- 10.** ******During the NumX Performance, how much time did you spend as the controller of the interface (compared with the other people you shared a device with)? (Rate on a scale from 1–10)

Table 1: Audience survey questions. Those marked with ****** were added to the survey for Performance 2.

4.2 Q2: What is the effect of the sonification interface design on audience experience?

How can we design interfaces for interactive performance sonification that are both usable and engaging by people with varying levels of expertise in sound de-

sign? The granularity of control over sonification seems likely to impact many factors of audience experience, including understandability and expressivity. We therefore created three different interaction modes, described below, to provide audience members with varying levels of control granularity within Movement 2.

Interaction Mode 1: Changing All Parameters At Once: Mode 1 offers very high-level control, in which users can only choose among preset sonification parameter settings defined by composers. This mode allows the composer to create multiple “versions” of the performance by specifying all of the sonification parameters for the performers. The audience can then choose among these versions using a simple interaction (i.e., choosing a preset from a list). This is similar to changing the channel on a radio and hearing the same song performed by different combinations of instruments.

Interaction Mode 2: Change All Parameters For Each Player: Mode 2 offers slightly more control, in that an audience member can change the sonification applied to each performer separately. The set of available performer presets is defined by the composer of the piece, and the audience members can choose which combination of presets to employ.

Interaction Mode 3: Change All Parameters Individually: Mode 3 gives users the most control, allowing them to directly modify multiple sonification parameters for each performer. Here, the composer decides which parameters to expose to the audience and the range of values of each parameter. Our user interfaces for Mode 3 allow audience members to continuously manipulate each control parameter using sliders. While Interaction Modes 1 and 2 allow the audience to choose what they hear from the composer-defined example sets, they simply take discrete samples from the entire design space created by the composer. Interaction Mode 3 opens up this design space so that the audience can have continuous control over the values of various sound control parameters.

Note that our aim in providing these three modes is not to identify the “right” level of granularity. Different users may have different preferences; furthermore, allowing an individual to exercise both high-level and lower-level control within the same interface could offer advantages. For instance, prior work shows that providing users with example designs can both facilitate inspiration [10] and assist users in exploring a design space in order to create new designs [16]. Because the audience is allowed to move freely between interaction modes throughout Movement 2, they can take an example design from Interaction Modes 1 or 2 into Mode 3, then vary the parameters slightly from those the composer has predefined.

In the two performances, we used a brief intermission between Movement 1 and Movement 2 to familiarise the audience to these three modes using a tutorial that demonstrated how each one worked. By the beginning of Movement 2, then, audience members should already understand how to use the modes, and our log data would reflect their intentional interactions with the different modes in the interface.

The post-concert questionnaire included questions related to interaction mode preference (Table 1, questions 3a and 3b) and the usefulness of the controls (question 6). We also collected log data (with consent) from each audience member’s device, capturing the details of their interactions with the different modes (including timing information for all mode switches and all sonification parameter changes).

4.3 Q3: What factors motivate audience members to interact with the personalised sonification control?

To inform this question, the post-concert survey included Likert-scale questions in which audience members rated their level of agreement with statements related to their designs and the design process (Table 1, question 4). Audience members were also asked an open-ended question about their motivations for changing their designs (question 5).

5 Performance 1: MalLo March

The first performance took place at the International Conference on Auditory Display (ICAD) held in Canberra, Australia on July 6th, 2016. This venue gave us access to audience members who were highly experienced with sonification, and therefore likely to be interested in exploring a personalised interactive performance sonification, as well as likely to easily understand how to use the sonification interfaces. Here we describe the compositional and performance details of the piece (titled *MalLo March*), describe the data we gathered during and after the performance, and present an analysis of this data structured by our research questions.

5.1 Composition and Performers

MalLo March was composed and designed by Reid Oda and the first author of this article. It was written for three MalLo Leap Motion controllers. Sonification events were triggered by performers’ “air-drum” hits, using MalLo’s prediction-driven synchronization. No other information about performer gestures, including speed and position of hits, was used by the sonification.

The user interface for Movement 2 supported three types of instrument sounds for each of the three performers: Drum Corps, Pitched Percussion, and Electronic. Table 2 lists the controls for each instrument.

Drum Corps	Pitched Percussion	Electronic
<ul style="list-style-type: none"> ● number of drums ● pitch ● dynamics ● amount of sustain ● amount of echo ● volume 	<ul style="list-style-type: none"> ● pattern of pitches ● vibrato ● glimmer ● amount of sustain ● amount of echo ● volume 	<ul style="list-style-type: none"> ● pattern of notes ● pitch ● grunge ● amount of sustain ● amount of echo ● volume

Table 2: The interface for Mode 3 of Performance 1 enabled audience members to choose one of three instrument types for each performer, and to manipulate these sonification parameters for each instrument type.

Performers followed the same structured improvisation for both movements, with little variation between movements. This improvisation involved sections of slow and fast sequences, as well as sections in which specific subsets of performers played alone or together. Each performer wore a headband and wristbands whose colour matched their corresponding sonification controls in the audience members’ user interface (i.e., green, blue, and pink in Figure 2 (b),(c)).

5.2 Performance

The *MalLo March* performance was the fourth of six pieces in the conference concert. As performers set up on stage, the audience was instructed on how to join the wireless network and load the web application. Headphones were handed out, and headphone jack splitters were available so people could share devices if needed. Upon opening the web application, audience members were given a consent form for log data collection followed by an interface to test whether their device properly received MalLo messages and played audio. If they did not hear the test audio, we recommended that they share a device with a neighbour in the audience.

During Movement 1 the audience had no control over the sonification of the performers on stage, and the web application instructed them to “Sit back and enjoy the piece”. The composers’ choice of sonification parameters were employed throughout the movement. At the conclusion of Movement 1 (which lasted 4.5 minutes), performers stopped playing, and the conductor web application was used to trigger the audience web

applications to display an “intermission” tutorial. This tutorial triggered a repeating sequence of notes simulating each performer playing a note in turn. This simulation allowed the audience to explore the three interaction modes, guided by an optional tutorial. After 1.3 minutes of exposure to this interface, the conductor interface was used to trigger the transition to the Movement 2 interface by all audience devices.

In Movement 2, the performers played the same structure of improvised notes as in Movement 1. However, now the audience had control over the sound of the performers on stage using the three interaction interfaces (Figure 2). Audience members could move freely among these interfaces. Movement 2 lasted 4 minutes.

5.3 Data Collection

As it was not feasible for participants to complete the survey immediately after the performance, we provided attendees with a link to the online survey via email and via the conference website. Nine people completed the survey. Figure 3 shows the nine respondents’ responses to the Likert-scale and preference survey questions.

For those 53 audience members who consented to log data collection, we logged all mode changes and sonification parameter changes. Of these 53 devices, logs showed that 37 had the web application open and receiving performers’ note data for the whole performance, so we use these in all subsequent analysis. Figure 4 shows the events logged by each device.

5.4 Discussion of Performance 1

5.4.1 Q1: How does personalised control over performer sonification affect the audience’s experience?

To answer this question, we compare the audience’s experience between the two movements. In the post-concert survey, seven of the nine respondents preferred Movement 2, one preferred Movement 1, and one did not respond (Figure 3a). Closer inspection of the surveys revealed that the one who preferred Movement 1 encountered technical difficulties in Movement 2.

When asked to explain their preference, responses from people who preferred Movement 2 included “novelty”, “because I could participate”, “the ability to interact”, “the variety of musical layers available... added more interest for the listener”, “because it was fun to customize the sound”, and “I could actively affect what I hear. Interaction made me more immersed”. One respondent replied that in Movement 2 they had “figured

out what was going on”, while another respondent commented on the composer’s sound choice for Movement 1 as “stuck on a drum sound ... and I was unable to change it until Movement 2”. When asked to describe how their experience was different between Movement 1 and Movement 2 two respondents used the words *active* when referring to Movement 2 and *passive* when referring to Movement 1: “In movement 1 I was passive listener. In Movement 2 I felt more like [an] active participant” and “active rather than passive involvement was engaging”. Other words that participants used to describe Movement 2 were: “more sense of control”, “more interesting”, and “[I] had the ability to modify the sounds to my personal liking”.

With only nine responses it is difficult to draw strong conclusions, but these responses suggest that the novelty and active participation in Movement 2 seemed to add a sense of engagement and immersion to the performance and made it a more enjoyable experience. The Likert-scale agreement ratings in Figure 3d reveal that seven of the nine participants agreed with the statement “I felt more engaged in Movement 2 than in Movement 1”. Additionally, all nine respondents agreed (three somewhat and six strongly) that “I would attend another performance of MalLo March”.

