Buildings as Audio Visual Synthesisers: Experiments Performing Live Music on Wirelessly Networked Multi-Speaker Media Architectures

Oliver Bown[1], Kurt Mikolajzcyk[2], Sam Ferguson[2], and Benedict Carey[1]
[1]Interactive Media Lab, UNSW Art and Design, University of New South Wales
[2]Creativity and Cognition Studios, School of Computer Science, University of Technology Sydney
Sydney, Australia
[o.bown|b.carey]@unsw.edu.au, [samuel.ferguson|kurt.mikolajzcyk]@uts.edu.au

Abstract

This paper presents an approach to expanding live music performance practices to encompass sonic media architectures. We demonstrate a method for creating a playable audio-visual synthesiser that incorporates the notion that the space itself is a medium for performance. We discuss the design concepts that inform this process, as well as detailing specific simulation tools and a creative workflow that facilitates development of performance experiments within architectural spaces.

Keywords

Media architecture, internet of sounds, spatial audio, live music technology.

Introduction

Live musical performance has traditionally been a significant driver of architectural design excellence, as well as media technology innovation. Opera houses and recital halls have been world-famous architectural sites throughout recent history, partially based on the social importance of live music for society, as well as the architectural excellence that may be shown in particular designs. Media technology development has also been propelled by musical performance contexts, and media architectures are a specific area of current research interest and technology innovation.

In this paper we seek to discuss an approach to combining these areas of interest, by augmenting music performance contexts with a media architecture that is controllable by typical musical DAW programs, and allows for creative experimentation with music, light, and space.

Background

Media architectures are physical augmentations of the built environment that bring audio, visual, tactile and sensing digital elements [6, 9]. Most work in media architecture focuses on visual elements – lighting and the design of non-standard screens that integrate with architectural forms [12, 11]. Researchers have also developed more complex systems that integrate interactive system design with media architectures [22, 13, 14] and have also investigated methods for incorporating kinetic and robotic capabilities into these systems [21, 16].

The incorporation of widely distributed speaker arrays into buildings is beginning to grow in prominence. Audio media architecture is challenging for a number of reasons: introducing sound into the built environment has greater potential to be considered a nuisance (although light pollution is a recognised issue in media architecture [9]) and less well associated with providing added value to the built environment; and technically, massively multispeaker sound design is harder to achieve [8]. Nevertheless, a number of teams are developing media-architecture oriented sound design systems, including our own work developing frameworks for rapidly developing and deploying distributed sound installations.

In the area of live music performance, whilst the academic, experimental and art music worlds have long embraced more-than-stereo audio systems, the vast majority of music performance contexts remain largely bound to PA systems with stereo configuration. Factors at play include the fact that the stereo format is widely adopted by venues and is considered to be more manageable with a constant turnover of touring acts, and that there is also no strong incentive to go beyond stereo (indeed large venue PAs offer only a limited stereo experience). The technology to render sound using multi-speaker audio for these contexts has been readily available for some time [19], and indeed has long been standardised for cinema – another common public sound reproduction context. But the limited embrace of multi-speaker performance practices, as well as unclear experiential benefits for audiences, means it would still be considered novel today to see a local band performing on anything other than a standard stereo PA.

At the intersection of these two fields, live music performance can be performed *via* situated media architecture systems, inviting new conceptions of music performance and composition, and continuing a long-standing tradition of dialogue and interdisciplinary creativity between music and architecture [15].

Sonic Media Architectures Beyond the PA

Live music is strongly associated with the use of PA systems: a set of loud speakers into which various sound producing or processing equipment can be connected to render music in realtime. Sonic media architecture systems such as ours do not necessarily offer themselves as PA systems through which sound is played, but more as instruments that can be played. As we discuss below, in our case this is because with the wirelessly connected technology we use there is no way to stream multiple channels of audio to each speaker in a low-latency

Figure 1: Performing with the *Mind at Work* system using a laptop running Ableton Live, and a MIDI controller. © The authors.

way that is suited to realtime music making. Instead, the system better suits another paradigm, well known to musicians, of the MIDI (musical instrument digital interface) studio, where it is synthesiser control messages rather than audio that is sent to the rendering devices. In a MIDI studio, the audio output of the rendering devices (synthesisers, drum machines and samplers) are typically mixed back into a stereo mix for the studio loudspeakers of venue PA, whereas in our system, each device has its own speaker, distributed in space.