5.4.2 Q2: What is the effect of the sonification interface design on audience experience?

We employed log and survey data to analyse how the audience spent their time in each interaction mode, and what aspects of the interfaces might have influenced their experience. In Figure 4, the top plot shows the number of actions that each audience member executed within each of the three interaction modes during Movement 2. Using a paired 2-sample t-test we see a statistically significant difference in the number of interactions between Mode 1 (M=10.6, SD=9.5) and Mode 2 (M=18.9, SD=18.9); $t(53)=2.028$, $p=0.02$, as well as between Mode 1 and Mode 3 (M=27.5, SD=17.4); $t(55)=2.028$, $p<0.001$ and between Mode 2 and Mode 3; $t(71)=2.028$, $p=0.047$. This is somewhat expected as the number of elements participants can interact with also increases with the modes.

The middle plot in Figure 4 shows how much time each individual device spent in each of the interface modes. There is a statistically significant difference between the amount of time spent in Mode 3 (M=122.7 seconds, SD=70.3) and Mode 1 (M=55.9 seconds, SD=61.4) $t(70)=2.028$, $p<0.001$ and between the amount of time spent in Mode 3 and Mode 2 (M=60.7 seconds, SD=59.9) $t(70)=2.028$, $p<0.001$. However, the difference between the amount of time spent in Mode 1 and Mode 2 is not

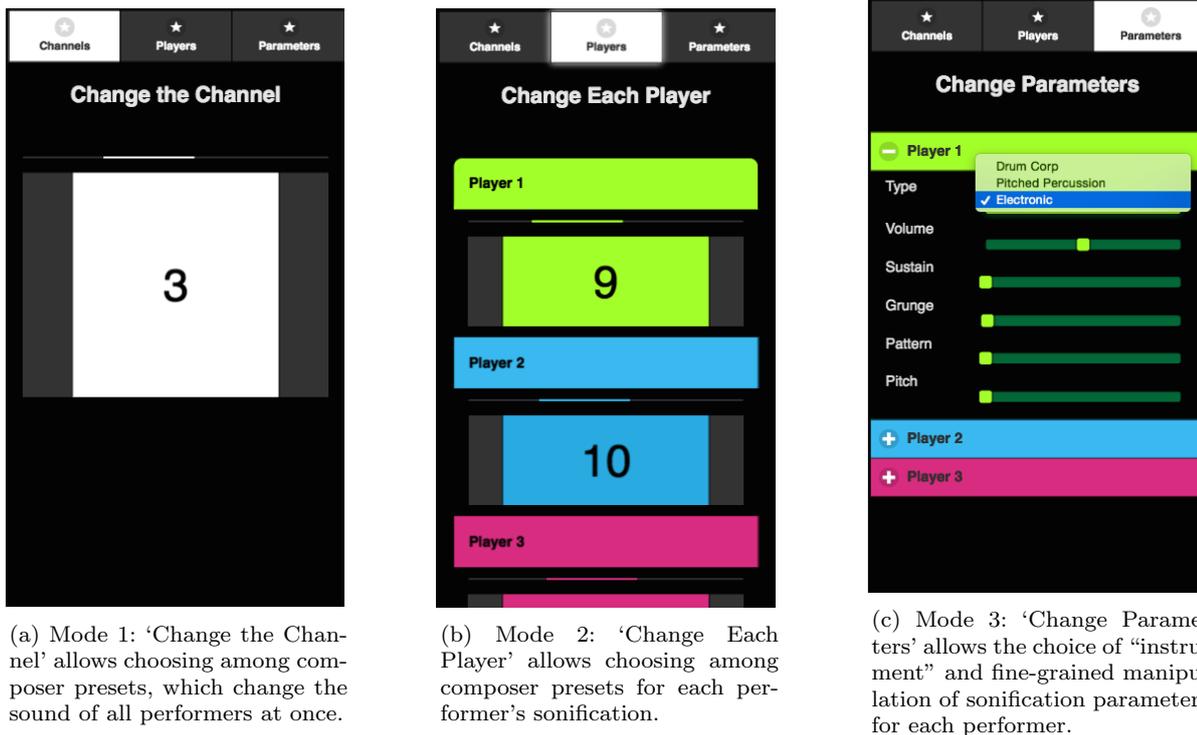


Fig. 2: User interfaces for the three interaction modes for Movement 2 of the *MalLo March* performance.

significant ($t(71)=2.028$, $p=0.732$). The bottom plot in Figure 4 shows the number of times each audience device switched between modes ($M=5.35$ times, $SD=3.33$). In summary, both the time and number of interactions in each interface varied substantially among audience devices, but it is clear that Mode 3 captured the majority of people’s time and interactions.

Figure 3b shows audience members’ preferred interaction mode, according to the survey. Six of the nine respondents preferred Mode 3. When asked why, these respondents reflected on the control that the interface gave (“*It provided nice way to fine tune [the] listening experience*”, and “*I like control*”), as well as the active engagement of the interface (“*it gave me the experience of being an active participant in the performance myself*”) and the tighter control and feedback loop (“*it is also much easier to discern the effect your interactions are having*”). Two respondents preferred Mode 2. One of these stated that it was the “*most easy to grasp without exploring the interface first... [Mode 1] lacked some visual meaning and [Mode 3] took a longish time to explore.*”. No respondent preferred Mode 1. One respondent preferred Movement 1 (no interaction), but this person reported experiencing technical difficulties during Movement 2. As Figure 4 shows, most users (all but nine) spent significant time in Mode 3.

We were curious whether other latent preferences or distinct approaches to interaction were also evident in the log data. We applied k-means clustering to the 37 interaction logs, representing each device’s log by a seven-dimensional feature vector containing the amount of time the device spent in each of the three modes, the number of interactions on the device within each of those modes, and the number of times it changed between modes. We applied the elbow criterion [15] to choose $k = 5$ clusters. Figure 4 shows the logs organised into the chosen clusters.

Cluster 1 contains devices that spent over 75% of their time in Mode 2. Users of these devices relied heavily on the examples in Mode 2 to create their sonifications, and spent very little time in Mode 3. We hypothesise that these users might have comparatively less experience in sound engineering/design and music performance/composition.

Cluster 2 contains devices that spent a large majority of time in Mode 1. Logs suggest that these devices may have experienced technical problems, causing their users to abandon their devices in Mode 1 (the default).

Cluster 3 contains devices that spent a clear majority of time in Mode 3 and had very few mode changes ($M=2.71$, $SD=1.11$). Logs show that five of these seven devices never returned to Modes 1 or 2 interfaces after entering Mode 3.

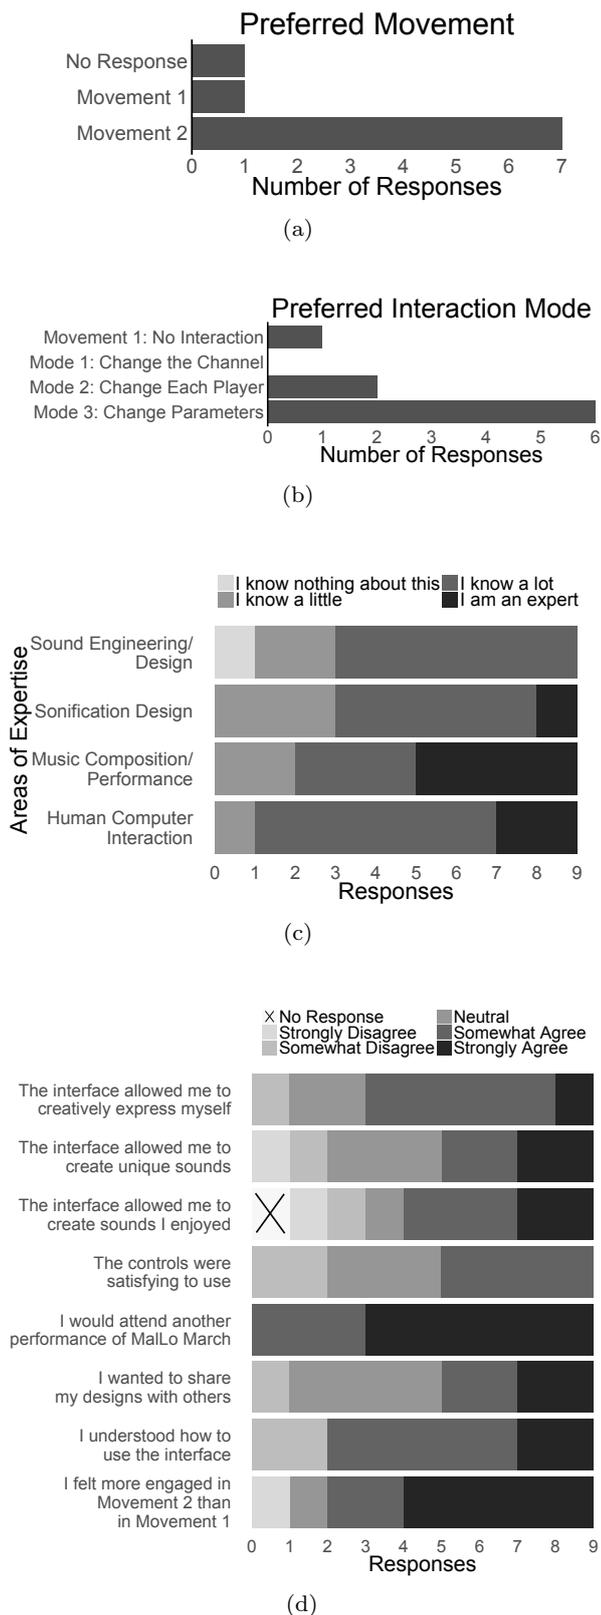


Fig. 3: Responses from the *MalLo March* audience survey showing (a) preferred movement, (b) preferred interaction mode, (c) self-reported areas of expertise, and (d) Likert-scale agreement ratings.

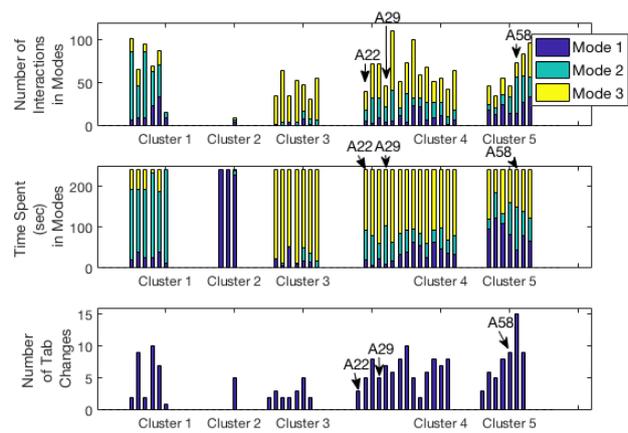


Fig. 4: The log data from Movement 2 of *MalLo March* showing the number of user interactions in each mode (top), the amount of time spent in each mode in seconds (middle), and the number of times the mode was changed (bottom). Each vertical bar represents one audience device that has been clustered using k-means clustering.

In Clusters 4 and 5, users split time among all three modes (and Cluster 4 spending slightly more time in Mode 3). These devices had a high number of mode changes ($M=6.29$, $SD=2.16$ for Cluster 4; $M=7.86$, $SD=3.85$ for Cluster 5). These users appear to be both making use of composer-provided examples and making refinements themselves.

The substantial expertise in sound design and auditory display of the conference audience (Figure 3c) makes it difficult to gauge whether the different control modes would be received differently by a more diverse audience. Within this audience, all six survey respondents who self-identified as expert in music composition/performance or sound engineering/design preferred Mode 3. One expert audience member noted that Modes 1 and 2 felt “restrictive,” and it may be that these feel too simplistic for experts.

5.4.3 Q3: What factors motivate audience members to interact with the personalised sonification control?

The survey asked respondents: “describe what motivated you to make changes to your design throughout the performance”. We applied an open coding process to analyse the responses. The responses exhibited two distinct themes, which we coded as *Explorers* and *Composers*. The *Explorers* were those who focused on the exploration of the interface and discovering how the interface affected the sounds, rather than exploring the

Explorer	Composer
I was curious to find out, how the control affet [<i>sic</i>] on sounds. It was fun to add my personal touch into to performance.	it was the desire to express my own thoughts about the arrangement of instrumental sounds and the timbral contrasts shaping the resulting musical composition (A22)
I explored Modes 1 and 2 briefly before diving into 3 where I spent almost all of my time. I changed instrument sounds first so I could get a sense of the raw sound, then starting playing with the effects to see what they did. I ran out of time to really explore to my satisfaction. (A29)	I watched the performers and decided early the kinds of sounds I would like them to be making - based on their playing style. e.g. Performer 2 was tapping with quite a regular pace, so I tried to make his sounds as percussive and deep as possible. I only wish that movement 2 had lasted a bit longer because I didn't quite finish arranging the sounds the way I had planned.
To see whta [<i>sic</i>] would happen.	I wanted to hear a greater variety of timbre between the instruments, e. g. a mix of percussive short attack notes vs long tones, harmonic sounds vs inharmonic etc
Exploration of the interface (A58)	

Table 3: Coded responses to the MalLo March audience survey prompt “describe what motivated you to make changes to you design throughout the performance”. The three responses marked with an identification number correspond to the labelled device logs in Figure 4.

sounds themselves, while the *Composers* were focused on what they heard, rather than what they were doing with the interface. Table 3 lists the seven responses from those who replied to the prompt and the codes they were assigned. Three researchers independently assigned codes to the responses; the assignments in Table 3 received an average pair-wise agreement of 81%.

We were able to match three audience members’ survey responses to their device logs (participants A58, A29, A22). These people’s design motivations are labelled in Table 3, and their logs are labelled in Figure 4. Participant A58 was coded as an Explorer by all three coders; interaction data for similar users shows a pattern of trying all possible modes and moving between them. A29 was coded as an Explorer by two coders and as a Composer by one, and A22 was coded as a Composer. Both their surveys indicated interest in the sound: “*I changed the instrument sounds first so I could get a sense of the raw sound...*” (A29), “*it was the desire to express my own thoughts about the arrangement of instrumental sounds and the timbral contrasts shaping the resulting musical composition*” (A22). It seems likely that the substantial time spent in Mode 3 by other members of Clusters 3 and 4 may be indicative of a similar participant interest in refining sonifications using the lower-level control of that interaction mode.

6 Performance 2: NumX

We conducted a second performance study, to collect data from an audience with broader interests and expertise, and to collect a greater number of survey responses. This performance took place 14 January 2017

as part of a showcase of student projects and performances at Princeton University. Practicalities of the event led to many audience members sharing devices with each other. In this section, we describe the piece, the performance, and the data collected, then analyse the data with respect to our framing research questions.

6.1 Composition and Performers

The music for this performance, titled *NumX*, was composed by student composer Joshua Collins specifically for this study. We asked him to write a piece that was similar to *MalLo March* in technical set-up and practically feasible within the open house setting. This piece was performed by three performers using MalLo Leap Motion controllers, played by holding wooden dowels and making percussive striking gestures. Each performer was a member of our university’s electronic music student ensemble; all were already familiar with performing new compositions with digital musical instruments.

The composer created the following: (1) a written score for the three performers; (2) code that uses the current sonification parameter settings to sonify incoming performer events; and (3) code to automatically control the sonification parameters in Movement 1. The composer also collaborated with our team to design the user interfaces for Modes 1–3 in the audience web application. Similar to *MalLo March*, performers repeat their actions from Movement 1 during Movement 2; the two movements are identical from their perspective.

In this composition, each performer’s sound is produced by three sine wave oscillators interacting with each other using frequency modulation. The perform-

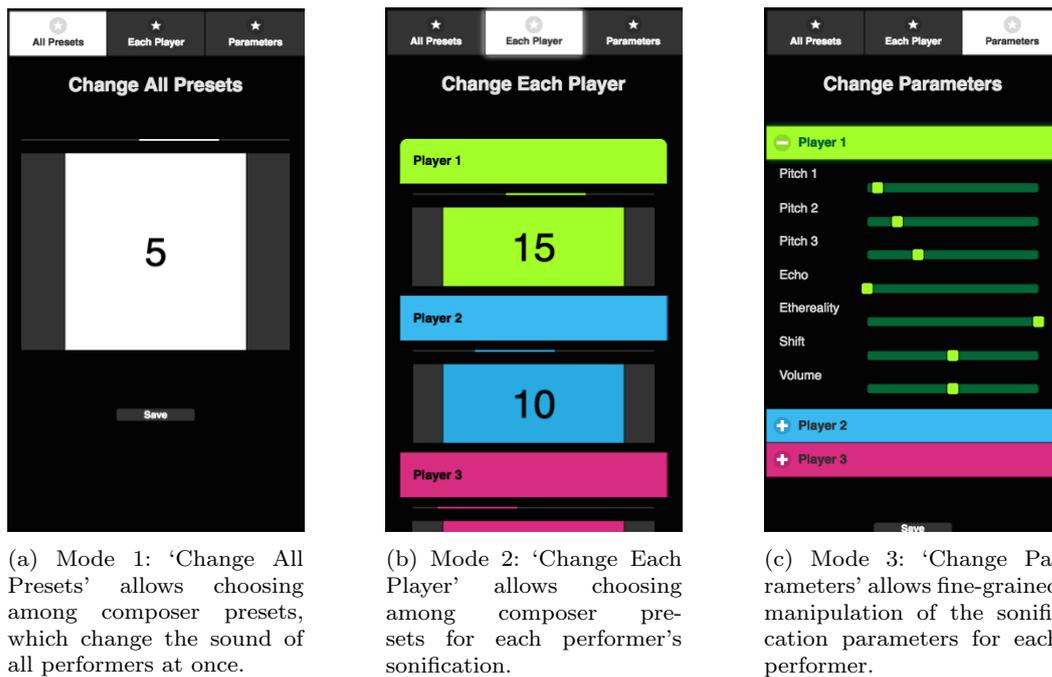


Fig. 5: Screen captures of the three interface modes of during Movement 2 of the *NumX* performance.