Historically, although the PA has been dominant for many years, a prominent alternative for musical rendering has been the world of mechanical musical instruments. Most commonly, mechanising musical instruments transforms them into acoustic rendering devices of digital content. Even after the appearance of radios and gramophones, mechanical instruments like player pianos remained popular, in part because the fidelity of recorded or broadcast media was so much more inferior than the sound of a genuine instrument in a space.

More recently the field of robotic music performance has extended how complex networks of acoustic rendering devices can be physically manifest and situated in space in interesting ways that relate questions of music performance and dramaturgy to questions of design. Taken to a theatrical extreme, the works of composer Heiner Goebbels, such as [1], situate mixed media including amplified and mechanical instruments into entirely mechanical theatre works. Robot artist Wade Marynowsky has similarly reconceptualised the operatic form as one that can be explored via interactive musical robots [18]. More generally, many music performers seek to engage with the built environment, to compose performances for specific spaces and to improvise with spaces. An elegantly simple example is the documented use of room acoustics by saxophonist Evan Parker, on records such as *Monoceros*, in which he constructs an improvised performance attuned to a specific, although relatively unremarkable, room [2]. This builds on experimental music's fascination with the creative contribution of space in musical processes of feedback most widely celebrated in Alvin Lucier's conceptual work *I am Sitting in a Room*. Back again in the digital world, the laptop and mobile phone orchestra movement (e.g, [23, 24, 20]) has also embraced the acoustic potential of situating distributed speakers, typically co-locating each speaker with a live musician and their sound generating equipment (laptop or other portable computing device). As with our work, communication between devices here is largely focused on small control messages and the challenge of keeping time between different computer clocks and metronomes.

In this paper we present initial steps in a creative project situated at this intersection. We approach this work both as a creative practice-based research project to explore the aesthetic possibilities of live music performance with distributed sound-and-light sculptural designs, and as a design research project into the creation of a system that simplifies and reduces the effort in creating musical performances on distributed systems.

Design Concepts

In a number of previous studies [4] we have outlined our system design and software architecture in detail. Briefly, our system is built from a set of small raspberry pi computers outfitted with hardware that allows the connection of various display WS2812B Neopixel LED arrays, as well as audio amplifiers that allow the driving of small loudspeakers. Each of these computers has enough computational resources to synthesize audio signals and perform audio sample manipulation. These devices run runtime software that allows them to be connected to a WiFi network and managed from a central control interface, while retaining their autonomy, and we can use this central interface to deploy arbitrary software sketches to each of the devices without requiring re-initialisation of the devices [5]. We have used this system to iterate a number of design concepts [18, 3, 17] related to our central research question: how can we make it quick and easy to create rich distributed audio-visual content for site-specific media architectures. These design concepts are summarised in the following subsections.

Physical flexibility

In our system a creative team can design a physical sculptural arrangement of lights and speakers for a given location in any 3D environment of their choosing. They can then load it into a Unity environment which simulates our runtime system, allowing them to directly program light and sound behaviours in simulation. Work in progress includes simulating the first-person audio experience of walking through such a simulation. Once the real physical system has been built, the same computational model can be loaded onto the computing hardware and will work exactly as in simulation.

From a design research perspective, we have been very interested in the question of when specific design decisions get locked in. For example, a media architecture designer may come up with an initial sketch of a specific physical design of an array of lights and speakers (in the work in Figure 1, the bottom half of a sphere). They must then consider what kinds of media content might play out on that system, as well as how the system will be installed and what it will cost. These interplaying factors might lead to new iterations of the physical design.

In our design research we found that it was quite common for a design decision (such as the spatial design of lights and speakers) to be locked in, only to be revisited at a later date given new knowledge or thinking. We subsequently termed such decisions 'fuzzy decisions', and have attended to thinking about how we could support this iterative process of settling aspects of a design in light of how fuzzy decisions arise and are resolved: either by helping better anticipate factors that would influence a design, or simply enabling design factors to remain flexible for longer.