ers can control the volume of the notes to a small degree with the velocity of their strikes, and also control the pitch slightly with the horizontal position of the mallet strikes. In the Mode 3 interface, the audience can control the base pitch of each of the three sine waves for each performer, as well as the amount of each performer’s volume, echo, pitch shift, and reverberation. Similar to *MalLo March*, the Mode 1 and 2 interfaces allow audience members to select among preset settings for these control parameters, for the whole ensemble (Mode 1) or for individual performers (Mode 2).

6.2 Performance

As in *MalLo March*, in Movement 1, audience members listened to the piece without any interactive control. Once Movement 1 ended, audience devices showed a tutorial that guided audience members through the control interfaces for Movement 2. During this “intermission”, the device simulated performer notes while the audience explored the Movement 2 interfaces. The Movement 2 user interfaces for *NumX* (Figure 5) were almost identical to those for *MalLo March* (Figure 2). The main difference was in Mode 3, where the choice and range of sliders was adapted to match the control parameters chosen by the composer.

6.2.1 Data Collection

The concert format did not allow time for audience members to complete a survey immediately after the performance. Therefore, audience members were asked to share their email addresses to receive a link to take the survey online. The surveys were emailed to the submitted addresses a few hours after the performance.

A few additions were made to the Performance 1 survey to reflect the different audience and performance conditions (Table 1). Many audience members shared devices during this performance, but they responded to the survey as individuals. We therefore added survey questions to ask each respondent whether they shared a device, and if so, to indicate how much of Movement 2 they spent as the person controlling the sonification interface. We also added survey questions about audience members’ expertise and background. As this was a university event, these included questions about their role (e.g., student, faculty) and academic department. Because the performance was slow to start due to technical difficulties with the local WiFi, we also added a question to gauge how much this impacted people’s enjoyment of the piece.

We received eighteen survey responses. Figure 6 shows respondents’ self-rated expertise and their agreement with various statements on a five-point Likert scale.

Users of 28 audience devices consented to log data collection. Of these, 22 devices had the web application

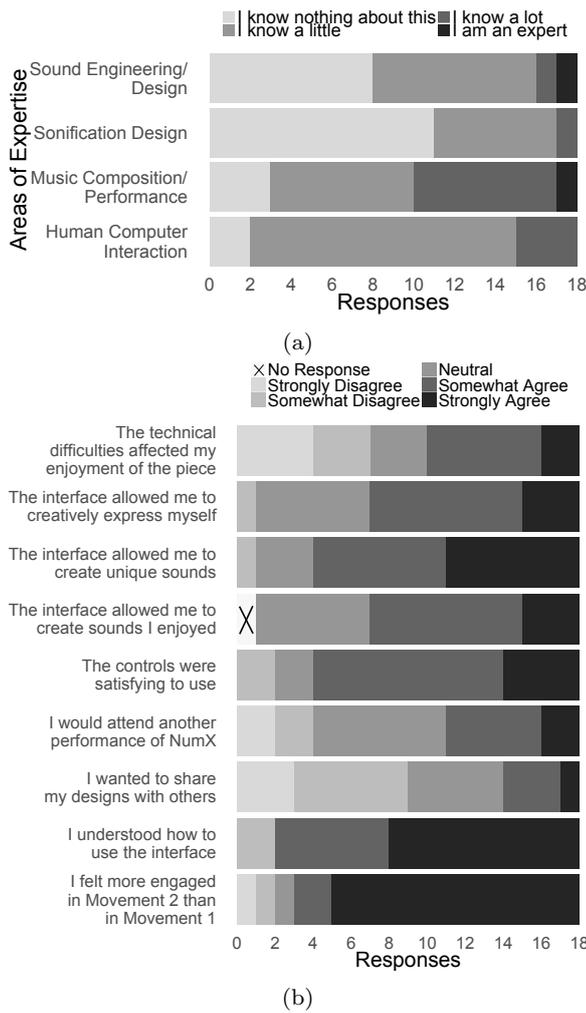


Fig. 6: The self-rated areas of expertise by *NumX* performance attendees (6a), and their Likert-scale agreement with various statements (6b).

open and receiving data for the whole performance, so we restrict our subsequent analysis to this set. Figure 7 shows a summary of audience interaction log events. Because of device sharing, we cannot know the number of people who contributed to the activity captured in a particular log. We were able to match eleven audience surveys with the log data collected from those respondents’ devices.

6.3 Discussion of Performance 2

6.3.1 Q1: How does personalised control over performer sonification affect the audience’s experience?

As shown in Figure 8c, only one of the eighteen survey respondents preferred Movement 1 over Movement

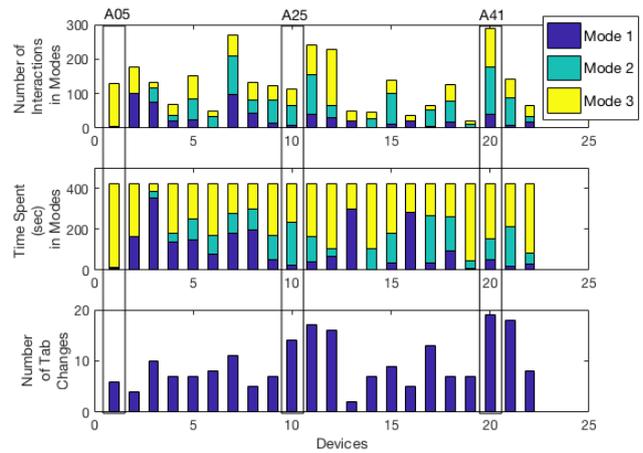


Fig. 7: The logged data from Movement 2 of the *NumX* performance showing the number of interactions in each mode (top), the amount of time spent in each mode in seconds (middle), and the number of times the mode was changed (bottom). Each vertical bar represents one audience device.

2, stating “*I enjoy listening to the performance rather than having to play around with all the different variables*”. The other seventeen respondents preferred Movement 2. Those respondents’ explanations for their preference fell roughly into three different categories:

1. **Some responses focused on the control of the interface:** “*Something to do*”, “*I liked the interactive*”, “*It was exciting to move the controls around during the performance*”, “*Movement 1 was long and with no interaction. Without the interaction, the piece wasn’t as interesting*”, and “*Because we had control over the interface.*”.
2. **Some focused on having control of the sounds and participating in the sounds of the performance:** “*I liked that I was able to customize the sounds so that it was less jarring*”, “*It was cool to have a choice in the performance, and to have the idea that everyone in the audience was hearing something different*”, “*Movement one was too quiet. Could barely hear it*”, “*It was much more dynamic, and allowed the user to have a hand in controlling the music itself*”, “*We were able to arrive at set of sounds that we found more harmonic than the ones in movement 1*”, “*It was more interesting to be able to control the sounds and see the effects of my choices*”, and “*Got to have some control over the sound myself*”.
3. **Other responses were more ambiguous:** “*more control*”, “*It got a lot more interesting after we were able to play with the effects*”, “*It was fun to play*”.

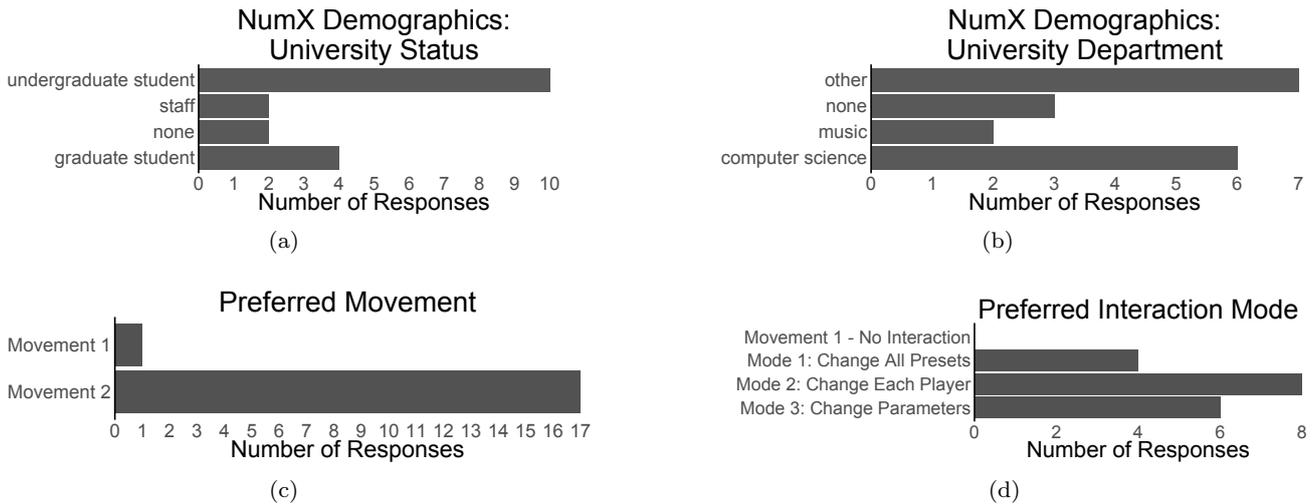


Fig. 8: Results from the *NumX* performance showing audience members’ status at the university (8a), their university department (8b), and their preferred movement (8c) and interaction mode (8d).

around with the different levels and see where that took us”, and “more interesting”.

When asked to describe how their experience differed between Movement 1 and Movement 2, some respondents discussed how their interaction changed between the two movements: *“It was different because of the controls as opposed to only watching the piece be performed”, “Movement 2 had the opportunity to experiment a bit with the qualities listed”, and “Movement 2 allowed you to control the sound that you were listening to”.* Two participants used the terms “passive” and “active” to describe how their experience differed, e.g., *“In movement 1, I was a passive listener, and in Movement 2 I felt like I was able to guide my own experience, even though there was a live performance element controlled by others”.* The one respondent who preferred Movement 1 found Movement 2 to be stressful, but they still preferred one of the Movement 2 interaction modes over the Movement 1 interface as it seemed to offer *“a good balance between completely reinterpreting the piece each time and the composer’s original intent”.*

Other respondents discussed their attention and focus. Some audience members felt distracted by the interactive interfaces, but others felt the interaction and control made them more involved in what the performers were playing. One person stated that *“I was much more focused on the interface and not as much on what the performers were actually playing”;* another said *“it might be argued that we focused less on the performance itself and more on the parameter tuning”.*

On the other hand, others noted that in Movement 2 they had *“increased involvement and attention”, and “more attention to which performer was playing what,*

and how to change the experience”, while others mentioned a lack of attention in Movement 1 (*“less attentive” and “a bit sleepy”*). Fifteen of the eighteen participants agreed that they felt more engaged in Movement 2 than in Movement 1, with thirteen of those agreeing strongly. However, the survey did not force users to specify whether this engagement was with the interface or with the performers/performance.

6.3.2 Q2: What is the effect of the sonification interface design on audience experience?

Figure 7 shows a summary of all devices’ log data for Performance 2. As in performance 1, we use a paired 2-sample t-test and find a statistically significant difference between the number of interactions in Mode 1 ($M=9.4$, $SD=4.8$) and Mode 2 ($M=27.2$, $SD=29.1$) $t(22)=2.080$, $p=0.010$, and between the number of interactions between Mode 1 and Mode 3 ($M=47.7$ and $SD=40.3$) $t(21)=2.080$, $p<0.001$. However, the difference between the number of interactions in Mode 2 and Mode 3 is not statistically significant ($t(38)=2.028$, $p=0.060$).

We also see a significant difference between the amount of time spent in Mode 3 ($M=225.9$ seconds, $DS=91.2$) and Mode 1 ($M=104.4$ seconds, $SD=102.3$) $t(41)=2.028$, $p<0.001$ and the amount of time spent in Mode 3 and Mode 2 ($M=91.0$ seconds, $SD=69.6$) $t(39)=2.028$, $p<0.001$. However, as in performance 1, the difference in amount of time spent in Mode 1 and Mode 2 is not statistically significant ($t(37)=2.028$, $p=0.615$).

The survey asked whether the respondent shared a device, and if so how much time they spent as the person controlling the device. With these responses, we

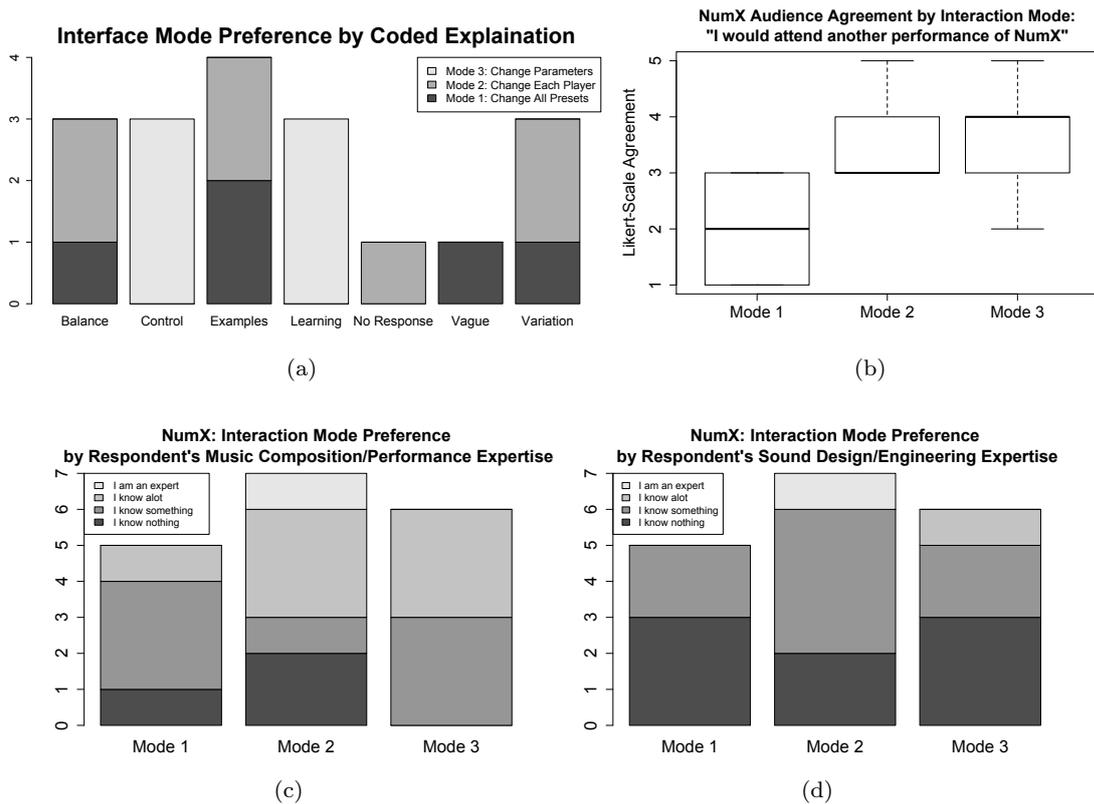


Fig. 9: Survey analysis results from the *NumX* performance showing: (a) the preferred interaction modes by the coded explanation of why those modes were preferred, (b) the summary of responses on a Likert scale from 1 (strongly disagree) to 5 (strongly agree) to the statement “*I would attend another performance of NumX*”, and the areas of expertise in (c) Music Performance/Composition and (d) Sound Engineering/Design plotted by the preferred interaction mode.

were able to associate three surveys with complete device logs for which the respondent reported being the sole person controlling the device. These devices are marked in Figure 7 (device IDs A05, A25, A41).

As shown in Figure 8d, four respondents reported that they preferred interaction Mode 1, six preferred Mode 2, and five preferred Mode 3; none preferred Movement 1 (no interaction). The survey also asked people to explain their preference for their chosen mode. Applying an open coding approach to the responses, five themes arose, and three researchers independently applied these codes to each response (with an average pair-wise agreement of 82.4%). Below we list the description of the themes and responses assigned this theme in coding. In addition to the responses below, one survey participant did not respond, and one responded vaguely (“*Preferred sound and was interesting*”).