Technical constraints

Our approach is based on the idea of managing distributed sound by having multiple low-cost computers processing audio, and communicating with those computers over WiFi [7, 10]. This presents two key technical constraints: the processing power of the individual computers, and the network performance of the overall system (bandwidth, latency and stability). The network constraint means that our systems cannot behave like multi-channel audio systems, where you can stream sounds directly to speakers from a central computer.

In light of this, we have found it preferable to conceptualise the system as a "distributed synthesiser", that receives global control data (notes and control parameters) from a central computer, and render sound over the array of speakers. Given this concept, we then treat the light components as just another output medium of the synthesiser. The latency of the network (around 100ms) is low enough to support basic realtime control from a central computer, but not as low as is typically desired in digital musical instruments.

The processor constraint limits the number of audio DSP processes that any one device can run. Thus within the conceptualisation of the system as a distributed synthesiser, each speaker is considered to be monophonic, and polyphony is achieved by distributing sounds across the array of speakers. For example, if you wish to play a chord on the system, a polyphony management process will distribute the notes in that chord across speakers.

Software flexibility

Our software approach is based on the idea of a runtime system that runs on each device, to which we can throw live-coded sketches in realtime over WiFi. In our latest version of the runtime, we have developed an approach of virtualising outputs (individual speakers or lights), so that we focus the programmer's attention on writing code that describes how one individual output behaves. That behaviour can be easily deployed to the distributed array of computers, which automatically works out which device controls which output.

Much of our design work then focuses on how to conceptually separate out the added task of thinking spatially about the system as a whole: how sounds and lights move around and how audio and light parameters vary across space. Since the user has direct access to a programming API for programming the devices, they have the freedom to hack any part of this design. However, our design research has strongly pointed us to the need to impose some default framing concepts as a way of narrowing the space of possibilities whilst still allowing a vast array of creative options. Conceptually, we frame this in terms of the concept of 'space enablers': specific configurations of creative tools that place a rapidly accessible array of creative opportunities in front of a creative practitioner. Space enablers are designs that hit the sweet spot between the power of creation and the ease of creation. Our standardised architecture arises from integrating knowledge about the technical constraints of our system and the basic requirements of thinking spatially with a distributed computing system.

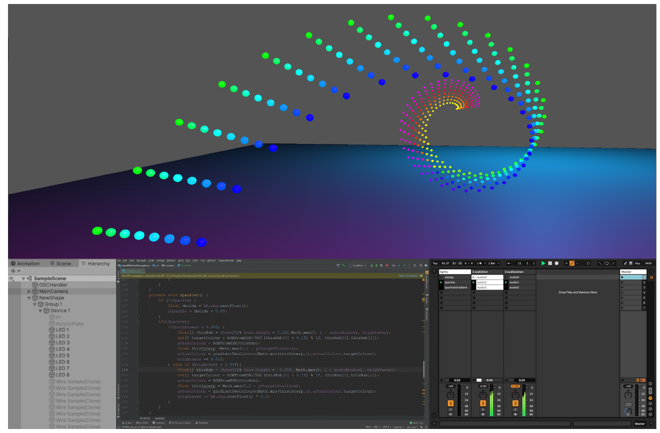

Figure 2: (Top) A screen grab of the composition environment, comprising a virtual model in Unity, a coding interface in IntelliJ, and an Ableton Live project for live and timeline-based control of parameters. (Bottom) a more detailed view of a light configuration in simulation, with a light gradient mapped across the 2D surface. © The authors.

Figure 2 shows a workspace in which a physical arrangement is modelled in Unity, the code describing the distributed synthesiser behaviour is being programmed in IntelliJ, and controller data is being composed in Ableton. When deployed in the final system, the Ableton project stays open to control the system. The code is deployed to the runtime system running on multiple distributed Raspberry Pis. The spatial model is also loaded into the Raspberry Pis so that each device is aware of the physical positions of the lights and speakers it is controlling.