Balance: Described a balance or trade-off the preferred interface struck between usability and control, or between composer and personal intent (“*I could make the sounds more want I wanted but didn’t get too*

confused”, “*It gave enough control to make some difference but not enough to get me lost*”, and “*Seems like a good balance between completely reinterpreting the piece each time and the composer’s original intent*”)

Control: Described a preference for the control or active involvement the preferred interface provided (“*It seemed like there was more control involved*”, “*Because you could use the delay whenever you wanted*”, and “*It provided the most significant results and the most control, which was interesting*”)

Learning/Effects: Described being able to learn from the effects of changing the interface (“*Could learn more*”, “*It was cool to see which parameter changed the sound in which way*”, and “*allowed me to pick out the individual instruments better*”)

Examples: Described using the interface to find desirable sounds without needing to change many things (“*It gave me a good baseline to then fuss with the parameters [in Mode 3]*”, “*sounded better than me messing with everything*”, “*It changed the sound most*

dramatically, but also gave you control (unlike the presets [in Mode 1]), and “we tried all of them, and while I liked the detail allowed in [Mode 3], it was actually more fun to listen while occasionally changing the individual preset instruments [in Mode 2] - it allowed me to focus more on the listening experience and less on the details of the sound”)

Variation: Described the variation in sounds/controls that the interface provided (“I liked hearing the difference of each setting”, “largest variation”, and “I could notice the difference in sound the strongest”)

A chi-squared test found a significant relationship between the mode respondents preferred and their stated reason for this preference, based on the assigned code above ($\chi^2(12) = 24.43, p < 0.05$). As shown in Figure 9a, respondents who preferred Mode 3 exclusively responded that they preferred either the control the interface offered or being able to learn from the interactions with the interface. Those who preferred Mode 2 and Mode 1 liked the balance the interfaces had to offer, the variety of sounds the interfaces contained, and the presentation of example designs that could be developed without much effort.

We further explored the connection between interaction mode preference and other survey question responses. We ran a Kruskal-Wallis rank sum test for all Likert-scale statement agreement questions (Question 4 in Table 1). We found that participants’ level of agreement with the statement “I would attend another performance of NumX” significantly differed among participants with different mode preferences ($H = 7.306, p < 0.05$). Those who preferred Mode 3 (which offered most fine-grained control) were more likely to want to attend another performance of NumX, while those who preferred Mode 1 (most course-grained control) were more likely to disagree (Figure 9b).

We did not find any significant effects of expertise (Kruskal-Wallis rank sum test) or university position and department (chi-squared test) on audience members’ preferred interaction mode.

6.3.3 Q3: What factors motivate audience members to interact with the personalised sonification control?

The codes of *Explorer* and *Composer* arising from analysis of Performance 1 surveys again seemed relevant to explain Performance 2 respondents’ descriptions of what motivated them to make changes to their sonification designs. The coded survey responses appear in Table 4. These were coded by three researchers with 87.5% average pair-wise agreement.

We applied Kruskal-Wallis rank sum and chi-squared tests to investigate whether audience member exper-

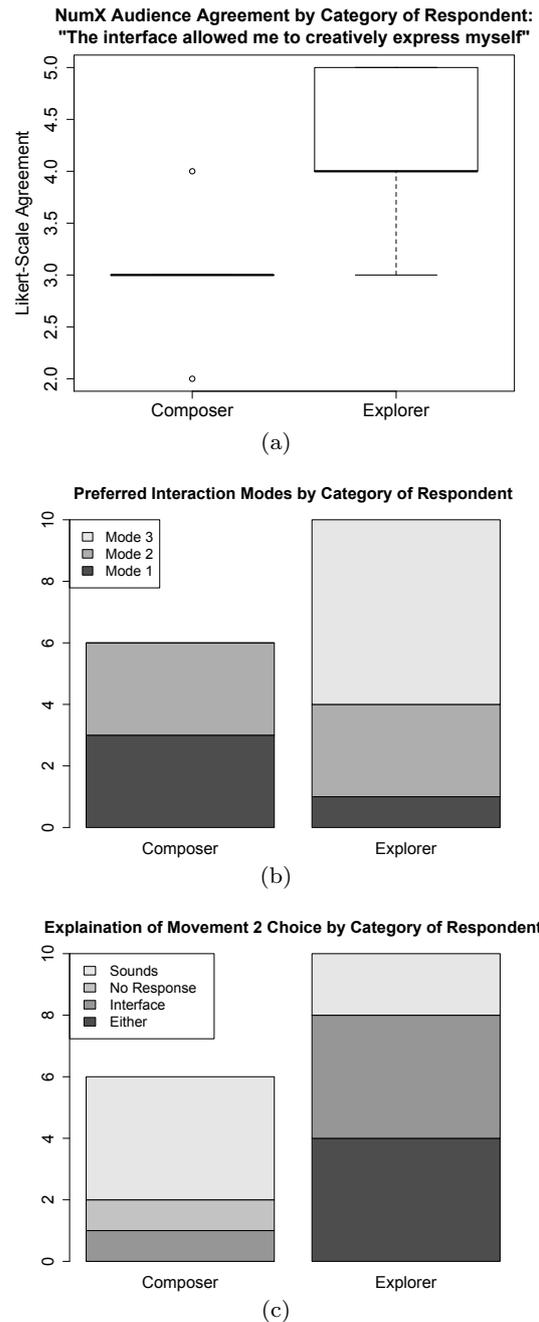


Fig. 10: The Likert-scale agreement of Composers and Explorers to the statement “The interface allowed me to creatively express myself”, and the responses that were coded as Explorers and Composers plotted by their: (b) preferred interaction modes and (c) explanation of why they preferred Movement 2 over Movement 1.

tise or university status, respectively, had an effect on Explorer/Composer coding; no significant effects were found. We applied Mann-Whitney tests to examine whether Composers and Explorers differed on the Likert-

Explorer	Composer
Curiosity was a main motivator, but also since the sound quality was really bad and quite quiet (maybe the earbuds were at fault) I was trying to raise the volume so I could hear the intricacies of the piece best.	I liked the deeper sounds and didn't like the beeping/echoing. But I wanted each sound to be fairly different. So I changed the settings to get that.
Curiosity	I liked hearing the individual changes of each setting
I just wanted to explore the various controls (A41)	I would get bored with an unchanging set of sounds
curiosity	To hear the different sounds. (A5)
Mostly curiosity, just act of observing how change in specific preset or parameter would affect the sound. Once setting on a single performer sound it was fun to look for sounds that matched the first one well.	Usually when I was bored with the current state of the synthesis and wanted to have a change for musical contrast.
Because it seemed fun.	I just changed it until I settled on something that sounded somewhat pleasant. (A25)
Identifying which controls were which	
curiosity, pursuit [<i>sic</i>] of new sounds	
just wanted to spice things up	
Just curiosity	

Table 4: Coded responses to the *NumX* audience survey prompt “describe what motivated you to make changes to your design throughout the performance”. The three responses marked with an identification number can be matched to the corresponding logs in Figure 7.

Statements	Results
The technical difficulties affected my enjoyment of the piece	$U = 23, p\text{-value} = 0.3225$
The interface allowed me to creatively express myself	$U = 9, p = 0.01198 *$
The interface allowed me to create unique sounds	$U = 21.5, p = 0.23955$
The interface allowed me to create sounds I enjoyed	$U = 22.5, p = 0.283$
The controls were satisfying to use	$U = 16, p = 0.07544$
I would attend another performance of <i>NumX</i>	$U = 16.5, p = 0.09675$
I wanted to share my designs with others	$U = 23.5, p = 0.3473$
I understood how to use the interface	$U = 28, p = 0.6067$
I felt more engaged in Movement 2 than in Movement 1	$U = 28, p = 0.5434$

(a) Mann-Whitney tests for differences between *Composers* and *Explorers* in the Likert-scale responses to the survey statements.

Questions	Results
Which interaction mode did you prefer?	$\chi^2(2) = 7.1477, p = 0.02805 *$
Why did you prefer that interaction mode? (coded)	$\chi^2(6) = 8.9722, p = 0.1751$

(b) Chi-Squared tests for differences in response between *Composers* and *Explorers* on their interaction mode preference, and on their coded response as to why they preferred that interaction mode.

Table 5: Tests for significant differences in survey responses between *Composers* and *Explorers*.

scale survey questions. Results appear in Table 5a. We found agreement with the statement “*The interface allowed me to creatively express myself*” was stronger for *Explorers* (Mdn = 4) than *Composers* (Mdn = 3), $U = 9, p = 0.01198$ (Figure 10a). The difference was not dramatic (average response from *Composers* was “neutral” and for *Explorers* was “somewhat agree”), yet it gives some insight into how well the interface supported activities that these two groups understood as expressive.