Jams in the Hueosphere

Jams in the Hueosphere is a series of live performance events held in Sydney, funded from a local government Covid relief fund to stimulate live music after the impacts of lockdown. The project was pitched to create a novel experience for audiences and also give a series of local artists the chance to experiment and develop their repertoire to new forms of digital media performance. Artists were invited to give live performances and to develop their own creative uses of the *Jams in the Hueosphere* technology. A challenge for the development of *Jams in the Hueosphere* was the limited time that would be available to work with the artists. Not only would the artists have limited time to experiment and develop their performances using the system hardware (due to other constraints such as access to the performance space), but given that this was a new software that was not fully user-tested, the

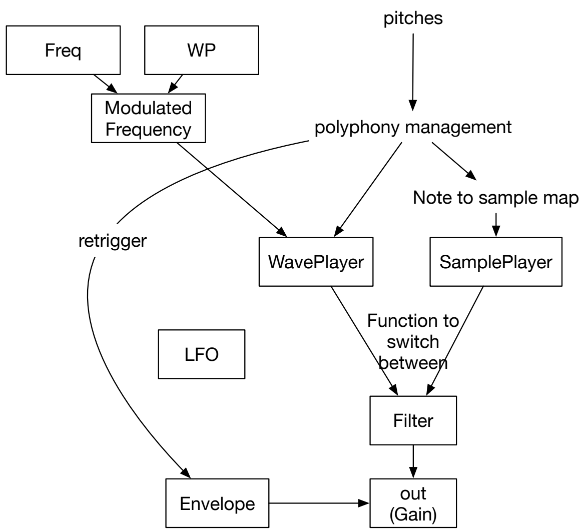

Figure 3: Our standard synthesiser architecture and Ableton Live controller device. © The authors.

potential to iterate the system in response to user feedback was limited. This is the context in which a general-purpose synthesiser architecture was developed.

System

Our development framework, *Happybrackets*, allows a program to be sent to a number of remote devices which then respond to incoming network messages from a controller computer. An audio API allows the creation of custom synthesiser signal chains, and equivalent processes for controlling lights. Typically we send the same synthesiser configuration to all devices and then work with global messages that control the devices *as one*. We then use the devices' physical positions as additional variables that can be used to map the global synthesiser parameters to create spatial effects. For example, we can easily create gradients of hue, frequency or filter strength, running across the space of the devices.

For the *Jams in the Hueosphere* project, we created a single synthesiser with a large number of variables and possible configurations (Figure 3). We then created a MIDI instrument device for Ableton Live which sends control messages to the network of devices. The result, from the musician's perspective, is something that looks and feels like a regular softsynth, except that instead of playing sound on the controller computer, sound and light is rendered on the system. Our Unity simulator can be used to view the design of the lights on the controller computer. The audio synthesiser includes functions for FM synthesis, sample playback, drum kits, wavetable synthesis, ADSR envelopes and LFO modulation: a standard suite of options for a synthesiser. The main difference with a real softsynth is how polyphony is managed: due to the computational limits of the Raspberry Pi, each speaker in the system is limited to monophonic playback. Polyphony is achieved by distributing the various notes in a chord across the system.

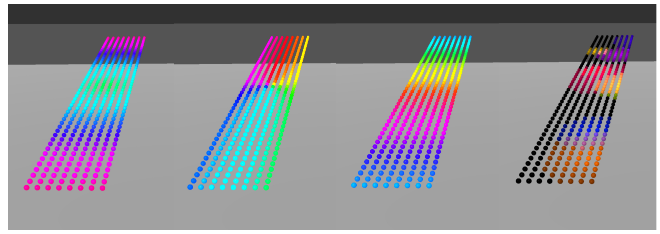

Figure 4: Spatial mapping strategies controlling colour hue. From left to right: variation radiating linearly from centre; variation mapped to angle around centre point; variation attached to discrete bands arranged linearly; specific zones of intensity. © The authors.

When a chord is played, each device chooses one note from the chord to play, according to a specified policy: either random, or based on the spatial position of the speakers (e.g., the lowest note in the chord would play at one end of the system, and the other notes would be located at different intervals along one dimension of the speaker array).

In addition, the user can select and modify a range of spatial strategies, including spatial envelopes, that dictates where on the system sounds and lights should be active, and spatial mappings, that dictate how some parameter varies across the space (e.g., hue or filter frequency). Spatial envelopes act like the envelopes used in sound synthesis, but are spatial rather than temporal. When the envelope is "fully open", all of the sounds and lights are active. The user can then choose to narrow the active zone, specifying a centre point and size parameters. In this way they can make sounds and lights move around the space. In reality, the sounds and lights are not moving, but simply being turned up and down depending on whether they are inside or outside of the active zone. A range of spatial mapping strategies are presented in Figure 4. Ongoing work includes refining this suite of functions to create an intuitive that can easily be manipulated to create spatial effects.

Creative Workflow

In our performance set up, a musical artist or band plays on stage with a standard PA system, but additionally, our system of lights and speakers is installed to flow up from the stage and over the heads of the audience. The musicians play as normal, using Ableton Live and any other instrumentation they like. Our synthesiser can be added into their existing Ableton Live rig so that as well as playing sounds out of the PA, they can control the distributed audio visual synthesiser. The parameters available in the interface of our Ableton Live device – including all of the synth settings, light mappings, the spatial control parameters and polyphony policy – work like any other Ableton Live parameters: the musician can pre-program control of these parameters into Ableton clips which can be triggered via the Ableton session interface, or they can connect any MIDI controller in and assign which dials and sliders will control which parameters (Figure 5 - See

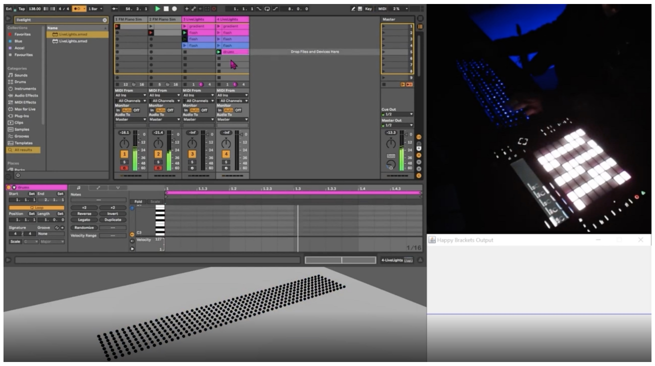

Figure 5: Our system incorporates commercial and well-known DAW software and hardware interfaces with simulation techniques and custom software that links these existing music paradigms to media architectures. © The authors.

our videos for details [1] [2]).

As an example, if the musician wanted to have a synth pad emanate from the stage in waves, they could prepare a looping clip that updates the parameter that controls the position of the active zone to fly from one end of the space to another. They could then manually play chords on a MIDI keyboard. Running the clip would cause these chords to fly out from the stage. Alternatively, they could assign an XY joystick to control this position variable, and manually move the sounds around the space.

We have found this setup makes effective use of the rapid workflow available to musicians in Ableton. Rather than programming more functionality into our system, resulting in a more complex interface and codebase to manage, we leverage as best as possible the creative freedom that the Ableton interface allows. For example, it is easy in Ableton to map an LFO to any synth parameter, making it very easy to rapidly mix together different processes to create generative effects. Thus, for example, making moving waves of hue intersecting with moving waves of timbral variation becomes easy.

In addition, our code can easily be hacked, so that the synthesiser design can be updated and redeployed to the network of devices. For example, if the musician wanted the array of synths to have microtonal variation of their frequencies, this is not a feature available in our standard synth, but could easily be added by adding a random variable at the right point in the expression calculating the frequency.

Aesthetic Considerations: How do we play buildings as audio visual synthesisers?

Although we have designed the *Jams in the Hueosphere* system to be as seamlessly integrated as possible into a musician's workflow, sonically the system has very different properties to a regular PA, the speakers are small and have poor frequency response below 200Hz, and each one is powered by just a 3W amplifier. They are housed in custom 3D printed spherical enclosures that amplify well. This should also not be thought of as similar to existing wavefield synthesis systems as the audio signals to each of the speakers are not tightly synchronised to create specific audio effects. The sound field of the speaker array creates a special acoustic canopy that is distinct in the way that the distributed array creates multiple interfering sources. Rich pads, noisescapes and clusters of sound are rendered with a spatial richness that is distinct from a standard PA. Reverb and delay-type effects can be created through the movement and slight temporal delays of sounds. Careful sound design can create playful effects where the sound leaps from the stage. The experience can also be more interactive for audiences, who by moving around the space, can seek different 'stances' on the soundscape.