We applied a chi-squared test (Table 5b) to examine whether participants’ interaction mode preference and preference rationale coding (i.e., Balance, Control, Learning, Examples, Variation) differed between *Explorers* and *Composers*. There was a significant difference in the preferred interaction mode ($\chi^2(2) = 7.1477, p = 0.02805$). As seen in Figure 10b, *Explorers* included all respondents who preferred Mode 3, three of the six respondents who preferred Mode 2, and only one of the four respondents who preferred Mode 1. This suggests

that audience members who were motivated to change their designs based on the sound, rather than on exploring the interface, did not prefer the Mode 3 interface that gave the finest granularity of control over the sounds. On the other hand, those who preferred Mode 3 were the audience members who stated a clear desire to explore the interface.

7 Discussion: Audience Preferences, Engagement, and Interaction

Next, we examine themes that arise from the comparison and synthesis of the two performance studies.

7.1 Users' Preference for Interaction

All but one survey respondent in each performance preferred Movement 2 to Movement 1, which suggests that interactive sonification was almost exclusively preferred to passive listening by both audiences for these pieces.

What drove the nearly unanimous preference for Movement 2? The opportunity for taking an active role seems to be one clear factor. The term “*passive listener*” was used several times by participants in both performances to describe the experience of Movement 1. In Performance 2, audience members commented that their experience with Movement 1 was just “sitting and listening”, while others mentioned that it felt long and repetitive, and that the different sounds were very similar. One participant even mentioned that it made them “*a bit sleepy*”.

Some people explicitly called out the personalised interaction with sound as driving their enjoyment and preference for Movement 2. One respondent from Performance 1 preferred Movement 2 “*Because it was fun to customize the sound and set the mood I wanted for the performance*”. When asked what was their favourite part about the performance, four out of the nine Performance 1 respondents directly referred to being able to create their own unique music and sounds.

Most respondents' comments about their interaction with the interface also indicated that this positively influenced their experience. In Performance 1, participants commented that the interaction made them more immersed and engaged in the performance, and in Performance 2 participants noted that they found the interface control exciting, fun, and interesting. One respondent from Performance 1 preferred Movement 2 because of “*The ability to interact. I love touch screen sound interfaces and use them daily*”. The same respondent replied that their favourite part of the performance

was: “*Definitely a buzz working wirelessly*”. In our experimental setup, it must be noted that Movement 1 and Movement 2 differed by both the *presence* of any interaction at all, and the specific ability to interactively control sonification. The differential impact of the ability to interact at all and the ability to interact with sound cannot be determined from this experimental design, and the audience responses above suggest that some of people's preference for Movement 2 may be driven by the mere presence of interaction. Nevertheless, if the primary goal of a performance is an engaged audience, interactive sonification seems to be one method for achieving this.

There is a danger that interactive technologies may be experienced as disruptive to live performance, however. Work by [12] found that smartphones, in particular, could be obtrusive to audience members, even when used to facilitate audience participation by controlling panning of guitars on stage. Furthermore, giving the audience an interactive interface may take the audience's attention away from the performers and from the performance in general. For instance, one *NumX* audience member commented that “*In movement 2, I was much more focused on the interface and not as much on what the performers were actually playing*”. In the next section we discuss how the control of the sounds and interface contributes to the audience's level of attention and engagement.

First, though, it is interesting to examine the one participant who preferred Movement 1 in Performance 2. (The one who preferred Movement 1 in Performance 1 experienced a technical difficulty with their interface, so they are less interesting.) They stated a desire to simply listen to the piece without the stress they felt when interacting with the interface (in the *NumX* performance). This explanation is interesting as this respondent could have chosen at any time during Movement 2 not to use the interface and leave it on one of the preset modes instead of being “stressed out” by the option of interaction. This suggests that there may be a sense of obligation to use an interactive interface, and that it may be beneficial for audience enjoyment to provide an explicitly interaction-free mode at all times.

7.2 Interaction, Engagement, and Attention

While both audiences generally agreed that they felt more engaged in Movement 2 (average rating 4.1 out of 5 for Performance 1 and 4.4 for Performance 2), it is unclear which aspect(s) of the performance drove this feeling of engagement. One audience member from Performance 2 highlights a possible tension between attention to interacting with the sonification and attention to

the performance in their description of their experiences in the two movements: *“Movement 2 was more engaging as we were playing around with the parameters. At the same time it might be argued that we focused less on the performance itself and more on the parameter tuning”*. This statement suggests that audience members may experience engagement with the sonification control and engagement with the performance as being distinct and even competing phenomena.

On the other hand, one Performance 2 audience member commented that *“In movement 1 I was just listening, while in movement 2 I was paying more attention to which performer was playing what, and how to change the experience”*. For this person, it appears that increased engagement with the sonification control deepened their engagement with the performance. Another audience member commented that their favourite part of the performance was *“the unusual experience of being in control of a performer’s sound”*. Another commented that *“It was more interesting to be able to control the sounds and see the effects of my choices”*. Thus, interactive sonification of performance can not only foster engagement by giving the audience “something to do”, but can also elucidate connections between the audience members’ actions, the sounds they are hearing, and the actions of the performers.

Additionally, the personalized interaction allowed audience members to design for themselves. Based on the responses to what motivated people to change their designs, we see that the purpose of the designs differed across audience members (some were interested in exploring the space of the sounds, while others were interested in specific aesthetic goals). This ties closely with Grond and Hermann’s conceptualization for sonifications where the sound can serve different purposes, which motivates different aesthetic design decisions [8].

7.3 Supporting Use (and Learning) by Users with Varying Expertise

The log data (Figures 4, 7) show that participants in both performances spent significantly more time in Mode 3 than in the other interaction modes. However, comparing the two performances, we see a difference: only in Performance 1 was there significantly more interactions in Mode 3 than in Mode 2. One explanation may be that the audience in Performance 1 consisted largely of experts in sonification, sound design, and music composition, whereas the Performance 2 audience was relatively less expert.

Furthermore, the survey responses regarding users’ mode preferences can be summarised thus: all but one

audience member (one who experienced technical difficulties) preferred Mode 1, 2, or 3 rather than Movement 1 / No interaction. Audience members who preferred Mode 3 preferred it for the amount of control it provided, its interactive nature, and because the audience felt they were able to learn from the effects of changing the interface. Those who preferred Modes 1 or 2 explained this preference in terms of three main criteria: (1) access to desirable sounds without needing to change too many things in the interface, (2) a wide variation in sound and controls, and (3) a balance or trade-off that was struck, e.g., between understandability and the time it took to explore the interface.

It appears, then, that giving audience members access to example sonifications in Modes 1 and 2 was useful to audience members (especially for Performance 2, where the audience had less expertise in sound design and music performance/composition). In Performance 2, the three audience members who knew nothing about music performance/composition preferred Modes 1 and 2. Choosing among example sonifications in Modes 1 and 2 also served as a stepping stone to more sophisticated interaction with the sound for some users. For instance, when asked what aspects of the control were most useful, one respondent from Performance 1 stated *“At the beginning of movement 2, I found that [Mode 2] was quite useful, but I found it a bit restrictive after a couple of minutes so I changed to [Mode 3].”*

Some audience members from Performance 2, where there were fewer experts in attendance, additionally noted the value of Mode 3 in enabling them to learn about sound and sonification. Participants appreciated being able to easily notice the sonic effects that their parameter changes had, and mentioned this being useful to learning. In this performance, the Mode 3 interface used wording that may be unfamiliar to many people: specifically, the control sliders were named Pitch 1, Pitch 2, Pitch 3, Ethereality, Echo, Shift, and Volume. Presumably, the fact that changing most of these controls has a simple and immediate effect on the sonification supports effective exploration and learning of these concepts; other control interfaces that expose higher level or less perceptible sonification parameters might not provide this benefit.

7.4 Timbre, Perception, and Understanding

Both performances used the same experimental setup with two movements and three different interaction modes in Movement 2. Both performances had very similar user interfaces that allowed for the same types of physical interactions and the same visual format (as can be seen by the screen captures in Figures 2 and 5).

However, the nature and breadth of sounds accessible within Mode 3 differed appreciably between the two performances. In Performance 1, audience members could choose a type of instrument sound (Electronic Drum, Pitched Percussion, Drum Corp), each of which exposed control over different sound parameters (Table 2). In Performance 2, though, the audience could only control sounds within a single timbral space.

In Performance 1, the audience reported that having three different timbral spaces was one of the most useful aspects of the control: *“The different instruments make it easier to work on one sound while others are playing”* and *“individually shaping those timbral features I thought made for interesting contrasts between instruments”*. Another Performance 1 participant was motivated to change their design because they *“wanted to hear a greater variety of timbre between the instruments, e.g. a mix of percussive short attack notes vs long tones, harmonic sounds vs inharmonic etc”*. In contrast, some Performance 2 audience members suggested that a wider variety of timbre would have been useful: *“I think main problem I had was sometimes telling which performer made which sound, especially in movement 1. I think providing either visual queues in the interface when specific performer makes a sound or having more distinct sound samples would make this easier”* and *“I could have used a wider range of timbre, or some kind of binary timbre change as opposed to a single sound with subtle parameter ranges. It could have been interesting for the piece to take advantage of pitched material possibilities”*.

One of our motivations for building a system for interactive, personalised performance sonification was to allow audience members to take on a more active and creative role within music performance. One would therefore expect to see audience comments evaluating the available sounds and timbres to relate to whether they support the creation of interesting or musically pleasing sounds. While we do see such comments in the surveys, other comments suggest that audience members’ interest in changing the sound additionally stems from a motivation to better understand and distinguish among the performers.

Furthermore, some audience comments reveal that distinct timbres are important for supporting another aim: that of obtaining efficient feedback about one’s own interactions. Being able to hear the changes to the sound as they manipulate sonification controls is important for the audience to form an understanding of how they are influencing the sound. We hypothesize that this is especially true for those who are new to sound design and music composition, as hearing the audio feedback can teach them by example the meaning

of particular sound design/audio engineering terminology used within the interface. A more trivial aspect is that this audio feedback confirms that the interface is working and that there are no technical problems.

7.5 Implications for Future Performances

This approach offers an interesting opportunity for future performances, in that data about performers’ actions can be logged and used to re-create the performance without their physical presence. The stream of performer actions can be re-played in real-time, allowing for interactive audience sonification using the same interface. It becomes a performance across time, with the performers playing in the past, and their actions being sonified in the present on a listener’s device. The log data from an audience member’s device also allows that person’s experience to be exactly re-created at a later time. This could enable an entirely new audience to experience the original performance through the sonification designed by members of the original audience, who have become the “producers” of the piece.

8 Conclusions

We presented a unique approach to interactive sonification, in which each audience member of a musical performance interactively modifies the sonification of live performers’ actions. We described how this approach can be implemented using a DSMI called MalLo and a web application interface used by the audience to control the sounds of the MalLos on stage. We used two live performances to begin to understand how personal, interactive sonification control was used by the audience and how making this interaction available impacts the audience experience.

Survey responses from participants suggest that using an interface to interactively control the sounds of the performers on stage can make for an enjoyable and engaging experience. Interfaces for interactive sonification can support multiple approaches to engaging with sound, including exploring alternative designs as well as crafting a personalised sonic representation of the performance. Interaction can also create a connection with the performers on stage, help people learn about sonification, and make audience members feel like more active participants.

Different audience members are likely to prefer different granularities of interaction, depending on their prior expertise with sound and music as well as other factors. Providing example sonifications as starting points

can help users efficiently explore the design space; create a balance between levels of control, exploration, and time commitment; as well as provide a wide variation in sounds and designs. More expert users may eventually find coarse-grained control too simplistic, and may enjoy access to manipulating lower-level parameters.

Different types of sonification interfaces can also assist in facilitating design for audience members with different design goals. We found that in the second performance, those who were motivated by exploring the interface liked the interaction mode that had the most complex interface and gave the finest level of control over the sounds, while those who were more motivated by the sounds themselves appreciated the example designs that were supplied by the composers of the performance. Being able to refine preset sonifications, by switching from a coarse-grained preset-selection mode into a finer-grained control mode, also facilitates the design process. When the audience is given control over the sounds of multiple performers, we found that having different timbral spaces available has multiple benefits. The audience can more easily differentiate each performer from the others, and they can more easily hear their changes to the sonification parameters for particular performers.

In summary, we have proposed a new approach to enabling audience participation in live musical performance, which offers the audience interactive control over the sound of the performers on stage. While we have only explored two performances with similar setups, there are many other possible ways such a performance could be designed. This work contributes to a broader perspective on interactive sonification, showing how a musical performance can incorporate benefits of interactive sonification that have been explored in non-creative domains. Our approach shows that creative expression and exploring new relationships with data can exist simultaneously as dual consequences of providing interactive audience control over sonification in a musical context.

Acknowledgements This material is based upon work supported under a National Science Foundation Graduate Research Fellowship under grant number 1148900. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Dr. Wolf is also grateful for the support of the Microsoft Research Womens Fellowship and the Gordon Y.S. Wu Fellowship from Princeton University's School of Engineering and Applied Science in enabling this work.

References

1. Barthet, M., Thalmann, F., Fazekas, G., Sandler, M., Wiggins, G., et al.: Crossroads: Interactive music systems transforming performance, production and listening. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI (2016)
2. Blaine, T., Fels, S.: Contexts of collaborative musical experiences. In: Proceedings of the 2003 conference on New interfaces for musical expression, pp. 129–134. National University of Singapore (2003)
3. Dahl, L., Herrera, J., Wilkerson, C.: Tweetdreams: Making music with the audience and the world using real-time twitter data. In: Proceedings of the International Conference on New Interfaces for Musical Expression (NIME) (2011)
4. Dahl, S.: Striking movements: A survey of motion analysis of percussionists. *Acoustical Science and Technology* **32** (2011)
5. Davison, B.K., Walker, B.N.: Sonification sandbox reconstruction: Software standard for auditory graphs. In: Proceedings of the International Conference on Auditory Display (ICAD) (2007)
6. Eriksson, M., Halvorsen, K.A., Gullstrand, L.: Immediate effect of visual and auditory feedback to control the running mechanics of well-trained athletes. *Journal of sports sciences* **29**(3), 253–262 (2011)
7. Goudarzi, V., Vogt, K., Höldrich, R.: Observations on an interdisciplinary design process using a sonification framework. In: Proceedings of the International Conference on Auditory Display (ICAD) (2015)
8. Grond, F., Hermann, T.: Interactive sonification for data exploration: How listening modes and display purposes define design guidelines. *Organised Sound* **19**(1), 41–51 (2014)
9. Hermann, T., Hunt, A.: Guest Editors' Introduction: An Introduction to Interactive Sonification. *IEEE MultiMedia* **12**(2), 20–24 (2005)
10. Herring, S.R., Chang, C.C., Krantzler, J., Bailey, B.P.: Getting inspired! Understanding how and why examples are used in creative design practice. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (2009)
11. Hill, E., Cherston, J., Goldfarb, S., Paradiso, J.: ATLAS data sonification : a new interface for musical expression and public interaction. *PoS ICHEP2016*, 1042 (2017)
12. Hödl, O., Kayali, F., Fitzpatrick, G.: Designing interactive audience participation using smart phones in a musical performance. In: Proceedings of the International Computer Music Conference (ICMC) (2012)
13. Hunt, A., Hermann, T.: The Sonification Handbook, chap. Interactive Sonification, pp. 1–6. Logos Publishing House (2011)
14. Jin, Z., Oda, R., Finkelstein, A., Fiebrink, R.: Mallo: A distributed, synchronized instrument for internet music performance. In: Proceedings of the International Conference on New Interfaces for Musical Expression (NIME) (2015)
15. Kassambara, A.: Practical Guide to Cluster Analysis in R: Unsupervised Machine Learning, vol. 1. STHDA (2017)
16. Lee, B., Srivastava, S., Kumar, R., Brafman, R., Klemmer, S.R.: Designing with interactive example galleries. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 2257–2266 (2010)
17. Lee, S.W., de Carvalho Junior, A.D., Essl, G.: Crowd in C[loud]: Audience participation music with online dating

- metaphor using cloud service. In: Proceedings of the Web Audio Conference (WAC) (2016)
18. Madhavan, N., Snyder, J.: Constellation: A musical exploration of phone-based audience interaction roles. In: Proceedings of the Web Audio Conference (WAC) (2016)
 19. Miranda, E.R., Wanderley, M.: *New Digital Musical Instruments: Control And Interaction Beyond the Keyboard* (Computer Music and Digital Audio Series). A-R Editions, Inc. (2006)
 20. Oh, J., Wang, G.: Audience-participation techniques based on social mobile computing. In: Proceedings of the International Computer Music Conference (ICMC) (2011)
 21. Smus, B.: *Web Audio API*. O'Reilly Media, Inc. (2013)
 22. Weger, M., Pirrò, D., Wankhammer, A., Höldrich, R.: Discrimination of tremor diseases by interactive sonification. In: 5th Interactive Sonification Workshop (ISon), vol. 12 (2016)
 23. Weinberg, G., Gan, S.L.: The squeezables: Toward an expressive and interdependent multi-player musical instrument. *Computer Music Journal* **25**(2), 37–45 (2001)
 24. Zhao, H., Smith, B.K., Norman, K., Plaisant, C., Shneiderman, B.: Interactive sonification of choropleth maps. *IEEE MultiMedia* **12**(2), 26–35 (2005)