With the resistance of the stereo PA convention to disruption from new technological possibilities, we view this field of experimentation as an aesthetic niche first and foremost, rather than a contender to challenge how live music is rendered in the mainstream. However, this paper has shown how the technological affordances, and associated constraints, of rendering sound through massive networks of remote computational devices, both invite and demand new ways of thinking about digital music and its situatedness in space. Beyond our own experimentation, the system is awaiting its first deployment in a public space, having been deferred due to pandemic response restrictions in NSW, Australia. Our initial thoughts on the effectiveness of our designs and the creative possibilities that will be most effective, are being tested soon with a cohort of live musicians who have been commissioned to create new works with the system. Our philosophy of curating such performances is grounded in an ongoing examination of the relation between sound, space and musical aesthetic. We will explore how our system works with different sized PAs, musical minimalism, free improvisation, acoustic instrumentation and different configurations of our system throughout original spaces.

Conclusion

We have presented an approach to expanding live music performance to new forms of sonic media architectures. The technological solution space we have explored, driven by cost, flexibility and scalability, actually pushes us from thinking about such systems as distributed audio visual synthesises rather than PAs and displays. When embedded in spaces, these systems in turn evoke the idea that the space itself becomes a playable audio visual synthesiser.

Acknowledgments

This paper is based on research supported by ARC Linkage Project LP180100151 with partners BitScope Designs, Squidsoup, the Casula Powerhouse Art Centre, 107 Projects, Create NSW, City of Sydney Council, and ArtworksRActive. We thank all the partners for their input into the production of these creative works.

References

[1] Birringer, J. 2013. Choreographic Objects: Stifters Dinge. *Body, Space & Technology* 11(2):1–16.

[1] https://cloudstor.aarnet.edu.au/plus/s/ySXyKl9veb1Nj2h

[2] https://cloudstor.aarnet.edu.au/plus/s/2sqoK689DpwL9Rx

[2] Blancarte, T. 2020. Charon as Muse - The Ferrying of Voices in Evan Parker Solo Saxophone Music to the Double-Bass as Creative Authorship. *Rhythmic Music Conservatory Copenhagen*.

[3] Bown, O., and Ferguson, S. 2016. A Musical Game of Bowls Using the DIADs. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, 371–372.

[4] Bown, O., and Ferguson, S. 2018. Understanding media multiplicities. *Entertainment Computing* 25(March 2018):62–70.

[5] Bown, O.; Loke, L.; Ferguson, S.; and Reinhardt, D. 2015. Distributed Interactive Audio Devices: Creative strategies and audience responses to novel musical interaction scenarios. In *Proceedings of the 21st International Symposium on Electronic Art*.

[6] Caldwell, G. A., and Foth, M. 2014. DIY media architecture. In *Proceedings of the 2nd Media Architecture Biennale Conference on World Cities - MAB '14*, volume 19-22-Nove, 1–10. New York, New York, USA: ACM Press.

[7] Ferguson, S.; Rowe, A.; Bown, O.; Birtles, L.; and Bennewith, C. 2017a. Networked pixels: Strategies for building visual and auditory images with distributed independent devices. *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition* 299–308.

[8] Ferguson, S.; Rowe, A.; Bown, O.; Birtles, L.; and Bennewith, C. 2017b. Sound Design for a System of 1000 Distributed Independent Audio-Visual Devices. *Proceedings of the International Conference on New Interfaces for Musical Expression* 245–250.

[9] Foth, M., and Caldwell, G. A. 2018. More-than-human media architecture. In *Proceedings of the 4th Media Architecture Biennale Conference*, MAB18, 66–75. New York, NY, USA: Association for Computing Machinery.

[10] Fraietta, A.; Bown, O.; Ferguson, S.; Gillespie, S.; and Bray, L. 2019. Rapid composition for networked devices: HappyBrackets. *Computer Music Journal* 43(2-3):89–108.

[11] Haeusler, M. H.; Tscherteu, G.; and Tomitsch, M. 2012. *New media facades: A global survey.* Ludwigsburg, Germany: AV Edition.

[12] Haeusler, M. 2009. *Media Facades-History, Technology, Content*. Ludwigsburg, Germany: AV Edition.

[13] Hallam, J.; Zheng, C.; Posner, N.; Ericson, H.; Swarts, M.; and Do, E. Y. L. 2017. The Light orchard: An immersive display platform for collaborative tangible interaction. In *UbiComp/ISWC 2017 - Adjunct Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*, 245–248. ACM.

[14] Hauesler, H.; Barker, T.; and Beilharz, K. 2010. Interactive Polymedia Pixel and Protocol for Collaborative Creative Content Generation on Urban Digital Media Displays. In *Proceedings of the International Conference on New Media and Interactivity*, 1–7. Marmara University.

[15] Hornecker, E., and Stifter, M. 2006. Learning from interactive museum installations about interaction design for public settings.

[16] Kuznetsov, S.; Paulos, E.; and Gross, M. D. 2010. WallBots: Interactive wall-crawling robots in the hands of public artists and political activists. In *DIS 2010 - Proceedings of the 8th ACM Conference on Designing Interactive Systems*.

[17] Loke, L.; Bown, O.; Ferguson, S.; Bray, L.; Fraietta, A.; and Packham, K. 2018. Your move sounds so predictable! In *Proceedings of the 2018 Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, CHI PLAY '18 Extended Abstracts, 121–125. New York, NY, USA: Association for Computing Machinery.

[18] Marynowsky, W.; Ferguson, S.; Fraietta, A.; and Bown, O. 2020. 'The Ghosts of Roller Disco', a Choreographed, Interactive Performance for Robotic Roller Skates. In *Proceedings of the Fourteenth International Conference on Tangible, Embedded, and Embodied Interaction*, 631–637.

[19] Mauzey, P. 1962. Control Console for Multi-Speaker Presentations of Electronic Music. *Journal of the Audio Engineering Society* 10(4):338–348.

[20] Ogborn, D. 2014. Live coding in a scalable, participatory laptop orchestra. *Computer Music Journal* 38(1):17–30.

[21] Sato, M. 2008. Particle Display System: A Real World Display with Physically Distributable Pixels. In *Proceeding of the 26th Annual CHI Conference on Human Factors in Computing Systems - CHI '08*, 3771–3776.

[22] Seitinger, S.; Perry, D. S.; and Mitchell, W. J. 2009. Urban Pixels: Painting the City with Light. In *Proceeding of the 27th Annual CHI Conference on Human Factors in Computing Systems - CHI '09*, 839–848.

[23] Trueman, D.; Cook, P.; Smallwood, S.; and Wang, G. 2006. PLOrk : The Princeton Laptop Orchestra, Anatomy of a PLOrk Meta-Instrument. In *Proceedings of the International Computer Music Conference*, 443–50.

[24] Wang, G.; Bryan, N.; Oh, J.; and Hamilton, R. 2009. Stanford laptop orchestra (SLOrk). In *Proceedings of the 2009 International Computer Music Conference, ICMC 2009*.

Author Biographies

Oliver Bown is an Associate Professor at the School of Art & Design, UNSW Sydney, and co-director of their Interactive Media Lab. He researches creative technology practice and is an electronic music maker. He is the author of "Beyond the Creative Species: Making machines that make art and music" (MIT Press, 2021).

Kurt Mikolajczyk is a musician, composer, and creative coder currently undertaking a Ph.D. in interaction design at the University of Technology Sydney. In 2019 he completed a Masters of Music at the Sydney Conservatorium of Music, developing software tools for composing polytemporal music and a portfolio of works for jazz ensemble and laptop.

Sam Ferguson is a Senior Lecturer within the School of Computer Science and co-director of the Creativity and Cognition Studios at the University of Technology Sydney. He has a background in music performance, cognitive science, and psycho-acoustics and acoustics. He focuses on sound and music and their relationship with creativity and human experience, in contexts such as installation art, creative coding, and machine learning, as well as focusing on cognitive science. He has more than 80 publications in areas as diverse as spatial hearing and loudness research, data sonification, emotion, and tabletop computing.

Benedict Carey's work branches aspects of computer science, musicology and cognitive science, focusing on analysis and creation of unique forms of music notation. Currently he teaches interactive media, rapid prototyping, music and sound production at the University of Technology Sydney and University of Sydney, is a research assistant in the Interactive Media Lab at the University of New South Wales, and is completing a doctorate in musicology at the University of Music and Drama Hamburg.